

Article

The Effect of Appearance of Virtual Agents in Human-Agent Negotiation

Berkay Türkgeldi ^{1,2}, Cana Su Özden ^{2,*} and Reyhan Aydoğan ^{2,3,*}

¹ Vestel Elektronik Sanayi ve Ticaret A.Ş. Manisa 45030, Türkiye

² Department of Computer Science, Faculty of Engineering, Özyeğin University, İstanbul 34794, Türkiye

³ Interactive Intelligence Group, Department of Intelligent Systems, Faculty of Electrical Engineering, Mathematics, and Computer Science (EEMCS), Delft University of Technology, 2628 CD Delft, The Netherlands

* Correspondence: su.ozden@ozu.edu.tr (C.S.Ö.); reyhan.aydogan@ozyegin.edu.tr (R.A.)

Abstract: Artificial Intelligence (AI) changed our world in various ways. People start to interact with a variety of intelligent systems frequently. As the interaction between human and AI systems increases day by day, the factors influencing their communication have become more and more important, especially in the field of human-agent negotiation. In this study, our aim is to investigate the effect of knowing your negotiation partner (i.e., opponent) with limited knowledge, particularly the effect of familiarity with the opponent during human-agent negotiation so that we can design more effective negotiation systems. As far as we are aware, this is the first study investigating this research question in human-agent negotiation settings. Accordingly, we present a human-agent negotiation framework and conduct a user experiment in which participants negotiate with an avatar whose appearance and voice are a replica of a celebrity of their choice and with an avatar whose appearance and voice are not familiar. The results of the within-subject design experiment show that human participants tend to be more collaborative when their opponent is a celebrity avatar towards whom they have a positive feeling rather than a non-celebrity avatar.

Keywords: celebrity vs. non-celebrity; experimental study; familiarity to opponent; human-agent negotiation; interactive intelligent systems; virtual agents



Citation: Türkgeldi, B.; Özden, C.S.; Aydoğan, R. The Effect of Appearance of Virtual Agents in Human-Agent Negotiation. *AI* **2022**, *3*, 683–701.
<https://doi.org/10.3390/ai3030039>

Academic Editors: José Manuel Ferreira Machado and Kenji Suzuki

Received: 11 May 2022

Accepted: 12 August 2022

Published: 16 August 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

With the ongoing development of Artificial Intelligence, the use of intelligent agents has become more and more prevalent in different parts of our lives. In some cases, those agents may need to communicate and collaborate with each other to achieve a common goal. In the case of any conflict between them, negotiation—the process of resolving conflicts—may take place to come to an agreement [1]. This kind of interaction may be necessary with their human counterparts. Compared to humans, software agents can be far more efficient at optimizing negotiation bids [2]. Combined with recent advances in human-computer interaction, these agents can have the ability to negotiate with a human via a natural dialogue. One can foresee a future where humans negotiate with AI agents or along with AI agents [3]. Therefore, it is vital to design and develop a human-agent negotiation framework that involves different challenges [4].

Designing such interactive systems may require considering human factors and benefiting from theories introduced by different fields such as psychology, economics, behavioral sciences, and cognitive science. Especially understanding psychological factors affecting human negotiators' behaviors and attitudes plays a key role in developing human-agent negotiation systems. The best course of action to build such systems would be studying human-human negotiations and deriving the critical factors influencing their decisions [5–7]. Another way would be to develop a virtual human-agent negotiation framework that can allow researchers to study the effect of the chosen factor. While

some works investigate the effect of facial expressions and emotions on negotiation outcomes [8,9], some focus on the effect of gender [10–12]. Accordingly, this work pursues studying the effect of knowing their opponents with limited knowledge (e.g., familiarity to their opponent's appearance and voice) on the negotiator's attitude and consequently on negotiation outcome.

As far as studies in the field of psychology are concerned, it has been shown that familiarity plays a vital role in negotiation outcomes in human-human negotiations [5,7]. For instance, Druckman and Broome show that a decrease in familiarity with the opponent causes a decrease in the human negotiator's willingness to act collaboratively and consequently reach an agreement. Moreland et al. investigate the effect of visual familiarity on perceived similarity [7]. According to the experiments conducted in that study, participants believed that they shared similar preferences and values with the people with whom they were more visually familiar than those they saw for the first time. Those works encourage us to explore whether the effect of familiarity in human-human negotiations could be observed in the human-agent negotiation in which agents present themselves as humanoid avatars. It is worth noting that familiarity can be investigated from several perspectives. For example, it can be interpreted as knowing someone with limited knowledge (e.g., through media) or knowing someone closely (e.g., knowing their personalities and attitudes, interacting with them on a daily basis, etc.). In this work, we study familiarity from the former perspective. In other words, this study empirically investigates the effect of limited knowledge about negotiation partners. Therefore, the aim of this research project is to design effective artificial negotiation systems that enable us to conduct experiments involving human negotiation with an avatar that is either familiar, or not familiar.

To sum up, this work introduces a human-agent negotiation framework where a human negotiator can negotiate with a virtual avatar so that researchers can investigate the effect of the human negotiator's limited familiarity with their opponent on negotiation outcome and process. In other words, the main aim of this work is to study the effect of familiarity with their opponent (i.e., celebrity versus non-celebrity unknown avatar) on their negotiation attitude and the outcome. A user experiment was conducted where participants negotiated with an avatar that looks like a chosen celebrity and also negotiated with a non-celebrity avatar. Results show that participants tend to be more collaborative when avatars represent celebrities for whom subjects have positive feelings, rather than unfamiliar avatars, as evidenced in their tendency to reach agreement at lower personal utility than when negotiating with unfamiliar avatars for whom they report no or negative feelings.

The remainder of this article is organized as follows. Section 2 presents the related work. Section 3 explains our human-agent negotiation framework, while a novel negotiation strategy employed by our agent during the user experiments is introduced in Section 4. Afterwards, Section 5 briefly provides our research methodology and road map. Section 6 explains the user experiments that we conducted to study the effect of human negotiators' familiarity with their opponent and analyzes our findings. Finally, Section 7 concludes our work with future research directions.

2. Related Work

Automated agent negotiation has been the focus of attention for several decades, and a variety of research studies have been conducted in this field [1,13–20]. Moreover, the research community has been organizing an international competition in the field of agent-based negotiation to facilitate research and provide benchmarks for the community [21]. The community has primarily focused on agent-agent negotiations where one agent outperforms the other in one or many ways such as monetary gain, user interactions, etc. Recently, new leagues in this competition regarding human-agent negotiation have been introduced [22], which indicates that human-agent negotiation has become more and more attractive for the community in terms of the challenges required to build agents interacting with human negotiators. There are several challenges in human-agent negotiation [4,23]. That is, the designers of the agents should concern themselves with

the human factor [2,24,25]. In a limited number of interactions, agents should be able to reach a consensus with their human counterparts. One of the leading research questions is how a negotiating agent effectively interacts with human negotiators, which is the focus of this work. In the literature, human-agent interaction aspects of negotiating agents are being explored currently in many ways: finding a way to outperform people by considering cultural differences of the opponents in negotiation [26]; the use of facial expressions and emotions to explore the effect on outcomes [9,27,28]; exploring the effect of argument usage in negotiation [29]; including the use of multiple modalities to explore human-agent negotiation dynamics [2,8,30]; and using agents to train people for future negotiations [31–33].

In several of these use cases involving human-agent negotiation interactions, agents are embodied in various forms that fit best. Some include humanoid animated avatars that can see and hear the human participant [2,30], while some agents are accessed through a text-based chat window [34–36]. Divekar et al. use an immersive room where agents appear as human-scale animated avatars in an extended reality environment [30]. Their goal is to make users feel like they are elsewhere, negotiating with street market shopkeepers in a foreign country. An important aspect of their work is making the agents believable in the sense that users engage with them as if they are believed to be human beings. Especially in such settings, the effect of familiarity could be crucial to the human interlocutors, which is the focus of our work.

As research gears towards representations of agents where agents look and communicate like humans, we ask the question of whether such a representation will also mean that some characteristics of human-human negotiations might be observed in human-agent negotiations where the agent is humanoid. Lin et al. have stated that such familiarity plays an important role in negotiation outcomes [37]. Accordingly, we focus on a specific aspect of human-agent negotiations, specifically the effect of knowing your partner in terms of familiarity to their opponent on negotiation.

Yuasa et al. study the facial expression and history effect on the decision-making process in a negotiation game, a variation of the Prisoner's Dilemma [9]. In their study, participants are asked to negotiate a couple of times with software agents that can show different facial expressions (i.e., bow, avert, happy, angry, and cool). In some cases, they negotiate with the same opponent to study the effect of history. Their experimental results show that happy agents make participants act more collaboratively. Furthermore, the negotiation result of previous sessions with the same appearance agents affects the participant's current decisions since familiarity with the current opponent builds up a certain level of trust. For example, participants tend to cooperate more if their opponent acted collaboratively in the past. Although we investigated the effect of familiarity from a different perspective (i.e., negotiating celebrity versus negotiating with the same person repeatedly), our results support each other as mentioned in Section 6.

De Melo et al. study the effect of an agent's emotion particularly anger and happiness on negotiation [27]. Participants are asked to negotiate with virtual agents that adopt a different facial expression. The results show that, when participants negotiate with the virtual agent that expresses anger, they concede more than the case of negotiating with a neutral or happy opponent. During our experiment, we observed that participants took their opponent's facial expression into account while generating their bids. Therefore, it is very important to use the right facial expression while designing a virtual negotiating agent. Moreover, Mell et al. study how human counterparts can be affected by competitive or collaborative agent opponents [38]. Their results show that competitive strategies made the human participants concede more. On the other hand, a study advocates that people are more willing to renegotiate with warm agents, although there is no significant difference found on negotiation outcome in their experiments [38]. Another study that explores the effect of aggressive attitudes on human-agent negotiations shows that aggressive attitudes in virtual environments affect participants' emotional states, similar to the real environment [39]. The degree of this impact is lower in the virtual environment than in the

real one. Lin et al. measure the effect of gender in negotiation [10]. Their results show that negotiators change their strategy according to the opponent's gender.

On the other hand, the familiarity effect has been studied in different contexts [40,41]. Wauck et al. developed a search and rescue game in a virtual environment where the main characters look like participants or other avatar appearances [41]. Their experimental results show there is no significant effect of using self-similar characters on the performance in designed games. Another study investigates the familiarity effect on negotiation in a task-based game environment [40]. For this purpose, participants are divided into two groups: participants in the former group are not allowed to communicate with each other before the game begins, whereas participants in the latter group are allowed to communicate with each other. During the game, the participants do not see each other but they know with whom they are playing. In this virtual environment, each player is demonstrated by a robot-like avatar. The results show that the participants feel more comfortable working with a partner whom they had met face-to-face first, although there is no significant performance difference between groups. No significant difference in their performance may stem from the simplicity of the rules in the game. In line with their results, we also observed that participants collaborated more when they negotiate with the celebrity avatar that they like in our experiments.

Since there is an undeniable effect of humans' perception and their psychological tendencies on human-human relations, it is worthwhile to take a look at the studies in the field of psychology. Our findings are supported by some studies that conduct experiments on the familiarity effect in the field of psychology. Moreland and Zajonc measure the influence of the mere exposure effect on the perceived similarity [7]. They divided participants into two groups where photos of different people were shown to the first group while the photos of the same people were shown to the second group each week. Afterward, they studied participants' attitudes toward the people they saw in the photos and their beliefs about which degree these people are similar to them and share their values. They found a positive correlation between attraction, perceived similarity, and familiarity. The results show that the mere exposure to the same individual's photo each week increased their familiarity with this stranger. Hence, it increased attraction toward this person and participants' perceived similarity between their values and this individual's values. They considered this person's preferences more and more similar to their preferences. Furthermore, Reder and Ritter claim that the participants respond faster and exert less effort when they are familiar with the given problem [42]. In another study, Druckman and Broome aim to understand the effect of liking and familiarity on the negotiation outcome [5]. They design an experiment where the participants play a representative of a culture and negotiate with a representative of a totally different culture to reach an agreement on a variety of issues. They investigated three conditions: (1) high familiarity and high liking; (2) low familiarity, high liking, and (3) high familiarity, low liking. The results showed that decrease in either liking or familiarity was followed by a decrease in the participants' willingness in reaching mutual agreement and being collaborative. In contrast, we focused on whether having a positive feeling for your opponent affects your negotiation attitude (e.g., ending up an agreement with a lower utility) during the negotiation.

Furthermore, Stuhlmacher and Champagne analyze the effect of time pressure and being given additional information on the negotiation process and outcome [43]. In their setting, participants are asked to play a job applicant role and negotiate with a computer program to agree on working conditions such as salary, medical coverage, and so on. They conduct between-group designed experiments. To investigate the effect of time pressure, they introduce two conditions: high (i.e., the deadline is 45 min) and low time pressure (i.e., the deadline is 15 min). For the additional information, they define three conditions: (1) only payoff table for users is shown, (2) detailed payoff information indicating the importance of each issue and evaluation scores for each value shown, and (3) detailed payoff information for both sides (i.e., user's and computer agent's payoff function) is

shown. Regarding the user interface, our preference graphs and tables are very similar to the pie-chart and tables used in that study. Their results showed that additional information influences the negotiation outcome and process while time pressure did not significantly affect participants' scores. Interestingly, when the utility of the opponent's utility was shown, somehow participants acted collaboratively and conceded more to find an agreement. While their focus is on studying the effect of time pressure and additional information, ours is on the effect of familiarity to the opponent (i.e., celebrity or non-celebrity). In this aspect, they are complementary works. Sheffield studies the impact of communication way on an individual's negotiation performance [44]. Participants were asked to act as buyers/sellers and negotiate with each other. Three factors with two conditions (e.g., verbal communication, visual communication, and negotiation orientation) are tested in their experiments. They found out that the total amount of communication and joint profit was higher in the verbal mode than in the text-based mode. In addition, visual communication increased cooperativeness and consequently joint profit when the participants acted cooperatively. In our case, in a virtual environment, the participants can communicate their offers through text-based communication as well as being able to see their opponent visually while the computer agent can communicate both verbally and text-based. However, the focus of our study is not the effect of communication medium.

3. Virtual Negotiation Avatar Framework

To examine human-agent negotiations experimentally, we design and develop a Web-based negotiation framework where a virtual avatar agent negotiates bilaterally with a human counterpart. The framework allows us to design and integrate negotiation strategies and to change the avatar character interacting with the human. The framework adopts Alternating Offers Protocol, where negotiating parties make offers iteratively in a turn-taking fashion until reaching a termination condition (i.e., reaching a predefined deadline or reaching an agreement) [45]. In our framework, a human negotiator initiates the negotiation by making an offer, and the framework counts down the timer after sending her/his first bid. There are five main components: virtual avatar character, human offer generator, conversation history, time remainder, and preference profile chart.

A virtual embodied avatar character interacts with the user through his/her facial expression as well as speech. Facial expressions used in the experiment are shown in Figure 1. During the negotiation, this character can change his/her facial expression in line with the designed strategy. In addition to the textual offer information, this character provides the content of its offers via speech. Since our research focuses on analyzing the effect of knowing the opponent, particularly the effect of the familiarity with the opponent in negotiation, this component plays a crucial role in our setting. The virtual avatar is located in the center, and its size is bigger than the other user interface components so that participants can focus on the avatar's facial expression directly. Figure 2 shows the avatar that we used in the training session. It is worth noting that celebrities (e.g., some Hollywood celebrities, singers, and famous businesspeople from all over the world) and real human faces are used during real negotiation sessions. We use a software toolkit to give the real human effects such as periodically blinking effects, minor head movements, speaking effects, and emotional facial expressions (e.g., sadness or happiness) [46]. When participants negotiate with the chosen celebrity, the avatar's voice should align with that celebrity's own voice. To achieve this, we use a Web service that generates a celebrity's voice from predefined sentences [47].

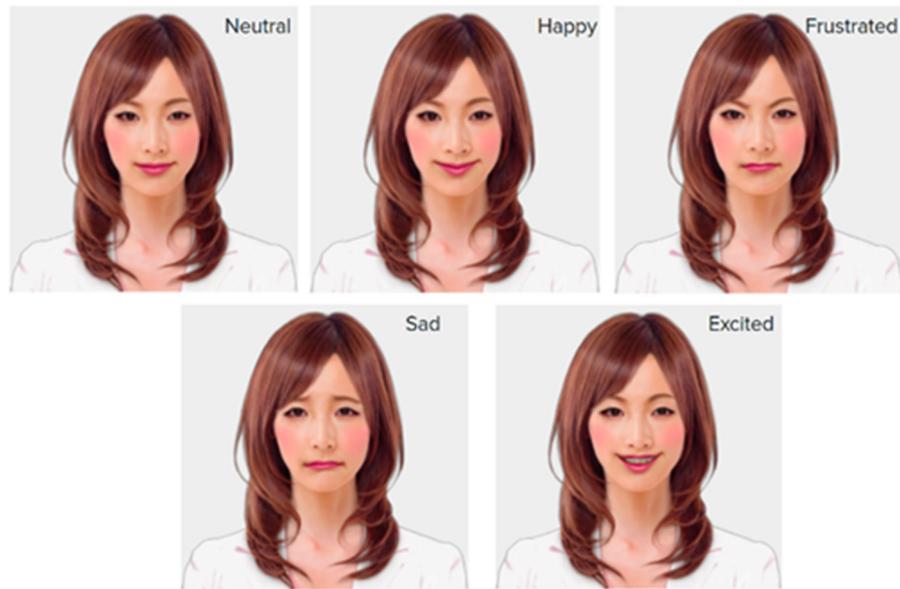


Figure 1. Facial Expressions [46].

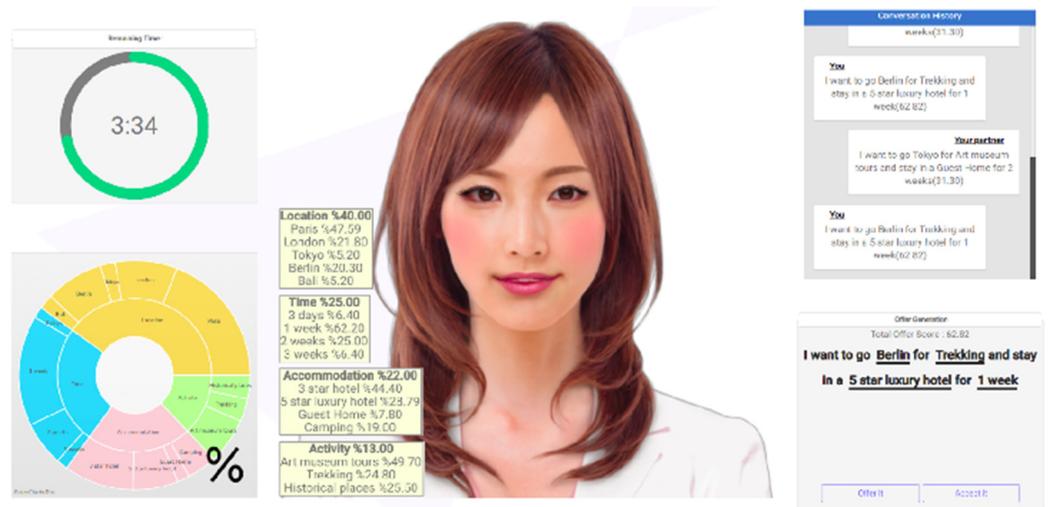


Figure 2. Our Human-Agent Negotiation Interface.

In our framework, human negotiators specify their offers through a textual interface. This interface provides a predefined offer sentence form that allows users to alter the content of the offers by using the drop-down in a user-friendly way. Suppose that the human negotiator wants to modify the content of the offer. In that case, s/he can click the related issue (e.g., issues in our holiday domain like location, accommodation, etc.) value which is underlined in the interface. It is worth noting that when the human negotiator enters a complete offer, a score box appears above the offer sentence and shows the utility score of that offer for the human negotiator. A utility score of an offer can take a real number between 0 and 100. In our framework, negotiating parties negotiate over multiple issues, particularly on their future holiday similar to the work [6]. According to the negotiation scenario, there are $I = \{1, 2, \dots, n\}$ negotiation issues whose possible values are defined in D . The set of all possible offers according to the negotiation domain is represented by O , and an offer from this set is represented by o . The representation of the agent’s preferences is given by using the additive utility function shown in Equation (1).

$$U(o) = \sum_{i \in I} w_i \times V_i(o_i) \tag{1}$$

Regarding the equation, w_i represents the importance of the negotiation issue i for the agent, o_i stands for the value for issue i in offer o , and $V_i(\cdot)$ represents the valuation function for issue i which denotes to what extent that issue value is preferred by the agent. The sum of weight values adds up to 1 represented by $\sum_{i \in I} w_i = 1$ and the domain of $V_i(\cdot)$ is between 0 and 100. for any i . Note that each negotiating party is aware of their own preferences and does not have access to the other side's preferences. On the one hand, this utility score bar enables human negotiators to see their score without making any calculations. That is, human participants can act according to their given roles during the experiments. On the other hand, being aware of the received score while making offers may influence their negotiation attitude. In order to send his/her offer to the avatar, the human negotiator should click the offer button while s/he should click the accept button to accept the given counteroffer by the avatar.

The framework provides a conversation history showing the exchanged offers with the utility scores for the human participants to enable participants to track their bidding history. This component is designed to give the impression of chat-based communication. As mentioned above, the negotiating parties have a deadline in terms of minutes (e.g., 5 or 10 min) to reach an agreement. For a fair negotiation, the framework provides a time reminder denoting the remaining minutes and seconds. The indicator is in green color at the beginning of the negotiation while it becomes yellow after a while and finally turns red when it reaches a critical time point (e.g., last minute). If time is up and no agreement occurs, negotiation ends with a failure. In this case, both sides will take a utility of zero points.

In experiments, human participants are given a particular preference profile and are asked to act according to the given preferences. Before their negotiation, they have enough time to study their profiles. However, due to the nature of humans, they may forget some details regarding these preferences, which may influence their negotiation significantly. Therefore, for our experiments, a user-friendly multi-level pie chart graph is created to show the preference profile visually, which is modeled in terms of an additive utility function (see lower left part of Figure 2). The inner part of this graph represents the importance levels of negotiation issues, while the outer part denotes preferences on issue values accordingly. The larger the segment associated with the issue becomes, the more important the issue is. For example, if Paris is more preferred over Tokyo for location, it is expected to have a larger segment than Tokyo. For the sake of simplicity, every issue and related issue values have the same color to distinguish categories easily. When the user hovers over the issue values in the chart, they can see the exact utility score of that issue. In addition, they can see the utility distribution of each issue by hovering over the percentage sign next to the graph. During the design stage of this framework, we did some pilot studies and got feedback regarding the design and location of those components. We finalized the design and location of those components based on that feedback.

4. Agent Design

Human-agent negotiations are fundamentally different from the agent-agent negotiations [2,22,48]. To investigate the effect of familiarity with the opponent in human-agent negotiation, we design a basic negotiating agent that takes human factors into account. One of the concerns regarding human-agent interaction is the time limitation and human patience. As our agent negotiates with its human counterparts, it does not have enough time to make random or same offers repeatedly to observe its opponent's behavior change. Moreover, repeating the same offer so often may bother the human partner and make them walk away. Therefore, our agent avoids falling repetition as much as possible. In addition, it is required to consider the opponent's behavior during the negotiation while generating the counteroffer and their preferences. Since the negotiating parties only know their own preferences and do not have any access to their opponent's preferences, our agent aims to learn them from the exchanged bids during the negotiation. In some studies [6,49], a negotiating party's behavior is classified into five categories: competing, avoiding, compromising, accommodating, and collaborating.

In our study, our agent also adopts this classification and tries to predict its opponent's attitude based on exchanged bids during the negotiation. Behavior classification is determined by the opponent's assertiveness and cooperativeness categorization as defined in [6]. To achieve this, we first calculate each participant's sensitivity to their opponent's preferences according to Equation (2) proposed by [50]. Here, sensitivity is calculated by considering the percentages of the negotiator's different moves. A move is determined based on the utility difference of the negotiator's subsequent offers for both sides. There are six different move types (fortunate, nice, concession, selfish, unfortunate, and silent). If sensitivity >1 , we consider the player as cooperative (C); if sensitivity <1 we classify the player as uncooperative (U). Otherwise, the player is considered as neutral (N) to the opponent's preferences. The assertiveness level of the opponent is determined according to the utility of the opponent according to the agent's own utility function as follows. We consider high, moderate, and low assertive if the utility of the bid is between [68–100], [34–67], and [0–33] respectively. The classification of opponent behavior is determined by the rules defined in Table 1. Accordingly, our agent adapts its bidding and communication way (e.g., facial expressions). To sum up, our agent follows a hybrid bidding strategy considering both the remaining time and behavior of the opponent. It considers three stages according to the normalized current time and adopts a different strategy. Stage is assigned to initial if $t < 0.1$; main if $0.1 \leq t \leq 0.9$, and final if $t > 0.9$.

$$Sensitivity_a(t) = \frac{\%Fortunate + \%Nice + \%Concession}{\%Selfish + \%Unfortunate + \%Silent} \quad (2)$$

Table 1. Behavior Classification.

Assertiveness	Cooperativeness	Behavior
High	Cooperative	Collaborative
High	Neutral	Competing
High	Uncooperative	Competing
Moderate	Cooperative	Accommodating
Moderate	Neutral	Compromising
Moderate	Uncooperative	Avoiding
Low	Cooperative	Accommodating
Low	Neutral	Avoiding
Low	Uncooperative	Avoiding

Our agent's negotiation strategy is illustrated in Figure 3. When the agent receives an offer from the opponent, it first checks if the offer is acceptable. Our agent adopts AC_{next} approach for this purpose [51]. It accepts the opponent's offer if the utility of that offer is higher than or equal to the utility of our agent's coming offer, $o^A_{current}$ (Lines 1–2). Otherwise, it generates its counteroffer as explained below. Note that the first offer of our agent is always the most preferred offer according to its own preferences. In the initial stage, the agent tries to recognize its opponent's attitude regarding cooperativeness. To achieve this, it simply compares the utility of the opponent's bids according to its own preferences. If the utility of the opponent's offer is greater than that of its previous offer, our agent considers its opponent as cooperative; otherwise, it is considered as uncooperative (Line 6). If the opponent is cooperative, then the agent also concedes and reduces its target utility value (Line 7); otherwise, it increases its target utility accordingly (Line 10). If the human opponent acts cooperatively, the agent demonstrates a happy facial expression (Line 8). Otherwise, it shows a frustrated facial expression (Line 11).

Algorithm 1 Next Offer Decision

o_{prev}^A and o_{prev}^H represents agent's and human's previous offer respectively while $o_{current}^A$ and $o_{current}^H$ denote their current offers. γ is the minimum acceptable value for the agent, also called as reservation value. s_c and θ_{beh} denote current state and estimated opponent behaviour respectively.

```

1: if  $U(o_{current}^H) \geq U(o_{current}^A)$  then
2:   acceptOffer()
3: else
4:   switch  $s_c$  do
5:     case initial
6:       if isCooperative( $o_{prev}^H, o_{current}^H$ ) then
7:          $u_{next} = \max(U(o_{prev}^A) - u_{dif}, \gamma)$ 
8:          $f_{exp} = happy$ 
9:       else
10:         $u_{next} = \min(U(o_{prev}^A) + u_{dif}, 1)$ 
11:         $f_{exp} = frustrated$ 
12:      end if
13:     case main
14:       switch  $\theta_{beh}$  do
15:         case competing
16:            $u_{next} = \max(U(o_{prev}^A), \gamma)$ 
17:            $f_{exp} = frustrated$ 
18:         case avoiding
19:            $u_{next} = getNextBestOffer(U(o_{prev}^A))$ 
20:            $f_{exp} = sad$ 
21:         case compromising
22:            $u_{next} = nearest(U(o_{prev}^A))$ 
23:            $f_{exp} = neutral$ 
24:         case accommodating
25:            $u_{next} = \max(U(o_{prev}^A) - u_{dif}, \gamma)$ 
26:            $f_{exp} = happy$ 
27:         case collaborating
28:            $\alpha_{next} = selectBestForBothFromEstimatedPareto()$ 
29:            $f_{exp} = excited$ 
30:       case final
31:          $\alpha_{next} = selectBestForBothFromEstimatedPareto()$ 
32:          $f_{exp} = frustrated$ 
33:       if  $\alpha_{next}$  is not assigned then
34:          $\alpha_{next} = pickOfferWithUtility(u_{next})$ 
35:       end if
36:       sendOffer( $\alpha_{next}, f_{exp}$ )
37:   end if

```

Figure 3. Offer Generation Algorithm.

In the main stage, our primary focus is to detect our opponent's attitude regarding the classification mentioned above and act accordingly. Our agent adopts a frustrated facial expression and makes an offer whose utility is the same as its previous one (Lines 15–17). For the avoiding opponent, it selects the next best offer according to the utility of its previous bid and shows a sad facial expression (Line 18–20). A compromising opponent makes a silent move, so our agent does too by making an offer whose utility is approximately the same as its previous offer and adopts a neutral facial expression (Line 21–23). For the accommodating opponent, a happy facial expression is adopted to express appreciation, and the agent concedes with the same concession value in a fair way (Line 24–26). When the opponent's attitude is detected as collaborating, the agent tries to make an offer which is good for both sides. To do this, the agents estimate Pareto-optimal offers and make the best offer from these candidates with a smiling facial expression (Line 27–29). The best

offer is determined by ordering multiplication of both sides' utility (Nash product) in the Pareto-optimal offers.

In the final stage of the negotiation (i.e., approaching the deadline), our agent orders the estimated Pareto-optimal offers according to the social welfare (i.e., the sum of utilities for both sides). The agent asks for offers in order at this stage. It shows a frustrated facial expression to create pressure on its opponent.

5. Research Methodology

To study the effect of familiarity with the opponent (i.e., celebrity versus non-celebrity unknown opponent) in human-agent negotiation, we conducted a user experiment. Accordingly, we suggest and examine three hypotheses as follows:

Hypothesis 1 (H1): *The negotiation outcome reached by the human negotiators would be significantly different when they negotiate with a celebrity virtual agent than that of when they negotiate with a non-celebrity agent irrespective of the hedonic tone of their feelings towards the chosen celebrity.*

Hypothesis 2 (H2): *The negotiation outcome reached by the human negotiators would be significantly different when they negotiate with a celebrity virtual agent for whom they have positive feelings than that of when they negotiate with a non-celebrity agent.*

Hypothesis 3 (H3): *The agent employing the designed negotiation strategy can receive higher utility compared to the utility of their human counterparts.*

In order to test the aforementioned hypotheses, we follow the road map illustrated in Figure 4. Participants are exposed to a single experimental condition in a between-group design, whereas in a within-group design, they are exposed to all the experimental conditions [52]. In a between-group experiment, the performance of one group of participants is compared with the performance of another group of participants. Therefore, individual differences may significantly impact the results. Furthermore, it requires a higher number of participants than the within-subject design. Consequently, we decided to follow a within-group design in this study. Accordingly, in the following section, we explain our experimental setting.

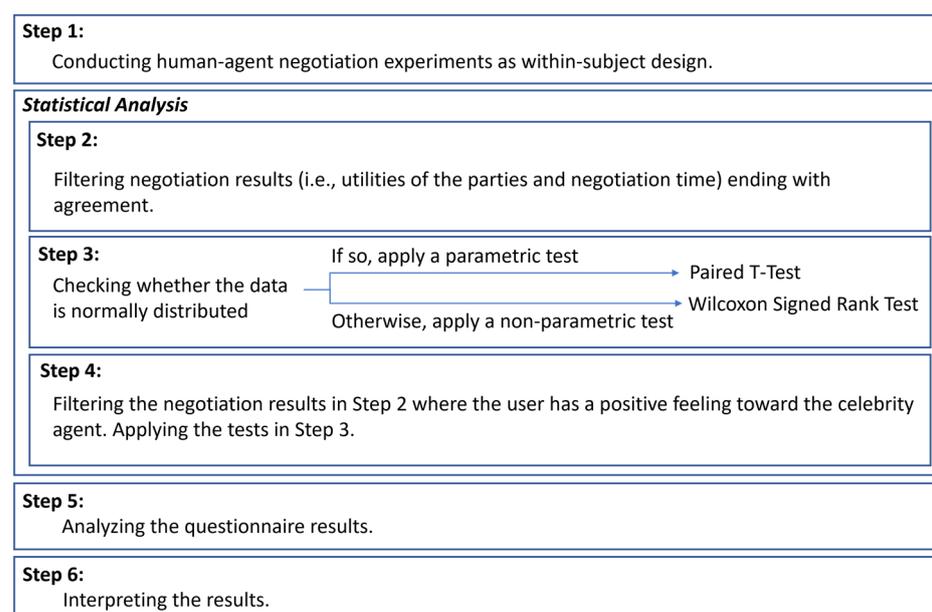


Figure 4. Empirical Research Road Map.

In our experiments, we aim to investigate the effect of negotiating with a celebrity virtual agent versus negotiating with a non-celebrity virtual agent on both the negotiation process and outcome. We have recruited 67 participants (i.e., some acquaintances and university students; 39 males, 28 females; median age: 25) for conducting our human-virtual agent experiments. Of the participants, 49 had an engineering background, while the rest had different backgrounds such as social sciences, business management, and law. It was a within-subject design. Therefore, each participant negotiated with the celebrity avatar and non-celebrity avatar in consecutive sessions. To reduce the learning effect, we divided the participants into two groups where in one group, they firstly negotiated with the non-celebrity avatar and then with the celebrity avatar. In the other group, the order of sessions is in the opposite way. Each participant has a 10-min break between their negotiation sessions.

In the experiment, a negotiation profile is given to each participant. As a role-playing game, they are asked to study their preference profiles elaborately before their negotiation. After studying their profiles, they are asked to answer three questions regarding their preferences to understand whether they grasp their preferences accurately. They specify the best and worst offers and determine which offer is more preferred among a given offer pair. In the beginning, all participants were informed about the experiment and asked to fill out a short survey form in which the demographics of the participants were collected. At this stage, the consent of the participants was taken. Before starting their real negotiation sessions, there was a training session for participants to experience the negotiation process and the framework. For the training session, a different profile from the real experiment was given to participants. The same process with the real experiment is used. Participants first watched a demonstration video and performed a five-minute negotiation on the given training scenario. During the training session, participants negotiated with a computerized virtual agent (Figure 2), which employed a simple random strategy. Note that the virtual character and agent strategies differ from those we used in the experiment.

After the training session, the participants were asked to choose one celebrity among six celebrities (two singers, two Hollywood celebrities, and two businesspeople) and specify the reason why they chose that celebrity. Afterward, the participants were asked to study the preference profile for their first real negotiation session. Note that the participants negotiate with a non-celebrity agent that has the same gender as their choice of celebrity agent in their session with a non-celebrity avatar to eliminate the effect of their opponent's gender. The deadline for both negotiation sessions is 10 min. In both negotiations, virtual agents employ the same negotiation strategy. Although they negotiate on the same negotiation problem, the preference profiles seem different in both sessions to reduce the learning effect. It is important to note that the overall utility distribution percentages for the preferences profiles are the same for both sessions. In contrast, the percentages of issue values and issues weights are rearranged with these predefined percentages. Consequently, we can compare the outcomes of both sessions fairly.

The time is initiated after the human participant's first offer. If participants cannot agree within 10 min, both parties receive zero points. Here, the goal of the participants is to maximize their individual utility scores, and participants could see their preference profiles at any time during the negotiation. It is also worth noting that participants were advised to look at the facial expression of the avatar agent during their negotiation.

According to our scenario, our participants are asked to negotiate with the virtual agent on their joint holiday plan like the one [6]. There are four negotiation issues: location (Tokyo, Bali, London, Paris, Berlin), accommodation (3-star hotel, guest house, 5-star luxury hotel, and camping), duration (3 days, 1 week, 2 weeks, and 3 weeks) and activities (Museum tours, trekking and historical places). That is, there are 240 possible outcomes. The bargaining power of the two parties is almost the same. Figure 5 shows the utilities of each possible bid in the given scenario as well as the agreement zone. After their negotiation sessions, the participants are asked to fill out two post-surveys. In the first survey, there are questions to compare their two negotiation sessions. After informing

them that they negotiated with the same negotiating agent in both sessions, they fill out the second survey consisting of general questions about their recent negotiation experience (e.g., their own negotiation strategy, their thoughts about their opponent's negotiation attitude) and experimental setup (e.g., clearance of instructions given during the experiment). In the following section, we report our findings.

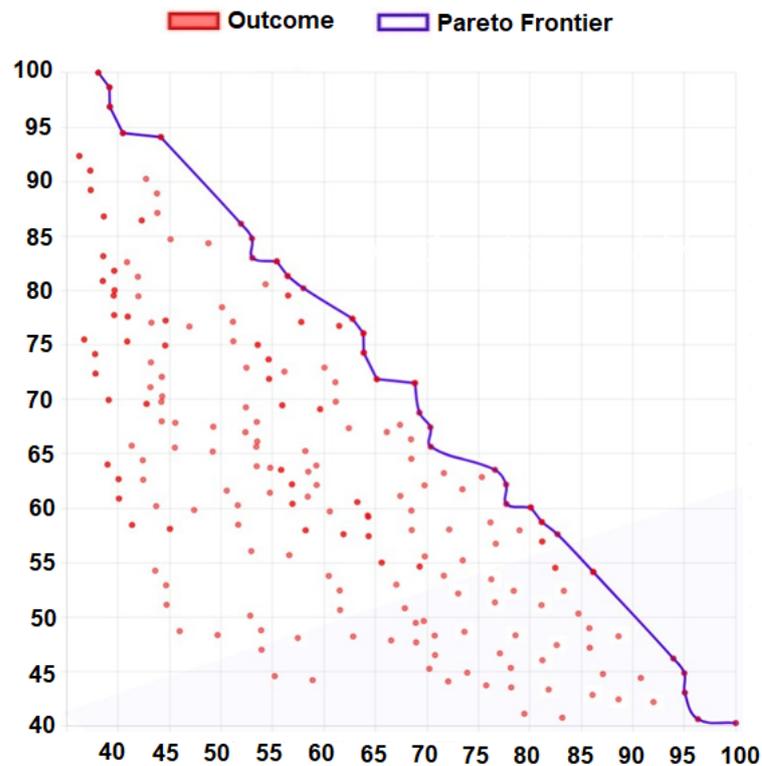


Figure 5. Outcome Space.

6. Experimental Results

After conducting the experiments, applying a convenient statistical test is essential. As specified in Figure 4, we also need to determine whether a parametric or non-parametric test should be applied. Before applying the normality tests to select the right statistical test, the negotiation results without agreement are eliminated from the data since they would act as an outlier. For more fine-grained analysis (i.e., considering results only when the participants negotiate with celebrities with whom they have positive feelings), we filtered the data one more time and apply a convenient statistical test. In addition to the negotiation results, we analyzed the responses to the questionnaire to obtain insights into how participants perceived their opponents. In Sections 6.1 and 6.2, we interpret those results.

6.1. Overall Experimental Results

We first investigate the number of agreements. Recall that 67 participants in total negotiated in both settings. Figure 6 shows the number of agreements in each setting. It can be seen that 4 of 67 sessions with non-celebrity avatars failed to find a consensus, while 3 of the 67 sessions ended with a disagreement for celebrity sessions. In 35 sessions with non-celebrity avatars and with celebrity avatars ended up with user acceptance.

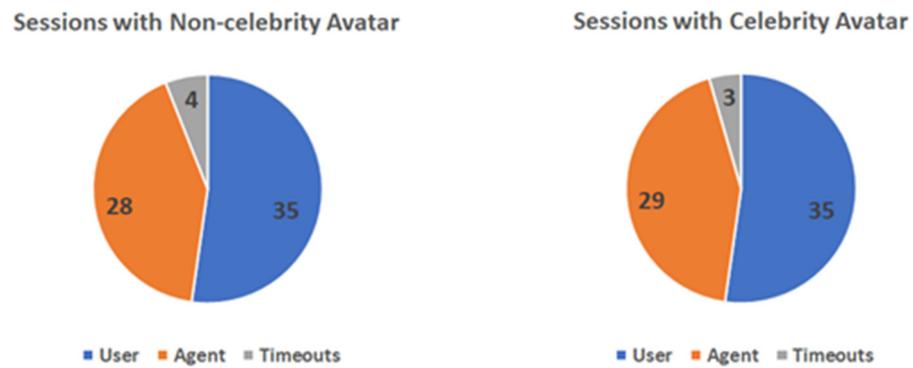


Figure 6. Number of Agreements in Each Negotiation Setting.

After analyzing the number of agreements, we study the negotiation results of the sessions ended successfully for both settings ($N = 60$) in terms of utility received by the user (i.e., user utility), the utility received by the agent (i.e., agent utility) and the negotiation time to reach an agreement. Note that when time out occurs and no agreement is reached, both negotiating parties receive zero utility. Since unsuccessful negotiations would act like outliers, we filtered them out in our analysis. To apply the proper statistical test, we first applied a normality test to see whether or not the data is normally distributed. To achieve this, we used Kolmogorov–Smirnov test on the data distribution regarding user and agent utility of the agreements and negotiation time. All data except the negotiation time ($p = 0.076$ for Non-Celebrity and $p = 0.200$ for Celebrity) are not normally distributed according to these statistical tests (see p values in the tenth column of Table 2). Therefore, we applied a non-parametric statistical test, namely the Wilcoxon signed-rank test.

Table 2. Statistics for only Successful Negotiation Sessions.

	N	Mean	Stdev	Median	Min	Max	1st Q.	3rd Q.	Normality
User Utility (Non-C.)	60	0.58	0.11	0.54	0.23	0.75	0.51	0.69	0.00
User Utility (C.)	60	0.56	0.13	0.54	0.25	0.94	0.50	0.71	0.00
Agent Utility (Non-C.)	60	0.82	0.11	0.83	0.62	01.00	0.71	0.94	0.00
Agent Utility (C.)	60	0.81	0.11	0.83	0.54	0.95	0.70	0.94	0.00
Total Time (Non-C.)	60	0.59	0.26	0.58	0.06	0.98	0.38	0.83	0.14
Total Time (C.)	60	0.57	0.26	0.55	0.08	0.99	0.31	0.82	0.07

Table 2 shows the statistics such as mean values, standard deviations, medians, minimum values, maximum values, first quartile, and second quartile for each performance metric in both non-celebrity and celebrity settings. An error bar chart graph was also provided based on the mean and standard deviation values in Figure 7. When the Wilcoxon signed-rank test is conducted, we found no significant difference ($p > 0.05$) between any of the performance metrics ($p = 0.14$, $z = -1.49$ for user utility, $p = 0.53$, $z = -0.62$ for agent utility and $p = 0.62$, $z = -0.50$ for negotiation time). Consequently, H1 is not supported by the given results. Furthermore, when we compare the utilities of the user and the agent, we found a statistically significant difference between agent and user utilities in both settings ($p < 0.01$), according to Wilcoxon signed-rank test ($z = -6.04$, $z = -5.59$ for non-celebrity and celebrity settings, respectively). As a result, H3 is supported by the collected negotiation utility data.

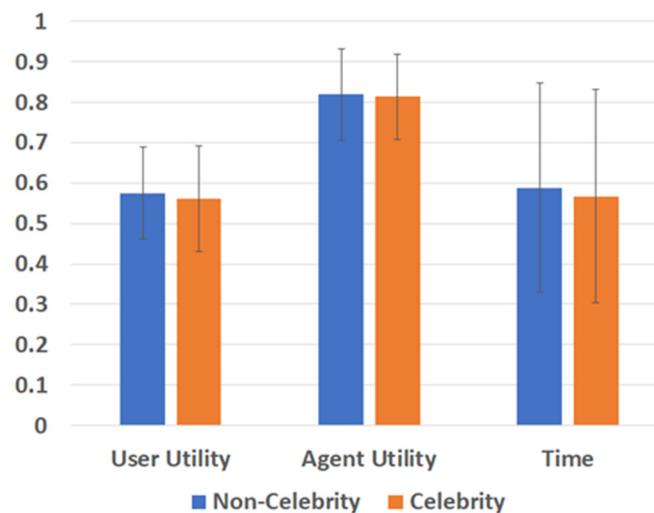


Figure 7. Mean with Error Bars of Negotiation Performance Metrics.

Moreover, we analyzed the participants' responses to the post questionnaires. In the first post-survey, we asked comparison questions between celebrity and non-celebrity avatars. A 9-point Likert scale is used where 1 and 9 represent Non-celebrity avatar and Celebrity avatar, respectively. These results supported that they felt more familiar with the celebrity avatar (average rating 6.3).

In the second post-survey, we asked the participants their perceptions of the avatar's strategy, the framework, and the experimental setup by using 9-point Likert scale survey questions. Regarding the responses to the experimental setup, it is seen that participants understood the given profile (average rating 8.0), and they took into account the avatar's facial expressions during their negotiation (average rating 6.9). We noticed that most of the participants found the virtual agent competitive in general (average rating 7.3). This is also supported by the negotiation results, where the agent receives higher utility than the human negotiator.

We also asked them the reason behind their celebrity avatar selection. There were some positive, negative and neutral answers in the given choices. We analyzed this specific question and filtered those who claimed a positive feeling for the chosen celebrity avatar. Accordingly, the following section investigates whether H2 is supported.

6.2. Experimental Results Investigating H2

A total of 29 participants specified that they have positive feelings towards their chosen celebrities. A further 13 participants claimed that they have negative feelings and 11 participants were neutral according to the questionnaire results. In addition, there were seven participants who couldn't be included in any of these groups. Regarding the acceptance rates, nine users accepted in 29 sessions with the non-celebrity avatar, while 11 participants accepted in 29 sessions with celebrity sessions (see Figure 8). All negotiations end with an agreement.

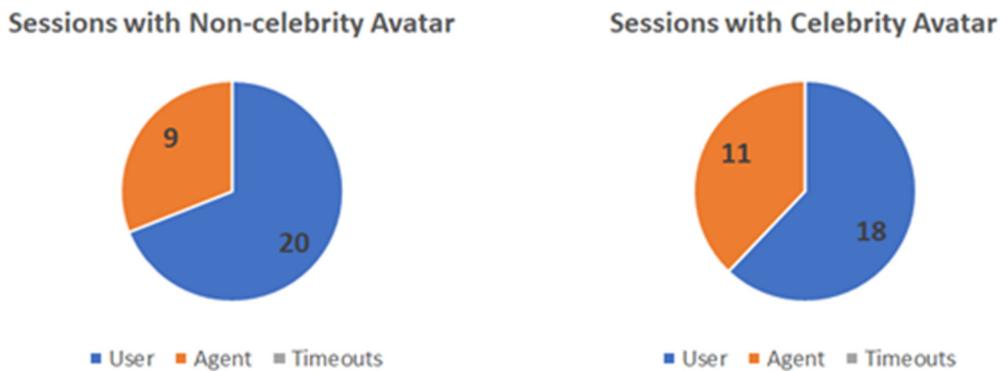


Figure 8. Acceptance Rates for H2 Related Data.

After analyzing the acceptance rates, we studied the negotiation results of the sessions that ended successfully (N = 29). We first checked whether the data is normally distributed. According to the Kolmogorov–Smirnov Normality Test, the negotiation time data is normally distributed ($p = 0.200$ for the Non-Celebrity and $p = 0.200$ for the Celebrity). However, the user and agent utilities are not normally distributed. Therefore, we applied the Wilcoxon signed-rank test again.

Table 3 shows the performance metrics in both non-celebrity and celebrity settings. An error bar chart graph was also provided based on the mean and standard deviation values in Figure 9. In terms of user utility, the results are statistically significantly different according to the Wilcoxon signed-rank test ($p = 0.049$ and $z = -1.972$). As expected, the average user utility received in the celebrity setting is lower than that in the non-celebrity setting (0.55 versus 0.58). It seems that participants tended to concede more against a celebrity opponent for whom they have positive feelings. The results show that participants accepted offers with lower utilities when they negotiated with the celebrity avatar. They may not consider the fairness of the negotiation outcome much.

Table 3. Statistics for only Successful Negotiation Sessions (participants with positive feelings).

	N	Mean	Stdev	Median	Min	Max	1st Q.	3rd Q.	Normality
User Utility (Non-C.)	29	0.58	0.12	0.54	0.23	0.75	0.52	0.69	0.00
User Utility (C.)	29	0.55	0.15	0.50	0.25	0.94	0.48	0.69	0.00
Agent Utility (Non-C.)	29	0.84	0.12	0.94	0.66	1.00	0.70	0.94	0.00
Agent Utility (C.)	29	0.82	0.11	0.83	0.54	0.95	0.73	0.91	0.00
Total Time (Non-C.)	29	0.561	0.24	0.65	0.13	0.98	0.41	0.77	0.20
Total Time (C.)	29	0.54	0.26	0.54	0.08	0.99	0.30	0.79	0.20

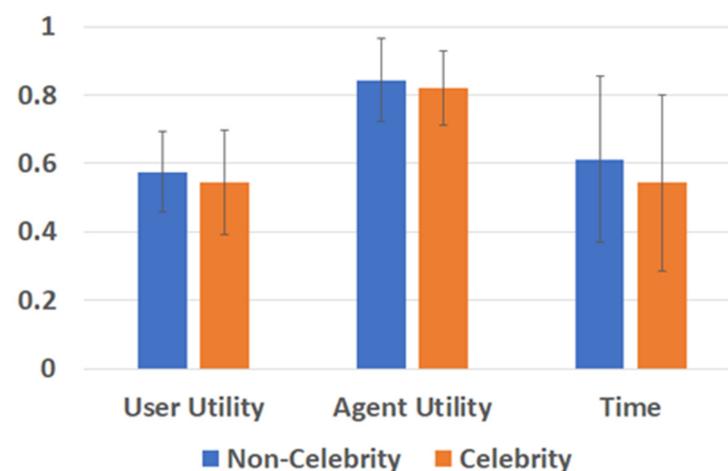


Figure 9. Mean with Error Bars of Negotiation Performance Metrics (participants with positive feelings).

Consequently, these results support **H2**. Apart from this, there is no statistically significant difference in terms of agent utilities between celebrity and non-celebrity settings. It is important to recall that the participants do not know their opponent's preferences. They may expect that conceding in their own utility would favor their opponent. Moreover, there is no statistically significant difference in negotiation time between celebrity and non-celebrity settings. This indicates that participants made similar efforts during their negotiations with celebrity and non-celebrity avatars.

Regarding the 9-point Likert scale questionnaire results where 5 refers to neutral and above 5 favors the celebrity avatar, when they compare their negotiations with celebrity and non-celebrity avatars, it can be observed that:

- They were more comfortable when they negotiated with the celebrity avatar (average rating $6.03 > 5$).
- They felt more familiar with the celebrity avatar (average rating $6.5 > 5$).
- They acted more collaboratively when they negotiated with the celebrity avatar (average rating $5.62 > 5$).
- They found the celebrity avatar more friendly to them (average rating $5.9 > 5$).

Questionnaire responses above also support **H2**. In the second post-survey, it is seen that participants understood the given profile (average rating 7.80), and they took into account the avatar's facial expressions during their negotiation (average rating 7.14). Lastly, we noticed that most of the participants found the virtual agent competitive in general (average rating 7.41).

7. Discussion

This work introduces a human-agent negotiation framework where a human negotiator can negotiate with a virtual avatar so that researchers can investigate the effect of the human negotiator's limited familiarity with their opponent on negotiation outcome and process. Our contributions are two-fold. Firstly, we introduce a new human-agent negotiation framework in which the virtual avatar employs a novel negotiation strategy inspired by the existing strategies. The proposed strategy considers the opponent's behavior and preferences, as well as the remaining negotiation time. Secondly, we conduct a user experiment in which participants negotiate with an agent whose appearance and voice are a replica of those chosen celebrities and also negotiate with an agent who is not familiar with them (non-celebrity). We compare the outcomes of those negotiations to find out the effect of limited knowledge about opponents during negotiation. To the best of our knowledge, our work is the first study pursuing this research question in human-agent negotiations. Furthermore, the designed strategy is able to beat human negotiators.

When we analyzed the experimental results, we observe that the negotiation outcomes reached by the human negotiators were not significantly different when they negotiated with a celebrity virtual agent than that of when they negotiated with a non-celebrity agent irrespective of the hedonic tone of their feelings towards the chosen celebrity (**H1**). On the other hand, it is seen that human participants tended to concede more against their celebrity opponent than a non-celebrity opponent when they had positive feelings for the chosen celebrity. That is, the participants acted more collaboratively when they negotiated with their favored celebrity avatar (**H2**). While designing a virtual agent negotiating with a human, it might be beneficial to build a virtual avatar that looks like a celebrity that the user likes. Regarding the performance of the designed negotiation strategy, experimental results show that the agent employing this strategy usually beats its human opponents (**H3**). Participants also specified that they found their opponent competitive. As future work, it might be interesting to study the effect of the opponent's competitiveness in human-agent negotiation settings.

The paradigm of conversationally negotiating with agents that appear like humans brings a naturalness to the interaction. In this new paradigm, many phenomena seen in human-human interactions could be translated into human-agent interactions. For example, in addition to visual similarity, rapport has shown to play a significant effect on the process

and outcome of negotiations between two parties [53]. It would be of interest to see how an AI agent is able to create such a rapport. Especially in a multi-modal paradigm where users also see the embodiment of the AI agent as well as converse with it, it would be of interest to explore how AI agents can express themselves using voice, emotion, facial expressions, gestures, etc. to develop rapport and come-off as a fair, trustworthy negotiation partner, an important prerequisite to successful negotiations.

Furthermore, it would be interesting to investigate how human negotiators act against a virtual agent whose appearance looks like a celebrity for whom they have negative feelings. However, we could not reach a sufficient number of participants for this particular case in our experiments. Therefore, we left this issue for future work.

Author Contributions: Conceptualization, B.T. and R.A.; methodology, B.T., R.A. and C.S.Ö.; software, B.T.; formal analysis, B.T., R.A. and C.S.Ö.; writing—original draft preparation, B.T.; writing—review and editing, B.T., R.A. and C.S.Ö.; supervision, R.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: The experiment protocol in this study was approved by the Ethics Committee of Özyeğin University on 21 June 2021.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: Due to the statements in the consent form, the subjects' data could not be shared publicly.

Acknowledgments: Berkay Türkgeldi was supported by Vestel Elektronik Sanayi ve Ticaret A.Ş. during his master studies. We would like to thank Rahul Divekar and our research group members particularly Onur Keskin, Anıl Doğru, Cihan Eran, Gevher Yesevi, Ertan Yildiz, Umut Çakan, Seda Çalikkocaoglu and Zeynep Ümitvar for their feedback and technical support.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Jennings, N.R.; Faratin, P.; Lomuscio, A.R.; Parsons, S.; Sierra, C.; Wooldridge, M. Automated negotiation: Prospects, methods and challenges. *Int. J. Group Decis. Negot.* **2001**, *10*, 199–215. [[CrossRef](#)]
- Aydoğan, R.; Keskin, O.; Çakan, U. Let's negotiate with Jennifer! Towards a Speech-based Human-Robot Negotiation. In *Advances in Automated Negotiations*; Ito, T., Zhang, M., Aydoğan, R., Eds.; Springer: Singapore, 2020; pp. 3–16.
- Aydoğan, R.; Keskin, O.; Çakan, U. Would you imagine yourself negotiating with a robot, Jennifer? *IEEE Trans. Hum. Mach. Syst.* **2021**, *52*, 41–51. [[CrossRef](#)]
- Oshrat, Y.; Lin, R.; Kraus, S. Facing the challenge of human-agent negotiations via effective general opponent modelling. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, Budapest, Hungary, 10–15 May 2009; pp. 377–384.
- Druckman, D.; Broome, B.J. Value Differences and Conflict Resolution: Familiarity or Liking? *J. Confl. Resolut.* **1991**, *35*, 571–593. [[CrossRef](#)]
- Güngör, O.; Çakan, U.; Aydoğan, R.; Öztürk, P. Effect of Awareness of Other Side's Gain on Negotiation Outcome, Emotion, Argument and Bidding Behavior. In *Proceedings of the Twelfth International Workshop on Agent-Based Complex Automated Negotiations*, ACAN 2019, Macao, China, 10–13 August 2019; pp. 377–384.
- Moreland, R.L.; Zajonc, R.B. Exposure effects in person perception: Familiarity, similarity, and attraction. *J. Exp. Soc. Psychol.* **1982**, *18*, 395–415. [[CrossRef](#)]
- Keskin, M.O.; Çakan, U.; Aydoğan, R. Solver Agent: Towards Emotional and OpponentAware Agent for Human-Robot Negotiation. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '21)*, Virtual Event, 3–7 May 2021; pp. 1557–1559.
- Yuasa, M.; Mukawa, N. The facial expression effect of an animated agent on the decisions taken in the negotiation game. In *Proceedings of the CHI'07 Conference on Human Factors in Computing Systems*, Denver, CO, USA, 6–11 May 2007; pp. 2795–2800.
- Lin, W.J.; Hu, C.H.; Lai, H. The Impact of Gender Differences on Response Strategy in e-Negotiation. In *Workshop on E-Business*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 192–205.
- Stuhlmacher, A.F.; Citera, M.; Willis, T. Gender differences in virtual negotiation: Theory and research. *Sex Roles* **2007**, *57*, 329–339. [[CrossRef](#)]

12. Van der Lubbe, L.M.; Bosse, T. Studying gender bias and social backlash via simulated negotiations with virtual agents. In Proceedings of the International Conference on Intelligent Virtual Agents, Stockholm, Sweden, 27–30 August 2017; Springer: Berlin/Heidelberg, Germany, 2017; pp. 455–458.
13. Aydođan, R.; Yolum, P. Learning Opponents Preferences for Effective Negotiation: An Approach Based on Concept Learning. *J. Auton. Agents Multi-Agent Syst.* **2012**, *24*, 104–140. [[CrossRef](#)]
14. Baarslag, T.; Kaisers, M.; Gerding, E.H.; Jonker, C.M.; Gratch, J. When will negotiation agents be able to represent us? The challenges and opportunities for autonomous negotiators. In Proceedings of the Twenty-sixth International Joint Conference on Artificial Intelligence, Melbourne, Australia, 19–25 August 2017; pp. 4684–4690.
15. Cao, M.; Luo, X.; Luo, X.R.; Dai, X. Automated negotiation for e-commerce decision making: A goal deliberated agent architecture for multi-strategy selection. *Decis. Support Syst.* **2015**, *73*, 1–14. [[CrossRef](#)]
16. Fatima, S.; Kraus, S.; Wooldridge, M. *Principles of Automated Negotiation*; Cambridge University Press: Cambridge, UK, 2014.
17. Fujita, K.; Ito, T.; Klein, M. Efficient issue-grouping approach for multiple interdependent issues negotiation between exaggerator agents. *Decis. Support Syst.* **2014**, *60*, 10–17. [[CrossRef](#)]
18. Marsa-Maestre, I.; Klein, M.; Jonker, C.M.; Aydođan, R. From Problems to Protocols: Towards a Negotiation Handbook. *Decis. Support Syst.* **2014**, *60*, 39–54. [[CrossRef](#)]
19. Razeghi, Y.; Yavuz, O.; Aydođan, R. Deep Reinforcement Learning for Acceptance Strategy in Bilateral Negotiations. *Turk. J. Electr. Eng. Comput. Sci.* **2020**, *28*, 1824–1840. [[CrossRef](#)]
20. Sanchez-Anguix, V.; Aydođan, R.; Julian, V.; Jonker, C. Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Electron. Commer. Res. Appl.* **2014**, *13*, 243–265. [[CrossRef](#)]
21. Jonker, C.M.; Aydođan, R.; Baarslag, T.; Fujita, K.; Ito, T.; Hindriks, K. Automated Negotiating Agents Competition (ANAC). In Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17), San Francisco, CA, USA, 4–9 February 2017; AAAI Press: Palo Alto, CA, USA, 2017; pp. 5070–5072.
22. Mell, J.; Gratch, J.; Aydođan, R.; Baarslag, T.; Jonker, C.M. The LikeabilitySuccess Tradeoff: Results of the 2nd Annual Human-Agent Automated Negotiating Agents Competition. In Proceedings of the 8th International Conference on Affective Computing and Intelligent Interaction (ACII), Cambridge, UK, 3–6 September 2019; pp. 1–7.
23. Mell, J.; Gratch, J. Grumpy & Pinocchio: Answering Human-Agent Negotiation Questions through Realistic Agent Design. In Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, São Paulo Brazil, 8–12 May 2017; pp. 401–409.
24. Gratch, J.; Lucas, G. Negotiation as a Challenge Problem for Virtual Humans. In Proceedings of the Fifteenth International Conference on Intelligent Virtual Agents, Delft, The Netherlands, 26–28 August 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 201–215.
25. Jonker, C.M.; Aydođan, R. Deniz: A Robust Bidding Strategy for Negotiation Support Systems. In *Advances in Automated Negotiations*; Ito, T., Zhang, M., Aydođan, R., Eds.; Springer: Berlin/Heidelberg, Germany, 2020; pp. 29–44.
26. Haim, G.; Gal, Y.; Gelfand, M.; Kraus, S. A cultural sensitive agent for human-computer negotiation. In Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, Valencia, Spain, 4–8 June 2012; pp. 451–458.
27. De Melo, C.M.; Carnevale, P.; Gratch, J. The effect of expression of anger and happiness in computer agents on negotiations with humans. In Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems, Taipei, Taiwan, 2–6 May 2011; pp. 937–944.
28. Prajod, P.; Al Owayyed, M.; Rietveld, T.; van der Steeg, J.J.; Broekens, J. The Effect of Virtual Agent Warmth on Human-Agent Negotiation. In Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, Montreal, QC, Canada, 13–17 May 2019; pp. 71–76.
29. Rahwan, I.; Ramchurn, S.D.; Jennings, N.R.; Mcburney, P.; Parsons, S.; Sonenberg, L. Argumentation-based negotiation. *Knowl. Eng. Rev.* **2003**, *18*, 343–375. [[CrossRef](#)]
30. Divekar, R.R.; Mou, X.; Chen, L.; De Bayser, M.G.; Guerra, M.A.; Su, H. Embodied Conversational AI Agents in a Multi-modal Multi-agent Competitive Dialogue. In Proceedings of the 28th International Joint Conference on Artificial Intelligence, Macao, China, 10–16 August 2019; pp. 6512–6514.
31. Broekens, J.; Harbers, M.; Brinkman, W.P.; Jonker, C.; Bosch, K.; Meyer, J.J. Virtual Reality Negotiation Training Increases Negotiation Knowledge and Skill. In Proceedings of the International Conference on Intelligent Virtual Agents, Santa Cruz, CA, USA, 12–14 September 2012; Springer: Berlin/Heidelberg, Germany, 2012; pp. 218–230.
32. Ding, D.; Burger, F.; Brinkman, W.P.; Neerincx, M. Virtual Reality Negotiation Training System with Virtual Cognitions. In *Intelligent Virtual Agents*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 119–128.
33. Gratch, J.; DeVault, D.; Lucas, G. The Benefits of Virtual Humans for Teaching Negotiation. In Proceedings of the International Conference on Intelligent Virtual Agents, Los Angeles, CA, USA, 20–23 September 2016; pp. 283–294.
34. Jonker, C.M.; Aydođan, R.; Baarslag, T.; Broekens, J.; Detweiler, C.A.; Hindriks, K.V.; Hultgren, A.; Pasma, W. An Introduction to the Pocket Negotiator: A General Purpose Negotiation Support System. In Proceedings of the 14th European Conference on Multi-Agent Systems, Valencia, Spain, 15–16 December 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 13–27.
35. Mell, J.; Gratch, J. IAGO: Interactive arbitration guide online. In Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, Singapore, 9–13 May 2016; pp. 1510–1512.

36. Rosenfeld, A.; Zuckerman, I.; Segal-Halevi, E.; Drein, O.; Kraus, S. NegoChat: A chat-based negotiation agent. In Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems, Paris, France, 5–9 May 2014; pp. 525–532.
37. Lin, J.; Huff, S.L.; Newson, E.F.P.; Amoroso, D. Efficiency in computer-mediated negotiation: The familiarity factor. In Proceedings of the ASAC Conference, Halifax, NS, Canada, 16–18 June 1988; pp. 1–12.
38. Mell, J.; Gratch, J.; Lucas, G. The Effectiveness of Competitive Agent Strategy in Human-Agent Negotiation. In Proceedings of the American Psychological Association’s Technology, Mind, and Society Conference, Washington, DC, USA, 5–7 April 2018; pp. 125–132.
39. Blankendaal, R.; Bosse, T.; Gerritsen, C.; de Jong, T.; de Man, J. Are Aggressive Agents as Scary as Aggressive Humans? In Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS ’15), Istanbul, Turkey, 4–8 May 2015; pp. 553–561.
40. Adams, H.; Thompson, C.; Thomas, D.; Sharis, F.; Jernigan, C.G.; Moore, C.; Williams, B. The Effect of Interpersonal Familiarity on Cooperation in a Virtual Environment. In Proceedings of the ACM SIGGRAPH Symposium on Applied Perception, Tübingen, Germany, 13–14 September 2015; Association for Computing Machinery: New York, NY, USA, 2015; p. 138.
41. Wauck, H.; Lucas, G.; Shapiro, A.; Feng, A.; Boberg, J.; Gratch, J. Analyzing the Effect of Avatar Self-Similarity on Men and Women in a Search and Rescue Game. In Proceedings of the CHI Conference on Human Factors in Computing Systems, New York, NY, USA, 21–26 April 2018; pp. 1–12.
42. Reder, L.; Ritter, F. What Determines Initial Feeling of Knowing? Familiarity with Question Terms, Not with the Answer. *J. Exp. Psychol. Learn. Mem. Cogn.* **1992**, *18*, 435–451. [[CrossRef](#)]
43. Stuhlmacher, A.; Champagne, M. The Impact of Time Pressure and Information on Negotiation Process and Decisions. *Group Decis. Negot.* **2000**, *9*, 471–491. [[CrossRef](#)]
44. Sheffield, J. The effect of communication medium on negotiation performance. *Group Decis. Negot.* **1995**, 159–179. [[CrossRef](#)]
45. Aydoğan, R.; Festen, D.; Hindriks, K.; Jonker, C.M. Alternating Offers Protocol for Multilateral Negotiation. In *Modern Approaches to Agent-Based Complex Automated Negotiation*; Fujita, K., Bai, Q., Ito, T., Zhang, M., Ren, F., Aydoğan, R., Hadfi, R., Eds.; Springer: Tokyo, Japan, 2017; pp. 153–167.
46. Inc. MotionPortrait. Motion Portrait. 2020. Available online: <https://www.motionportrait.com/> (accessed on 11 August 2022).
47. Headliner Voice. 2020. Headliner Voice. Available online: <https://voice.headliner.app/> (accessed on 11 August 2022).
48. Lin, R.; Kraus, S. Can automated agents proficiently negotiate with humans? *Commun. ACM* **2010**, *53*, 78–88. [[CrossRef](#)]
49. Thomas, K.W.; Kilmann, R.H. *Thomas-Kilmann Conflict Mode:TKI Profile and Interpretive Report*; CPP, Inc.: Sunnyvale, CA, USA, 2008; pp. 1–11.
50. Hindriks, K.V.; Tykhonov, D. Let’s dans! An analytic framework of negotiation dynamics and strategies. *Web Intell. Agent Syst.* **2011**, *9*, 319–335. [[CrossRef](#)]
51. Baarslag, T.; Hindriks, K.V.; Jonker, C.M. Acceptance Conditions in Automated Negotiation. In Proceedings of the ICT, Veldhoven, The Netherlands, 14–15 November 2011; pp. 95–111.
52. Lazar, J.; Feng, J.; Hochheiser, H. Experimental Design. In *Research Methods in Human-Computer Interaction*, 2nd ed.; Green, T., Ed.; Elsevier: Cambridge, MA, USA, 2017.
53. Cundiff, N.L.; Kim, K.; Choi, S.B. Emotional Intelligence and Negotiation Outcomes: Mediating Effects of Rapport, Negotiation Strategy, and Judgment Accuracy. *Group Decis. Negot.* **2014**, *24*, 477–493.