Let's negotiate with Jennifer! Towards a Speech-based Human-Robot Negotiation

Reyhan Aydoğan[1,2], Onur Keskin and Umut Çakan

[1] Department of Computer Science, Özyeğin University, Istanbul, Turkey
{reyhan.aydogan,onur.keskin,umut.cakan}@ozyegin.edu.tr
[2] Interactive Intelligence Group Delft University of Technology Delft, The Netherlands

Abstract. Social robots are becoming prevalent in our society and most of the time they need to interact with humans in order to accomplish their tasks. Negotiation is one of the inevitable processes they need to be involved to make joint decisions with their human counterparts when there is a conflict of interest between them. This paper pursues a novel approach for a humanoid robot to negotiate with humans efficiently via *speech*. In this work, we propose a speech-based negotiation protocol in which agents make their offers in a turn-taking fashion via speech. We present a variant of time-based concession bidding strategy for the humanoid robot and evaluated the performance of the robot against human counterpart in human-robot negotiation experiments.

1 Introduction

Nowadays, robots are becoming an intrinsic part of everyday life and in the near future we will presumably see them almost everywhere in our society (e.g. work, home, hospital, school) [11,30]. Even now, robots get to work in classrooms, hotels, airports and so on. As robots become more human-like, they will be integrated into the society more widely. To carry out our tasks, we may need to collaborate with them more often and even negotiate with them on task allocations or other daily issues. For this reason, there is an urgent need for designing negotiating agents that can interact with humans effectively.

Although automated negotiation have been studied for several decades and there already exist a variety of negotiation frameworks [8,25,2,9,3,1,14,18,17], the nature of "human-agent negotiation" requires considering different dynamics [16]. For example, it is possible to make hundreds of offers to reach an agreement in automated negotiation; however, this is not feasible for human negotiator. Moreover, the way of communication is another issue to be taken into account in human-agent negotiation as the rules of interaction. Therefore, a number of studies have been carried out on designing protocols for human-agent negotiation. For instance, *Avi et al.* propose a chat-based negotiation framework supporting to some extent using natural language processing and issue by issue negotiation [23]. Mell and Gratch have introduced IAGO framework, which allows a human negotiator to exchange offers, emotions (via emoji), preference statement, and free chat [19]. Jonker *et at.* develop a negotiation support tool, Pocket Negotiator, which aims to help human negotiator by means of some analytics and recommendation [12].

From the point of view of how aforementioned text-based negotiation frameworks function in HRI, it would be more convenient and effective to communicate via speech. Furthermore, as the reciprocal interaction of engagement is an inevitable part of human communication [27], a turn-taking fashion interaction would be appropriate for human-robot negotiation (HRN). Accordingly, this paper introduces a speech-based negotiation protocol, in which a humanoid robot negotiates with a human counterpart via speech by means of speech recognition and text-to-speech methods. A Nao humanoid robot was used to represent a human negotiator the action it should take was commanded remotely in [5]. In our case, our robot makes all decisions by itself in multi-issue negotiation. In that sense, this is the first humanoid robot agent negotiating with human autonomously.

To date, a variety of negotiation strategies have been developed for automated negotiation [29,13,4]. In this work, we propose a novel negotiation protocol using a set of arguments. We present a variant of a time-based bidding tactic, which changes its behavior stochastically between *Conceder* and *Boulware* [7]. A carefully designed user experiment has been conducted to study the performance of the proposed negotiation strategy in human robot negotiation.

The rest of the paper is organized as: Section 2 explains the proposed negotiation protocol and strategy elaborately while Section 3 presents our experimental design and empirical evaluations of our findings. A list of related work is given in Section 4Lastly, we summarize our contributions and discuss future work in Section 5.

2 Proposed Negotiation Framework

In the proposed framework, a humanoid robot Nao namely *Jennifer*[3] negotiates with a human negotiator to come up with a mutual agreement. In the following part, proposed speech-based negotiation protocol governing the interaction between Jennifer and a human negotiator, is explained.

2.1 Speech-based Negotiation Protocol

Communication medium (e.g. *speech*, *vision* or *text*) plays a crucial role in human communication [30]. “How we say” is just as important as “what we say”. Therefore, it is important to select the best way of communication while designing a negotiation protocol for HRN. Most existing work on human-agent negotiation have used a text-based communication; however, human-human negotiations are mostly carried out through speech. In addition to this, establishing a communication with a human through speech rather than using a text-based method is a more natural and fluent way. Therefore, we propose a **S**peech-based **H**uman-**R**obot **N**egotiation **P**rotocol, called **SHRNP**, which is a variant of Alternating Offers Protocol [24].

Figure 1 shows the FIPA representation of this protocol. According to this protocol, Jennifer initiates the negotiation by asking whether the human negotiator is ready to make an offer (Turn 1). Human negotiator should tell Jennifer that she/he is ready to make her/his offer (Turn 2). Note that she or he should say ready at this stage (e.g. “I

[3] In this paper, “Jennifer” is used to refer our humanoid robot.

am ready"). When Jennifer recognizes the word "ready", she says that she is ready to hear her/his offer (Turn 3). Then, the human negotiator should tell her/his offer (Turn 4). Jennifer can accept this offer or make a counter offer (Turn 5). If Jennifer makes a counter offer, she asks whether the human negotiator accepts this offer. The human negotiator should say "Yes" to accept this offer; otherwise, she/he should say "No" to continue negotiation with another offer (Turn 6). Afterwards, the process continues in a turn-taking fashion (Turn 1–6) until having an agreement or reaching a predefined deadline.

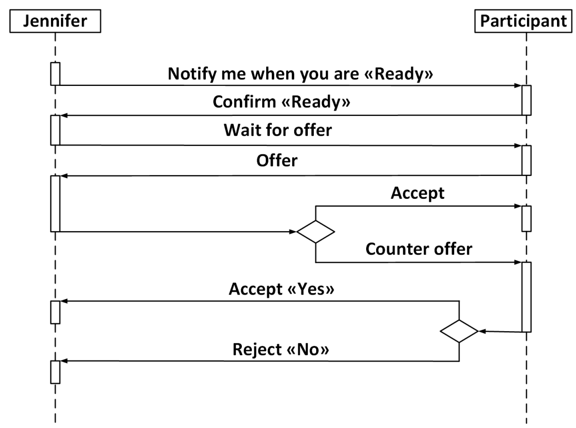

Fig. 1. FIPA Representation of Speech-based Protocol

The underlying protocol is flexible enough for enabling agents to perform different types of bilateral negotiation. While they may negotiate on their holiday (e.g. location, duration, activities), they can also negotiate how to allocate a set of resources between them. The proposed speech-based protocol consists of four fundamental phases as follows:

– **Notification Phase (Turn 1–3):** Human negotiators may sometimes think out loud and what they said can be perceived as an offer by Jennifer. To avoid such misunderstanding, SHRNP ensures when exactly the human negotiator makes her/his offer by confirming when they are ready to make their offers so that Jennifer can process the right utterances to gather her opponent's offer.
– **Offering Phase for Human Agent (Turn 4):** Human negotiator makes his/her offer.
– **Robot-Response Phase (Turn 5)**: Jennifer evaluates her opponent's offer and either accepts it or makes a counter offer.
– **Human-Response Phase (Turn 6):** If Jennifer makes a counter offer, the human agent should inform whether she/he accepts/rejects the given offer.

Jennifer uses speech recognition to perceive what the other party says, and text-to-speech technology to speak to the human negotiator. For the fluidity of the conversation, our framework aims not to restrict the user with a set of predefined words. Therefore, Jennifer uses dictation instead of grammar-based speech recognition. That is, Jennifer listens to the human negotiator until she or he stops her/his speech. With the help of the speech recognition tool, Jennifer translates her opponent's speech into a set of words. Afterwards, she processes the recognized words regarding the given phase. In notification and human-response phases, our agent tries to find a predefined keyword (e.g. "ready", "yes" or "no") in the given set and ignores other words. In the offering phase for human agent, the opponent may say her/his offer in a different way (e.g. use a different order). Therefore, our agent should process the given set of recognized words and convert them into a valid offer. In the following part, we introduce a negotiation strategy using a variant of time-dependent bidding tactic.

2.2 Time-dependent Stochastic Bidding Tactic

Time-dependent concession strategies have been widely used in automated negotiation [7]. When a negotiating agent employs such a concession strategy, its behavior changes with respect to the remaining time. That is, the agent has a tendency to demand more at the beginning and to concede over time. The target utility at a given time is estimated according to a time-dependent function and a bid that has a utility close to the estimated target utility is offered by the agent.

The proposed time-dependent stochastic bidding tactic (TSBT) defines time-dependent lower and upper bounds and randomly generates a bid between them. To estimate their values, we adopt to use a time-dependent concession function proposed by [28]. Equation 1 represents the adopted concession function where t denotes the scaled time $t \in [0,1]$ and P_0, P_1, P_2 are the maximum value, the curvature of the curve, and minimum value respectively [4].

$$TU(t) = (1-t)^2 \times P_0 + 2 \times (1-t) \times t \times P_1 + t^2 \times P_2 \tag{1}$$

It is worth noting that the adaptive lower and upper bounds correspond to *Conceder* and *Boulware* behavior respectively. Recall that Conceder agent concedes fast during the negotiation while Boulware agent hardly concedes until the deadline. As seen on Figure 2, our agent may switch its strategy between these tactics stochastically. Consequently, the human opponent may not easily guess our agent's behavior.

2.3 Negotiation Strategy

Algorithm 1 shows how Jennifer makes her decisions during the negotiation. Jennifer checks whether the deadline is reached; if so, she ends the negotiation (Line 1–2). Otherwise, she generates her offer according to her bidding tactic. (Line 3). If the utility of the opponent's current bid is higher than or equal to the utility of Jennifer's incoming offer, she accepts the given offer (Line 4–5). Otherwise, Jennifer makes her counter offer (Line 6).

[4] For the lower bound, P0, P1, P2 are 0.94, 0.5, 0.4 respectively and for the upper bound they are 1, 0.9, and 0.7 respectively in our experiments as seen in Figure 1

Algorithm 1: Jennifer's Decision Module

Data: $T_{deadline}$: deadline, t_{cur}: the current time,
T_{Θ}: warning time for deadline, $tactic$: Jennifer's bidding tactic
$O_h^{t_{cur}}$: human opponent's current offer,
$O_h^{t_{prev}}$: human opponent's previous offer,
$O_j^{t_{cur}}$: Jennifer's counter offer, R: reservation utility,
$U(O_h^t)$ the utility of the human opponent's offer at time t for Jennifer,

if $t_{cur} >= T_{deadline}$ **then**
    Behavior $\leftarrow$ End-Negotiation ;
else
    $O_j^{t_{cur}} \leftarrow$ generateBid($tactic$);
    if $U(O_j^{t_{cur}}) \leq U(O_h^{t_{cur}})$ **then**
        Behavior $\leftarrow$ Accept ;
    else
        Make $O_j^{t_{cur}}$;
        if $O_h^{t_{prev}}$ == *null* **then**
            isHurryUp $\leftarrow$ false ;
        else
            if *isHurryUp = false &* $T_{\Theta} <= t_{cur}$ **then**
                Behavior $\leftarrow$ Hurry-up ;
                isHurryUp $\leftarrow$ true ;
            else
                if $U(O_h^{t_{cur}}) < R$ **then**
                    Behavior $\leftarrow$ Offended ;
                else
                    $\Delta U \leftarrow U(O_h^{t_{cur}}) - U(O_h^{t_{prev}})$;
                    Behavior $\leftarrow$ getMood(ΔU,$O_h^{t_{cur}}$) ;
                end
            end
        end
    end
end

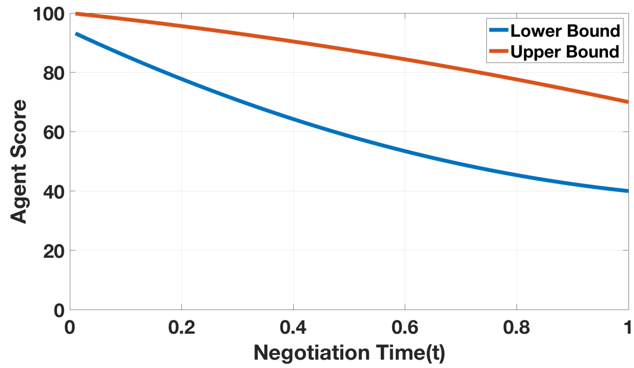

Fig. 2. Time-dependent Lower and Upper Bounds

Afterwards, Jennifer decides her attitude towards her opponent. Jennifer pushes her opponent as the deadline approaches (Line 9–11). There is a predefined time, T_Θ to warn the opponent once. If it is reached, Jennifer tells the opponent to hurry up as specified in Table 1. If the opponents make a humiliating offer, which is not acceptable at all, Jennifer feels "offended" (Line 12–13). That is when the utility of the given offer is less than the reservation utility – the minimum acceptable utility.

Otherwise, Jennifer calculates the utility change in her opponent's subsequent offers (Line 14) and decides her attitude considering the given offer and her opponent's move (e.g. concession, silent, selfish). As specified in Table 1, a mild behavior is adopted by Jennifer if Jennifer thinks that they are approaching a consensus. That happens when Jennifer employs TSBT and the utility of her opponent's offer for Jennifer, $U(O_h^{t_{cur}})$ is higher than or equal to the estimated lower threshold, LT.

When the opponent does not change the utility of its offer, it corresponds to a neutral behavior. If the opponent concedes ($\Delta U > 0$), Jennifer feels pleasant. On the other hand, if the opponent makes a selfish move (e.g. decreasing Jennifer's utility), she shows her dissatisfaction by saying that she did not like her opponent's offer. Table 1 indicates what Jennifer says to her human opponent in each case.

3 Evaluation

In order to evaluate the performance of the developed negotiation strategy with TSBT, we design a user experiment. Determining the basic design structure is a crucial task especially in HRN setting. In the following part, our experiment design and our findings are given.

3.1 Experimental Setup

We have recruited 30 participants (i.e. university students and faculty members; 19 Male, 11 Female; Median age: 23) for conducting our human-robot experiments. Our

Table 1. Argument Decision Matrix

Case	Behavior	Arguments
$U(O_h^{t_{cur}}) < R$	Offended	It is not acceptable!
$\Delta U < 0$	Dissatisfied	I dont like your offer. You should revise it.
$\Delta U = 0$	Neutral	Himm
$\Delta U > 0$	Pleasant	It is getting better but not enough.
$U(O_h^{t_{cur}}) >= LT$	Mild	I like your offer but you can increase a little bit.
$U(O_h^{t_{cur}}) >= U(O_j^{t_{cur}})$	Acceptance	Yes, I accept your offer!
$T_\Theta <= t_{cur}$	Hurry up	Hurry up! We need to find an agreement soon
$t_{cur} >= T_{deadline}$	Time's up	Let's stop ! We cannot reach an agreement.

main aim is to investigate whether the robot can perform at least as well as human negotiators in the given negotiation task. Therefore, we asked the volunteer participants to negotiate with Jennifer and then analyzed the negotiation outcomes elaborately.

In the experiment, a negotiation scenario is given to each participant and as a role-playing game, they are asked to study their preference profiles and the interaction protocol elaborately before their negotiation. Apart from the given negotiation scenario, an easy negotiation scenario has been created for the training session. The participants watched a video of a training negotiation session; afterwards, they do a five-minute negotiation training session.

After the training session, the preference profile for the negotiation session is studied by the participant and then they negotiate with Jennifer for up to 10 minutes. The deadline for each negotiation is set as 10 minutes. If there is no agreement within 10 minutes, both parties receive zero points. Note that the aim of the participants is to receive at least 30 points out of 100. The participants are encouraged to maximize their score by pointing out that the participants with the highest score will win a gift card from a well-known coffee brand. Thus, the participants take their negotiation seriously.

According to our scenario, our participants need to negotiate with Jennifer on resource allocation in order to survive in a deserted island. There are eight indivisible items: some of them will be given to the participant and the rest of them will be taken by Jennifer. Note that human participants ask for what they would like to get and Jennifer offers what items to be given to the participants in order to avoid misunderstanding. In other words, the negotiation is on what items would be given to the participant. Table 2 shows these items and their scores for Jennifer and her human counterpart. Figure 3 shows the utilities of each possible bid in the given scenario as well as the agreement zone. It can be seen that the bargaining power of two parties are almost the same.

It is worth noting that the participants only know their own scores and they are informed that Jennifer does not know their scores too.

As seen in Figure 4, participants are allowed to use a paper to take their notes and their phones to check the remaining time. They keep current preference profile and

Table 2. Preference Profiles for Negotiation Sessions

Items	**Jennifer's Profile**	**Human's Profile**
Hammer	6	13
Container	22	20
Knife	5	10
Match	20	22
Compass	13	5
Medicine	7	6
Food	17	7
Rope	10	17

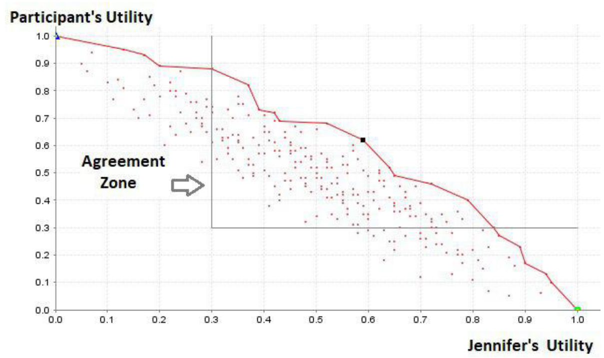

Fig. 3. Outcome Space for Negotiating Parties

Fig. 4. Experiment Setup

interaction flowchart with them during their negotiation. Their negotiation session is recorded so as to check the quality of speech recognition with our detailed log files. At the end of their negotiation, each participant is asked to fill in a questionnaire form about their negotiation experience with Jennifer.

3.2 Experiment Results

Out of 30 negotiations, 26 negotiations ended up with an agreement while only 4 of them failed. In other words, participants reached an agreement in 86.7 percent of negotiations. Table 3 shows the detailed results of each successful negotiation in our experiment. First column shows negotiation session id and the following five columns indicate the percentages of human negotiator's attitude perceived by the robot as it is described in Table 1. For instance, "offended" indicates the percentage of offensive offers made by the human negotiator. The seventh and eighth column show the score received by our agent (Jennifer) and the score gained by the human participant respectively. The last column shows the normalized agreement time [0,1].

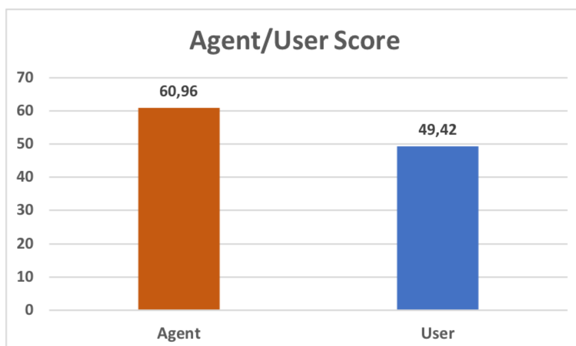

Fig. 5. Average Scores in Agreements

Jennifer beats the human participants in approximately 58 percent of the negotiations (15 out of 26). Furthermore, we test the following null-hypothesis, H_0**:** There is no difference between the gained score by human negotiator and the score collected by the robot. We applied t-tailed test for two paired samples on agent score and user score. Since it is observed that t=2.937 ¿ tc=1.708 and p=0.0035¡0.05, it is concluded that the null hypothesis is rejected. In other words, the scores for the agent and human participants are significantly statistically different at 95 percent confidence interval. That is, we can conclude that Jennifer outperformed human negotiators on average (60.96 vs 49.42) as also seen in Figure 5.

As far as the agreement time is concerned, it is seen that it took about 6-7 minutes to find an agreement with Jennifer. Due to the nature of our agent's bidding strategy, the human participants who were patient to wait longer, gained higher score. Recall that Jennifer concedes stochastically over time.

When the behavior of best performing human negotiator is studied, it is observed that 75 percent of his/her moves is dissatisfying ($\Delta U < 0$). On the contrary, the worst performing human negotiator made mostly pleasant and mild behavior (37.5 % and 50 % respectively).

Table 3. Analysis of Successful Negotiations

ID	Offended	Dissatisfied	Neutral	Pleasant	Mild	Agent Score	User Score	Agreement Time
1	0.0	7.5	0.0	37.5	25.0	79	40	0.757
2	0.0	50.0	0.0	50.0	0.0	77	36	0.505
3	0.0	66.7	0.0	33.3	0.0	75	32	0.272
4	0.0	0.0	0.0	0.0	100.0	75	32	0.473
5	0.0	12.5	0.0	37.5	50.0	72	30	0.846
6	0.0	100.0	0.0	0.0	0.0	69	34	0.492
7	0.0	25.0	0.0	50.0	25.0	68	37	0.466
8	0.0	40.0	40.0	20.0	0.0	65	49	0.478
9	0.0	20.0	0.0	40.0	40.0	65	49	0.644
10	0.0	0.0	0.0	66.7	33.3	65	49	0.743
11	0.0	20.0	20.0	40.0	20.0	63	47	0.545
12	37.5	12.5	12.5	25.0	12.5	62	50	0.688
13	0.0	44.4	0.0	44.4	11.1	60	49	0.618
14	14.3	28.6	14.3	42.9	0.0	59	62	0.693
15	0.0	60.0	0.0	20.0	20.0	59	62	0.553
16	66.7	0.0	0.0	22.2	11.1	59	62	0.702
17	0.0	50.0	0.0	50.0	0.0	57	58	0.556
18	0.0	44.4	33.3	22.2	0.0	57	44	0.647
19	0.0	11.1	22.2	44.4	22.2	57	58	0.749
20	0.0	45.5	9.1	27.3	18.2	55	43	0.724
21	0.0	75.0	0.0	25.0	0.0	50	66	0.755
22	0.0	62.5	0.0	37.5	0.0	48	59	0.807
23	0.0	71.4	0.0	28.6	0.0	48	59	0.845
24	9.1	54.5	0.0	18.2	18.2	47	59	0.982
25	76.9	0.0	0.0	15.4	7.7	47	65	0.924
26	23.1	23.1	0.0	53.8	0.0	47	54	0.838
MEAN	8.75	36.72	5.82	32.77	15.94	60.96	49.42	0.665

Figure 6 shows the average ratings given by the users to our questionnaire consisting of 9-point scaled questions after their negotiation (1 for strongly disagreement whereas 9 for strongly agreement). The resulting average ratings of the more positively-structured statements such as “Her gestures were mostly consistent with the situation.” were satisfyingly high. Besides, some of the negatively-structured statements such as “I was frustrated with Jennifers attitude.”, had reasonably low scores (i.e. disagree) as positive feedback.

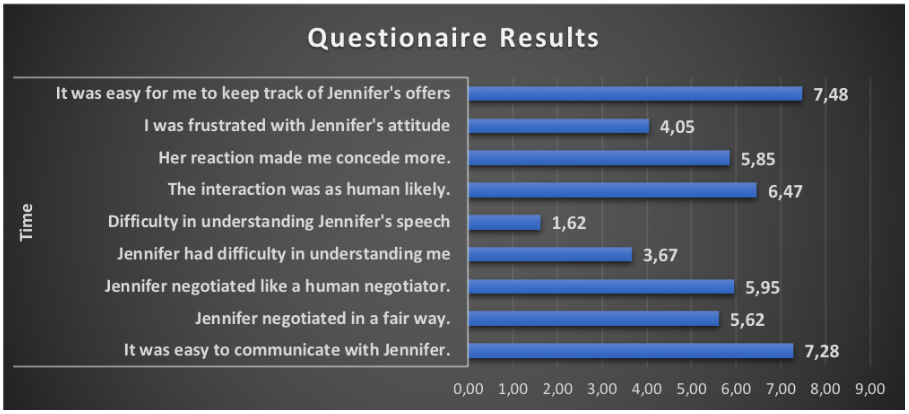

Fig. 6. Questionnaire

4 Related Work

In this section, we review most recent works on human-agent negotiation. Designing agents that are able to negotiate with humans requires considering some human factors and dynamics[16]. For example, the fairness of the offers might have a significant influence on human negotiator's decision making.

Most of the agent-based negotiation systems (e.g.,[19,23,12]) use text-based communication. However, communication medium plays a key role in human negotiations and mostly verbal communication is preferred. It is more natural and effective to communicate via speech in human-agent negotiation. Recently, DeVault, Mell, and Gratch worked on establishing a fluent conversation between a virtual agent and a human negotiator by using speech libraries collected from human-human negotiations [6]. The agent itself is not fully autonomous; the speech and high-level behavior of the virtual agent is controlled by two experts.

In agent-based negotiation framework, during the negotiation in addition to offers, arguments can be exchanged to persuade the other party [22,21]. In recent years, negotiation frameworks support to exchange arguments [19,23]. IAGO has been developed for human-agent negotiations in which parties can exchange offers, arguments, and emotional expressions. Note that the agent uses a predefined set of utterances during its negotiation. Moreover, Mell *et al.* investigate whether using the arguments indicating the appreciation to the opponent has an impact on the negotiation. [20]. In another study, by applying neural networks and reinforcement learning on the dialogues collected from human-human negotiations, it is aimed to learn how to use dialogues effectively in negotiation [15].

Although there are a variety of works on human-virtual agent negotiation, there are relatively less work on negotiating robots. Bevan and Fraser studied experimentally whether or not handshaking before the negotiation has a comprising effect in negotiation. [26] examined experimentally whether or not the use of guilty expression by a robot has an effect on its opponent's compromise. In almost all of those works [26], robots are remotely controlled by a human. To best of our knowledge, there is an urgent need to develop fully autonomous negotiating humanoid robots.

5 Conclusion

This work introduces a novel negotiation scheme in which a humanoid robot negotiates effectively with a human counterpart via speech. Our experimental results showed that our robot can negotiate at least as well as human counterpart on average. Even Jennifer managed to outperform the human participants although it employs a time-based concession strategy.

As a future work, we are planning to develop more sophisticated negotiation strategies and compare their performance with the performance of the time-based concession strategy. Moreover, we would like to investigate the effect of other factors such as "gesture" on the negotiation outcome. Furthermore, it would be interesting to study the cultural differences in human-robot negotiation as [10] do.

References

1. B. An, N. Gatti, and V. Lesser. Extending alternating-offers bargaining in one-to-many and many-to-many settings. In *Proc. of the 2009 IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology*, volume 2, pages 423–426, 2009.
2. R. Aydoğan, T. Baarslag, K. V. Hindriks, C. M. Jonker, and P. Yolum. Heuristics for using cp-nets in utility-based negotiation without knowing utilities. *Knowledge and Information Systems*, 45(2):357–388, 2015.
3. R. Aydoğan, D. Festen, K. V. Hindriks, and C. M. Jonker. Alternating offers protocols for multilateral negotiation. In K. Fujita, Q. Bai, T. Ito, M. Zhang, F. Ren, R. Aydoğan, and R. Hadfi, editors, *Modern Approaches to Agent-based Complex Automated Negotiation*, pages 153–167. Springer, 2017.
4. T. Baarslag, E. H. Gerding, R. Aydogan, and M. C. Schraefel. Optimal negotiation decision functions in time-sensitive domains. In *IEEE/WIC/ACM WI-IAT*, pages 190–197, 2015.
5. C. Bevan and D. S. Fraser. Shaking hands and cooperation in tele-present human-robot negotiation. In *HRI*, pages 247–254, 2015.
6. D. DeVault, J. Mell, and J. Gratch. Toward Natural Turn-Taking in a Virtual Human Negotiation Agent. In *AAAI Spring Symposium on Turn-taking and Coordination in Human-Machine Interaction*, pages 2–9, Palo Alto, California, Mar. 2015. AAAI Press.
7. P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
8. S. Fatima, S. Kraus, and M. Wooldridge. *Principles of automated negotiation*. Cambridge University Press, 2014.
9. S. S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.
10. G. Haim, Y. K. Gal, M. Gelfand, and S. Kraus. A cultural sensitive agent for human-computer negotiation. In *AAMAS*, pages 451–458, 2012.
11. M. Johansson, T. Hori, G. Skantze, A. Höthker, and J. Gustafson. Making turn-taking decisions for an active listening robot for memory training. In *ICSR*, volume 9979 of *Lecture Notes in Computer Science*, pages 940–949, 2016.

12. C. M. Jonker, R. Aydoan, T. Baarslag, J. Broekens, C. A. Detweiler, K. V. Hindriks, A. Huldtgren, and W. Pasman. An introduction to the pocket negotiator: A general purpose negotiation support system. In *EUMAS 2016*, pages 13–27. Springer, 2016.
13. S. Kawaguchi, K. Fujita, and T. Ito. Agentk: Compromising strategy based on estimated maximum utility for automated negotiating agents. In *New Trends in Agent-Based Complex Automated Negotiations*, pages 137–144. Springer, 2012.
14. G. E. Kersten and G. Lo. Aspire: an integrated negotiation support system and software agents for e-business negotiation. *International Journal of Internet and Enterprise Management*, 1(3):293–315, 2003.
15. M. Lewis, D. Yarats, Y. Dauphin, D. Parikh, and D. Batra. Deal or no deal? end-to-end learning of negotiation dialogues. In *EMNLP*, pages 2443–2453. Association for Computational Linguistics, 2017.
16. R. Lin and S. Kraus. Can automated agents proficiently negotiate with humans? *Commun. ACM*, 53(1):78–88, 2010.
17. R. Lin, S. Kraus, T. Baarslag, D. Tykhonov, K. Hindriks, and C. M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.
18. I. Marsa-Maestre, M. Klein, C. M. Jonker, and R. Aydoğan. From problems to protocols: Towards a negotiation handbook. *Decision Support Systems*, 60:39 – 54, 2014. Automated Negotiation Technologies and their Applications.
19. J. Mell and J. Gratch. IAGO: interactive arbitration guide online (demonstration). In *AAMAS*, pages 1510–1512, 2016.
20. J. Mell, G. M. Lucas, and J. Gratch. An effective conversation tactic for creating value over repeated negotiations. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, pages 1567–1576, 2015.
21. P. Pasquier, R. Hollands, I. Rahwan, F. Dignum, and L. Sonenberg. An empirical study of interest-based negotiation. *Autonomous Agents and Multi-Agent Systems*, 22(2):249–288, 2011.
22. I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *Knowledge Eng. Review*, 18(4):343–375, 2003.
23. A. Rosenfeld, I. Zuckerman, E. Segal-Halevi, O. Drein, and S. Kraus. Negochat: a chat-based negotiation agent. In *AAMAS*, pages 525–532. IFAAMAS/ACM, 2014.
24. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
25. V. Sanchez-Anguix, R. Aydogan, V. Julian, and C. Jonker. Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Electronic Commerce Research and Applications*, 13(4):243–265, 2014.
26. B. Stoll, C. Edwards, and A. Edwards. why arent you a sassy little thing: The effects of robot-enacted guilt trips on credibility and consensus in a negotiation. *Communication Studies*, 67(5):530–547, 10 2016.
27. A. L. Thomaz and C. Chao. Turn-taking based on information flow for fluent human-robot interaction. *AI Magazine*, 32(4):53–63, 2011.
28. R. Vahidov, G. E. Kersten, and B. Yu. Human-agent negotiations: The impact agents concession schedule and task complexity on agreements. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*, pages 412–420, 2017.
29. C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings. Iamhaggler2011: A gaussian process regression based negotiation agent. In *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pages 209–212. Springer, 2013.

30. J. Xu, J. Broekens, K. H. Hindriks, and M. Neerincx. Mood expression of a robotic teacher on students. In *Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2614–2620, 2014.