Deniz
A Robust Bidding Strategy for Negotiation Support Systems

Jonker, Catholijn M.; Aydoğan, Reyhan

DOI
10.1007/978-981-15-5869-6_3

Publication date
2021

Document Version
Accepted author manuscript

Published in
Advances in Automated Negotiations, ACAN 2018

Deniz: A Robust Bidding Strategy for Negotiation Support Systems

Catholijn M. Jonker[1,2] and Reyhan Aydoğan[1,3]

[1] Interactive Intelligence Group Delft University of Technology Delft, The Netherlands C.M.Jonker@tudelft.nl
[2] Leiden Institute for Advanced Computer Science, Leiden University, c.m.jonker@liacs.leidenuniv.nl
[3] Department of Computer Science, Özyeğin University, Istanbul, Turkey reyhan.aydogan@ozyegin.edu.tr

Abstract. This paper presents the Deniz agent, that has been specifically designed to support human negotiators in their bidding. The design of Deniz is done with the criteria of robustness and the availability of small data, due to a small number of negotiation rounds in mind. Deniz's bidding strategy is based on an existing optimal concession strategy that concedes in relation to the expected duration of the negotiation. This accounts for the small data and small number of rounds. Deniz deploys an adaptive behaviour-based mechanism to make it robust against exploitation. We tested Deniz against typical bidding strategies and against human negotiators. Our evaluation shows that Deniz is robust against exploitation and gains statistically significantly higher utilities than human test subjects, even though it is not designed specifically to get the highest utility against humans.

1 Introduction

Negotiation is part of our daily lives, informally at home, or formally in matters of business. We negotiate to reach a consensus if we have a conflict of interests ([20,12,15,21,7]). While some people are very good at negotiation, others have difficulty in reaching optimal outcomes and mostly end up with suboptimal outcomes ([12,21]). Improving on negotiation outcomes, can be done by training the human before they enter the negotiation, or by supporting them during the negotiation. Artificial Intelligence applications have been developed for both purposes. For example, Conflict Resolution Agent (CRA) is a virtual agent used to let humans train their negotiation skills [10,6,18], and another example can be found in [5].

Decision support systems are meant to support people during their decision making. For negotiation, the Pocket Negotiator is an example ([13]) that aims at helping human negotiators improve their negotiation outcomes by providing guidance throughout the negotiation process, and with a specialisation in bidding support. Another is the Social agent for Advice Provision (SAP), that interacts with the human and attempts to convince her to choose a certain option [2].

In this paper we focus on agents for negotiation support during the bidding phase. For these agents the negotiation strategy plays a key role. The literature on automated negotiations has produced a wealth of strategies that have proved themselves in the Automated Negotiating Agents Competition, for an overview see [14,9,8,1]. Even if when focusing only on bidding with complete bids without exchanging any other information with the opponent, these strategies are not directly transferable to human negotiations. The most obvious difference is that in the automated negotiation competitions the agents work with a deadline of some minutes, which is enough for most agents to exchange thousands of bids, whereas in human negotiations the number of rounds is low, although culture dependent, somewhere between 3 (USA) and 20 (Northern Africa), if you ask the experts. The consequence is that instead of considering big data, in human negotiations we only have small data available to get some inkling about the preferences and the bidding strategy of the opponent. That means that all agents that were developed with elaborate phases for exploring the reaction of the opponent by making random bids are not very suitable for human-human negotiation. Furthermore, we subscribe to the general aim of creating explainable Artificial Intelligence, so the strategy shouldn't be too complex either. This eliminates some more agents. With respect to outcome optimality, for Deniz we chose not only to optimize the utility of Deniz's side in the negotiation, but to optimize outcome in combination with acceptability for human negotiators. For example, in our experience, on average people don't accept bidding advice from a hardheaded agent, as it is too extreme to their taste. This entails that we need to look for a very good concession strategy. Finally, we have the responsibility of creating agents whose strategy is robust against exploitation attempts by the opponent. This means that a straightforward concession strategy is not applicable, but that it should have some aspect of Tit-for-Tat in it. With these criteria in mind we developed Deniz[4], an agent that can be used to support humans in their negotiations.

This paper is structured as follows. Section 2 presents the necessary definitions of bids, utilities and moves, that we need in our description and definition of Deniz. Section 3 defines Deniz, with an emphasis on the explanation of its strategies for determining its next move and a deepening of its concession strategy. The criteria of robustness against exploitation and its suitability for conceding at the right moments are discussed explicitly. The proof of its explainability, is in the reading of this section. Section 4 is devoted to the evaluation of Deniz, in which we focus on general performance when supporting humans, and on the criteria of robustness. The criterion of explainability to the user is left for future work, as it is too much tied into the explainability of the PN framework. We wrap up with our main findings, and conclusions in Section 5.

[4] Deniz is a gender independent name, i.e., it is used for both females and males, and means "sea"

2 Bids, Utilities and Moves

This section presents the notation and definitions for bids, utilities and moves as used in the remainder of this article. Furthermore, we would like to point out that our work is inspired on the work by many in the Automated Negotiating Agents Competition (ANAC), as reported in e.g., [14].

Let N be a set of at least two negotiators. For any negotiator $n \in N$, $o_n \in N$ denotes the opponent of n. If no confusion is possible, we drop the index n for the opponent. In the following our two main negotiators are Deniz agent D, and opponent O.

Let B denote the space of all possible bids and $b_n^i \in B$ denotes the bid made by negotiator n in round i. Let $u_n : B \mapsto [0, 1]$ denote the utility function of negotiator n. Then the utility functions of the negotiators map bids $b \in B$ to the two-dimensional utility space $[0, 1] \times [0, 1]$:

$$\langle u_n(b), u_{o_n}(b) \rangle$$

Note that the utility functions of D and O are denoted by u_D and u_O. It is worth noting that in real negotiations, the opponent profiles are typically not known, but estimated. So, in general, the reader should read u_O to be the estimated utility function. For anaylis sake, we had, of course, the true utility function available.

A move μ_n^i is a tuple (b_n^{i-1}, b_n^i) of two sequential bids of negotiator $n \in N$, made with respect to negotiation round $i > 1$[5].

For all $i > 1$, $n, m \in N$, and moves μ_n^i made by n, we define the following:

move size: $\Delta_m(\mu_n^i) = u_m(b_n^i) - u_m(b_n^{i-1})$, is the size of the move (i.e, difference in utility) according to m. As in bilateral negotiations m can be either n or o_n, we consider both Δ_n and Δ_{o_n} for any move.

silent moves: $\mathit{silent}(\mu_n^i)$ if $|\Delta_n(\mu_n^i)| = 0 \wedge |\Delta_{o_n}(\mu_n^i)| = 0$, which means that agent n made a silent move.

concession moves: $\mathit{concess}(\mu_n^i)$ if $\Delta_n(\mu_n^i) < 0 \wedge \Delta_{o_n}(\mu_n^i) \geq 0$, which means that negotiator n made a concession.

selfish moves: $\mathit{selfish}(\mu_n^i)$ if $\Delta_n(\mu_n^i) > 0 \wedge \Delta_{o_n}(\mu_n^i) \leq 0$ which means that n made a selfish move.

nice moves: $\mathit{nice}(\mu_n^i)$ if $\Delta_n(\mu_n^i) = 0 \wedge \Delta_{o_n}(\mu_n^i) > 0$ which means that n made a nice move, i.e., better for o_n, and for n the utility is the same.

fortune moves: $\mathit{fort}(\mu_n^i)$ if $\Delta_n(\mu_n^i) > 0 \wedge \Delta_{o_n}(\mu_n^i) > 0$ which means that n made a fortune move,i.e., a move which is better for both negotiators.

unfortunate moves: $\mathit{unfort}(\mu_n^i)$ if $\Delta_n(\mu_n^i) \leq 0 \wedge \Delta_{o_n}(\mu_n^i) < 0$, which means that negotiator n made matters worse for both sides.

uncooperative moves: $\mathit{uncoop}(\mu_n^i)$ if $\mathit{unfort}(\mu_n^i)$ or $\mathit{selfish}(\mu_n^i)$.

cooperative moves: $\mathit{coop}(\mu_n^i)$ if it is not uncooperative[6].

[5] Note that in the first round, move is undefined.

[6] Note that *silent* moves are thus considered here to be cooperative moves.

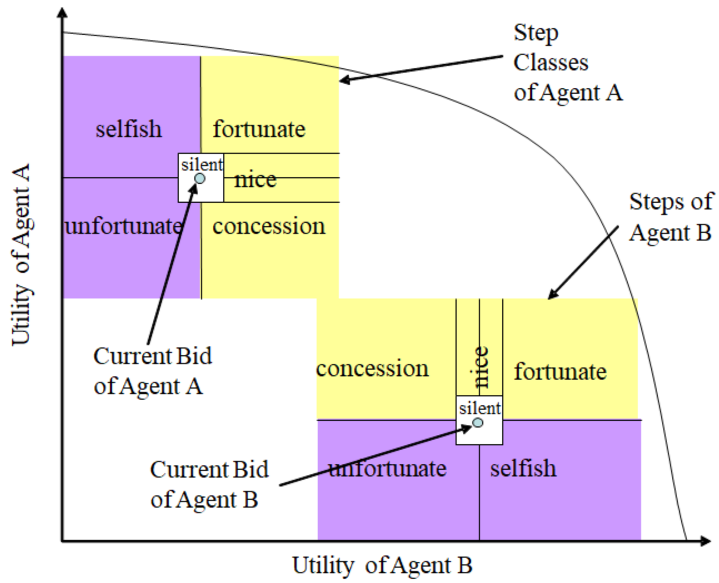

Fig. 1. Moves

These notions are taken from the Dynamics Analysis of Negotiation Strategies (DANS) framework of [11], although the cooperative and uncooperative moves are our additions. Note that in practice it is sometimes useful to take a margin ε around the silent and nice bids as depicted in Figure 1. The details of defining that precisely are left to the reader.

3 Deniz Strategies

The bidding strategy of the Deniz agent is based on the Greedy Concession Algorithm (GCA) as presented in [3], but with a twist. The twist is that a Tit-for-Tat flavor is added as GCA is to easily exploitable by opponents that recognize that GCA is pure concession strategy. Deniz's negotiation strategy includes a bidding and an acceptance strategy, in which Deniz reasons about the end of the negotiation according to two concepts, i.e., estimated number of negotiation rounds, and negotiation deadline.

The estimated number of negotiation rounds is culture-dependent variable which predicts how many rounds negotiations in a culture typically last. For example, in the USA the number is 3, whereas in Northern Africa the number is somewhere around 20, these numbers we got from talking to negotiation experts. The point for Deniz is not how to define that number, but that it indicates to Deniz when the opponent might lose patience and might end the negotiation without an agreement.

Some negotiations have a deadline; for example some types of auctions, or in negotiations about perishable goods, or in negotiations about transportation, where the departure time of the ship that you would like to give your cargo to is a natural deadline. For Deniz this is important, as it should concede before that deadline (in as far as it is willing to concede). Overall Deniz uses a three-phased negotiation strategy:

Initial phase Deniz initially offers a bid with maximal utility for itself.

Mid phase During most of the negotiation (until last phase), Deniz behaves as explained in Section 3.1.

Last phase The last phase starts when there are no rounds left (current round = estimated number of rounds plus 1) or we reached the deadline. Deniz will only make silent bids, and if takes too long to its taste, then it ends the negotiation without an agreement. Too long to its taste is determined by either the deadline, or with a probability that increases with the rounds.

Deniz's acceptance strategy is a thresholded ACNext [4], which means that Deniz will accept an offer if the utility of that offer is greater than or equal to the threshold and the utility of its next bid, if it would make a next bid.

3.1 Making a Next Move

If Deniz decides to make a countermove, then which countermove is decided according to the algorithm defined in this section. In the following we use M to denote the set of move types $\{concess, silent, fort, unfort, nice, selfish\}$, and $(\sigma_i)_{i\in\mathbb{N}}$, with $\sigma_i \in M$ to refer to a sequence of opponent move types in round $i > 1$, e.g., $(\sigma_{i-2}, \sigma_{i-1}, \sigma_i)$. Note that later moves are to the right of earlier moves. Let O^r denote the move type sequence of the last r moves by the Opponent. Then the next move of Deniz is determined by:

$$next\_move = \begin{cases} same & \text{if } O^3 = (uncoop, uncoop, selfish) \\ concess & \text{if } O^3 = (coop, uncoop, selfish) \\ silent & \text{if } O^2 = (coop, selfish) \\ concede\_or\_project & \text{if } O^1 = (unfort) \\ concess & \text{otherwise} \end{cases}$$

Note that Deniz is robust against exploitation in human negotiations, as the opponent needs to cooperate every now and then for Deniz to concede. Concessions are done according to the algorithm in Section 3.2. The *concede_or_project* procedure is presented in Algorithm 1. To project the opponent's bid to the Pareto Frontier Deniz searches for the bid b on the Pareto Frontier, see also Figure 2, such that $u_{o_n}(b)$ is as close as possible as to $u_{o_n}(b_t^{o_n})$. Here $b_t^{o_n}$ refers to the last bid made by the opponent, i.e„ the bid made in round $t = i-1$ if Deniz started the negotiation, and $t = i$ if the opponent started the negotiation[7].

[7] If negotiator n started the negotiation, then in every round i, n is the first to bid, and o_n is the last. So, when in round i, and referring to the last bid made the opponent, it can well be that that bid was made during the previous round.

Algorithm 1: Concede or Project to Pareto Procedure

Data: s_c is the current negotiation state, including b_O, which is the opponent's last offer

Result: b is the bid to be offered by our agent

1 $b_c \leftarrow Deniz\_GCA(s_c)$;

2 $b_p \leftarrow ProjectPareto(b_O)$;

3 $b = \text{argmax}\{u_n(b) \mid b \in \{b_c, b_p\}\}$;

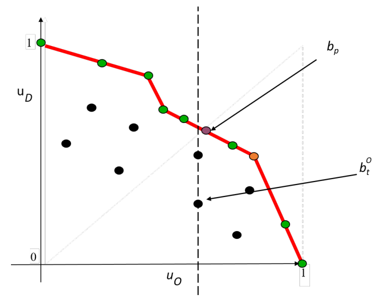

Fig. 2. Projection of the Opponent bid to the Pareto Frontier

3.2 Concession strategy

Deniz first applies GCA, a bidding strategy that determines the optimal concession for a given time t in two steps[3]. Firstly, $f^t(t)$ returns the optimal target conceding utility for that time t. Secondly, it finds a bid with the target utility on the Estimated Pareto Frontier. This gives one bid that is equal to the previous bid or a concession. Deniz improves on this approach by considering this optimal bid as well as nearby bids that may better fit the opponent's preferences. Deniz's next bid determination procedure is defined in Algorithm 2. Note that by explicitly responding to the opponent's last bid, Deniz is a behaviour-based variant of GCA. The algorithm uses the following notation:

b, $b' \in B$ are variables over the bidspace B, $d \in B$ denotes Deniz's previous bid, and $o \in B$ denotes the last bid by opponent O

$s, s_c \in S$: The set of negotiation states S, with variables s, and s_c, where s_c denotes the current negotiation state. A negotiation state refers to the bid history (of bids made so far, annotated by who made which bid), the deadline, current time, and the discount factor (if a discount factor is used)

$f^t : S \mapsto [0, 1]$: GCA's function to determine the target utility, on the basis of a negotiation state, see [3]

E: the Estimated Pareto Front, i.e., the Pareto Front corresponding to u_D and u_O. As u_O is typically estimated, also E is an estimation.

$h : BxB \mapsto \mathbb{N}$: the hamming distance between bid b and b', defined as the number of bid issues that have different values in b and b'.

$\mathcal{R} : \mathcal{P}(B) \mapsto B$ is a function that randomly picks an element from a set, here applied to B.

Algorithm 2: Deniz GCA Concession Algorithm

Data: s_c is the current negotiation state
Result: b_c is the bid to be offered by our agent

`/* Determine the target utility for the optimal bid */`
$tu \leftarrow f^t(s_c)$;
`/*` Determine B_{tu} as the bids on the Estimated Pareto Frontier that are closest to the optimal target utility for Deniz. Note that in discrete domains, more that one bid in E might exist that satisfies the constraint `*/`
$B_{tu} \leftarrow \{b \in E | \operatorname{argmin}_b |u_D(b) - tu|\}$;
`/*` tu' is the adapted target utility `*/`
$tu' \leftarrow u_D(b)$, where $b \in B_{tu}$;
`/*` Create the set C of all possible concessions between the target utility and Deniz's previous bid d. As a concession the bids should have an opponent utility that is higher or equal than that of Deniz's previous bid `*/`
$C \leftarrow \{b \in B | tu' \leq u_D(b) < u_D(d) \wedge u_O(d) \leq u_O(b)\}$;
`/*` If C is empty, then we fall back to bids that have Deniz's previous bid utility. This set is non-empty, as it always contains Deniz's previous bid d `*/`
if $C = \emptyset$ **then**
  $C \leftarrow \{b \in B | u_D(b) = u_D(d)\}$;
end
`/*` Return a randomly chosen bid from the bids in C that have the smallest Hamming distance to the opponent's last bid `*/`
$b_c \leftarrow \mathcal{R}(\{b \in C | \operatorname{argmin}_b h(b, o)\})$;

4 Evaluation

We evaluated Deniz by letting it play against itself, a hardheaded opponent, an explorative agent, and against agents playing some concession strategy. Furthermore, we experimented with human participants that we asked to play against Deniz and then asked about their experience. In all cases we used the Job domain, with default settings for the preferences of employer and employee. The number of expected rounds was set to 10 in all experiments. This number is used by Deniz in its bidding strategy (to know when to concede) and in its acceptance strategy (to know when to accept an acceptable bid). Also we gave Deniz a internal (private) deadline of 15 minutes, after which Deniz will end the negotiation without an agreement.

4.1 Job Domain

The job domain is that of an prospective young ICT professionals negotiating with his/her prospective employer about the job conditions.

The profiles for this domain were obtained in 2010 by interviewing young ICT professionals and HR officers from ICT companies [19]. This resulted in the issue descriptions provided in Table 1, in which fte stands for full time equivalent:

issue	value range	Employee Profile	Employer Profile
salary	[€2000 - €4000]	higher is better	lower is better
fte	0.6, 0.8, 1.0	1.0 > 0.8 0.6	1.0 > 0.8 0.6
work from home	0, 1, 2 days	2 > 1 > 0	0 > 1 > 0
lease car	yes no	yes > no	yes < no
permanent contract	yes, no	yes > no	yes < no
career opportunities	low, medium, high	high > medium > low	high < medium < low

Table 1. Issues with value range of the Job domain

Furthermore, we asked all participants about their preferences, and on the basis of that we set up a typical profile for the employees and one for the employer side, and created utility functions for them. For this paper, the exact procedure is not important, nor is the representativeness of the profiles. Note that the domain size is small. To be precise that is 540, if the salary options are done in steps of €500.

We also compared the outcomes of the negotiations to the Nash Product, which is the maximum of the product of the utilities over all possible bids. In general, for a set of utility functions U, and a bid space B the Nash Product $\eta(U, B)$ is defined as:

$$\eta(U, B) = max\{\prod_{u \in U} u(b) | b \in B\}$$

Similarly, we define $\beta(U, B)$ to a Nash Product bid, i.e., a bid in B such that its utilities form a Nash Product with respect to the set of utilities U. If no confusion is expected as U and B are fixed, we simply write η and β. In the Job domain there is only one Nash Product bid, which is described in Table 2.

4.2 Deniz against Other Negotiating Agents

Deniz agent when playing itself gets the result depicted in Figure 3 and elaborated in Table 2. As experimenter we played the Employee role and for every bid asked a recommendation of Deniz, which we then offered without any changes to the opponent (played by a copy of Deniz). The negotiation lasted 10 rounds.

We also negotiated against Deniz ourselves and to test if Deniz takes advantage of conceding opponents, we first played a negotiator that always concedes. As you can see in Figure 4, part a), Deniz also concedes, but follows the Greedy Concession Algorithm, which concedes much slower. To prevent a possible failure of negotiation, we accepted the offer by Deniz in round 10, which was the estimated deadline for the negotiation.

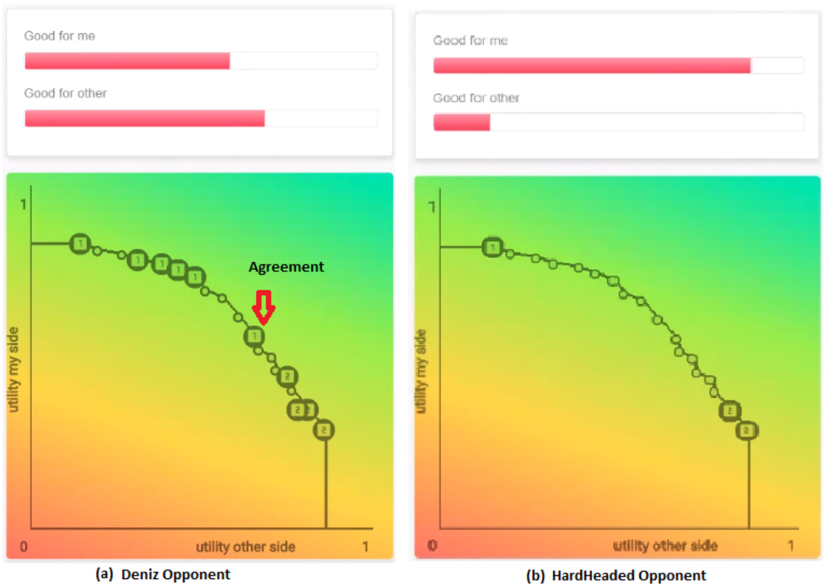

Fig. 3. Negotiation Dance of Deniz vs Deniz (on the left) and Deniz vs HardHeaded (on the right) in the Jobs Domain. Note that no agreement was reached here.

Then we played a very reasonable opponent that starts its bidding close to the Nash Product and is trying to reach a fair outcome, see part b) of Figure 4. In our offers we stayed close to what we considered fair and waited for Deniz to concede to us. The bids we offered had the following utility pattern with respect to our utility function was: 0.7, 0.72, 0.64, 0.7, 0.58, 0.64, 0.72, 0.58, 0.54, 0.58 and Deniz accepted our last offer (in round 10).

Finally, we played a somewhat explorative strategy, see part c) of Figure 4. In that negotiation we deliberately sometimes offered bids that are below the Pareto Optimal Frontier and also sometimes conceded and then went back up again. As can be seen, Deniz is not confused by this and slowly concedes after its opponent makes a concession. Furthermore, it also shows that Deniz agent will accept a reasonable offer when the estimated deadline is reached.

Note that Deniz can only make bids on the Pareto Optimal frontier if the estimation of the utility function of its opponent is correct. In these experiments Deniz (and both parties) always obtained the correct utility functions from the authors. Note, that the depicted Pareto Frontier is one that is computed by PN.

4.3 Deniz Negotiating against humans

In the previous section, we as authors of Deniz tried to exploit Deniz strategy and found it to be robust against our efforts. We could get some results (0.7 utility) if we did our best, and we had the help of PN overview of the bid space,

issue	Deniz-Deniz	Nash Product bid
salary	3000	4000
fte	1.0	1.0
work from home	0	0
lease car	no	no
permanent contract	yes	yes
career development opportunities	low	low
(Employer utility, Employee utility)	(0.68, 0,58)	(0.58, 0.70)

Table 2. The second column shows the agreement reached when both Employer and Employee were played by Deniz. The third column shows the bid corresponding to the Nash Product.

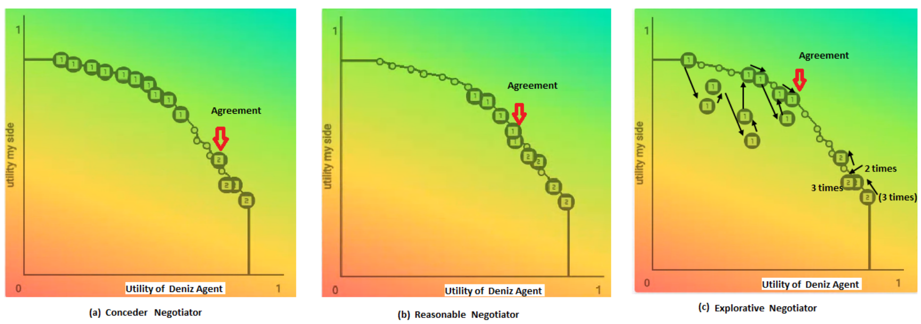

Fig. 4. Negotiation Dance of Deniz against (a) Conceder, (b) Reasonable, and (c) Explorative human negotiator in the Jobs Domain, The agreement configurations and utilities are presented in Table 3.

including the Pareto Optimal Frontier, see Figure 5. In this section we present and discuss the results of an experiment in which humans that were unfamiliar with the Deniz agent negotiate against Deniz. The humans were asked to play against Deniz in the PN framework, which makes it possible to toggle between a supported and an unsupported version. The experiment was approved by TU Delft's Human Research Ethics committee, the data of participants that gave informed consent, is stored according to the Data Management Plan approved by TU Delft's Data Steward.

In PN bidding support interface, the middle section allows users to put in their own choices per issue. The red bars in the right part indicate for a selected point i the bid space, how good that is from the perspectives of the user and the opponent. Note that this is done using the estimated opponent profile that can be constructed in another part of PN. The graph on the right shows an overview of the bid space with the Estimated Pareto Optimal Frontier, and all the bids made so far. Note that the user can click to select points on the Pareto Optimal Frontier. If done so, the selected bid's content is copied to the middle section of interface. Finally, note that at the bottom, the user can ask PN for a bid

issue	**conceder**	**reasonable**	**explorative**
salary	2000	3000	4000
fte	1.0	1.0	1.0
work from home	0	0	0
lease car	no	no	no
permanent contract	yes	yes	yes
career development opportunities	low	low	medium
(u_D, u_O)	(0.78, 0.46)	(0.68, 0,58)	(0.53, 0.72)

Table 3. Agreements of the negotiations of Deniz against different opponents

suggestion, and can also accept a bid from the opponent or walk away without an agreement by ending the negotiation.

In the unsupported version of PN, the users only have the middle section available, where they can enter their offers. The red bars, the graph, and the button to ask for a suggestion are not available.

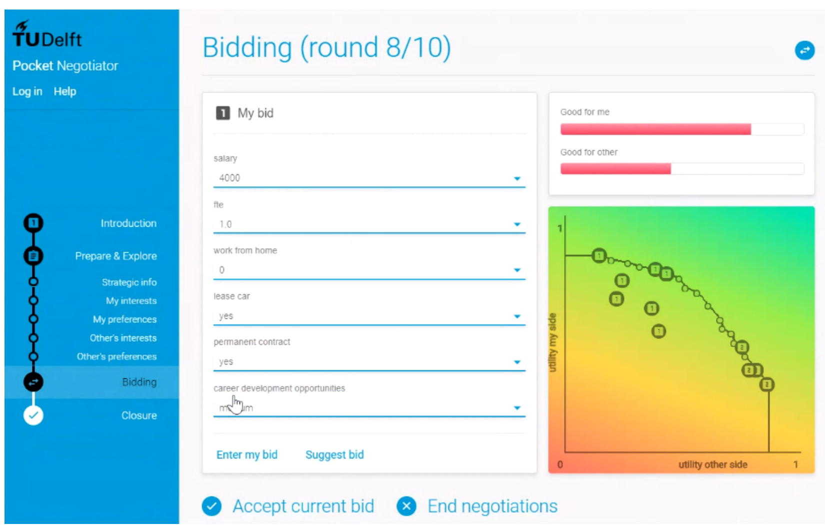

Fig. 5. Pocket Negotiator Bidding Interface.

For our experiment we gathered 78 participants from three classes of students. The first group consisted of Computer Science students, the second group of Industrial Engineering students. These groups of students studied at Özyeğin University (Turkey). The third group consists of business students at Erasmus University (The Netherlands). We did not inform the participants that Deniz agent was also their supporting agent in the condition that they received support from PN.

The average utility received by the participants when negotiating against Deniz without support from PN was 0.50, whereas with the help from PN they received an average utility of 0.53. In the support condition they had the option of asking for a recommendation on what to bid. The agent that would give them that advice was also Deniz. However, on average the participants only asked for a recommendation 2.3 times, and 29 participants never asked for a recommendation at all. Deniz agent on the other hand, in the no-support condition had an average of 0.69 utility, and in the support condition on average 0.68 utility (Std = 0.09). In all conditions, and also when considering all negotiation sessions together Deniz had a statistically significant higher average utility than the human participants, see Table 4.

	no support		**support**		**all**	
	Human	**Deniz**	**Human**	**Deniz**	**Human**	**Deniz**
average utility	0.50	0.69	0.53	0.68	0.50	0.66
standard deviation	0.09	0.10	0.09	0.09	0.13	0.15
t value	9.50		7.71		12.16	
p value	<.00001		<.00001		<.00001	

Table 4. Deniz against Human negotiators: Statistical Analysis

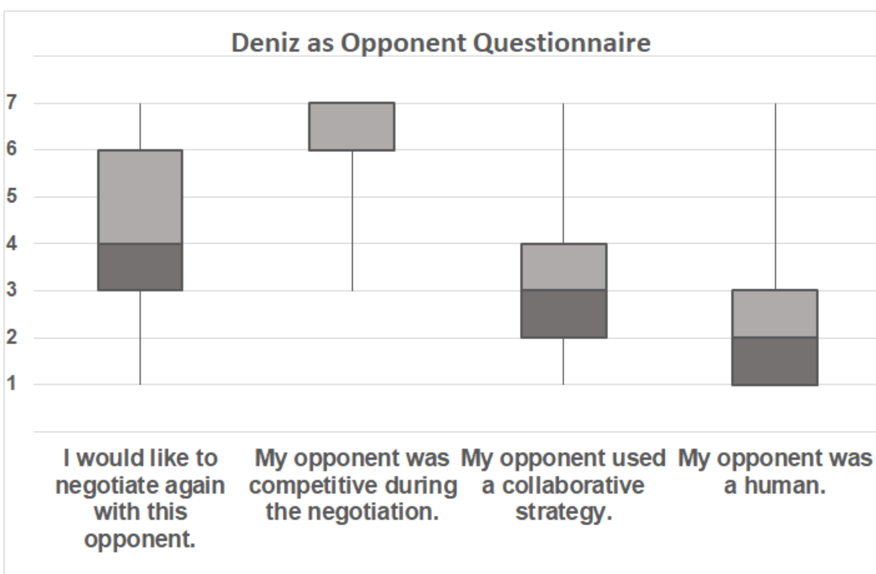

Fig. 6. Participants thoughts about Deniz agent as an opponent

We asked the participant how they felt about Deniz as an opponent through the following statements using a 5 point Likert scale:

1. I would like to negotiate again with this opponent sometime in the future.
2. My opponent was competitive during the negotiation.
3. My opponent used a collaborative strategy.
4. My opponent was a human.

The results are presented in Figure 6. The participants clearly found Deniz to be a competitive agent that did not seem human-like to them. The responses to the statement "I would like to negotiate again with this opponent" gives us more information on the differences between the groups, which was statistically significant p = 0.017; Fratio = 4.4 and dF = 2. The averages per group are as follows Group 1: 4.9; Group 2: 4.9, and Group 3: 3.4. Given that the Deniz agent strategy is a bit like Tit-for-Tat, meaning that non-conceding behaviour would result in no concessions by Deniz, and the fact that Group 3 participant are trained in not giving in too easily, they would find Deniz more tough to play against than the other groups and thus they would be less eager to play Deniz again. This can be seen from the analysis of the types of moves made by the different groups, see Table **??**. In particular, Table **??** shows that in Group 3 the participants (P) make more silent moves than the participants from the other groups. This causes the opponent (Deniz agent, marked by D) to also make more silent moves.

5 Conclusion

The literature on automated negotiation has produced a wealth of automated bidding strategies that have proved themselves in many negotiation domains and against many automated negotiating agents. However, these agents have not been designed for supporting humans in their negotiations. This paper presents a bidding strategy, called Deniz, that is based on the Greedy Concession Algorithm and deploys a Tit-for-Tat strategy. This construction allows Deniz to take the bidding behaviour of its opponent into account. The Greedy Concession Algorithm ensure an optimal concession strategy with respect to the expected number of negotiation rounds. We tested Deniz against typical bidding strategies and against human negotiators. Our evaluation shows that Deniz is robust against exploitation and gains statistically significantly higher utilities than human test subjects. Interestingly, we also observed that human test subjects don't always adhere to the recommendations of their bidding support agent for which we also used Deniz. By rigorously bidding according to Deniz's bidding advice, we obtained higher utility against Deniz than the average participant. Deniz agent should be tested in other negotiation domains and against other automated bidding strategies specifically designed to play against human. In particular, it would be interesting to see how Deniz would hold up in comparison to the agents presented in [17] and [16]. However, the agents do not run in the same platform, so this requires re-implementation of these agents and was therefore left for future work. Furthermore, it would be important to develop a host of bidding support strategies that match with different human negotiation styles.

Further directions of research are to combine our bidding support strategies with other work, such as that of Gratch and co-authors, on training people to become better negotiators.

Acknowledgments

We thank all participants that over the years helped us in our experiments. We also give specific thanks to scientific programmer Wouter Pasman, technician Bart Vastenhouw, Dimitrios Teskouras, professor of the course at Erasmus University that hosted part of our experiments, and our collaborator in many negotiation research projects Tim Baarslag.

References

1. Aydoğan, R., Baarslag, T., Jonker, C.M., Fujita, K., Ito, T., Hadfi, R., Hayakawa, K.: A baseline for non-linear bilateral negotiations: the full results of the agents competing in anac 2014 (2016)
2. Azaria, A., Rabinovich, Z., Kraus, S., Goldman, C.V., Gal, Y.: Strategic advice provision in repeated human-agent interactions. Institute for Advanced Computer Studies University of Maryland **1500**(20742) (2012)
3. Baarslag, T., Gerding, E.H., Aydoğan, R., Schraefel, M.C.: Optimal negotiation decision functions in time-sensitive domains. In: 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT). vol. 2, pp. 190–197 (Dec 2015). https://doi.org/10.1109/WI-IAT.2015.161
4. Baarslag, T., Hindriks, K.V., Jonker, C.M.: Effective acceptance conditions in real-time automated negotiation. Decision Support Systems **60**, 68–77 (Apr 2014). https://doi.org/10.1016/j.dss.2013.05.021, http://dx.doi.org/10.1016/j.dss.2013.05.021
5. Broekens, J., Harbers, M., Brinkman, W.P., Jonker, C.M., Van den Bosch, K., Meyer, J.J.: Virtual reality negotiation training increases negotiation knowledge and skill. In: International Conference on Intelligent Virtual Agents. pp. 218–230. Springer (2012)
6. DeVault, D., Mell, J., Gratch, J.: Toward natural turn-taking in a virtual human negotiation agent. In: 2015 AAAI Spring Symposium Series (2015)
7. Fisher, R., Ury, W.L., Patton, B.: Getting to yes: Negotiating agreement without giving in. Penguin (2011)
8. Fujita, K., Aydoğan, R., Baarslag, T., Hindriks, K., Ito, T., Jonker, C.: The sixth automated negotiating agents competition (anac 2015). In: Modern Approaches to Agent-based Complex Automated Negotiation, pp. 139–151. Springer (2017)
9. Fujita, K., Aydoğan, R., Baarslag, T., Ito, T., Jonker, C.: The fifth automated negotiating agents competition (anac 2014). In: Recent Advances in Agent-based Complex Automated Negotiation, pp. 211–224. Springer (2016)
10. Gratch, J., DeVault, D., Lucas, G.: The Benefits of Virtual Humans for Teaching Negotiation. In: Proceedings of the 16th International Conference on Intelligent Virtual Agents (IVA), 2016. Springer, Los Angeles, CA (Sep 2016)

11. Hindriks, K., Jonker, C.M., Tykhonov, D.: Let's dans! an analytic framework of negotiation dynamics and strategies. Web Intelligence and Agent Systems: An International Journal **9**(4), 319–335 (2011)
12. Howard Raiffa, J.R., Metcalfe, D.: Negotiation Analysis: The Science and Art of Collaborative Decision Making (2002)
13. Jonker, C.M., Aydoğan, R., Baarslag, T., Broekens, J., Detweiler, C.A., Hindriks, K.V., Huldtgren, A., Pasman, W.: An introduction to the pocket negotiator: a general purpose negotiation support system. In: Multi-Agent Systems and Agreement Technologies, pp. 13–27. Springer (2016)
14. Jonker, C.M., Aydoğan, R., Baarslag, T., Fujita, K., Ito, T., Hindriks, K.V.: Automated negotiating agents competition (ANAC). In: AAAI. pp. 5070–5072 (2017)
15. Lewicki, R.J., Saunders, D.M., Barry, B., Minton, J.W.: Essentials of Negotiation. McGraw-Hill, Boston, MA (2003)
16. Lin, R., Kraus, S.: Can automated agents proficiently negotiate with humans? Communications of the ACM **53**(1), 78–88 (2010)
17. Lin, R., Oshrat, Y., Kraus, S.: Automated agents that proficiently negotiate with people: Can we keep people out of the evaluation loop. In: New Trends in Agent-Based Complex Automated Negotiations, pp. 57–80. Springer (2012)
18. Mell, J., Gratch, J., Baarslag, T., Aydoğan, R., Jonker, C.M.: Results of the first annual human-agent league of the automated negotiating agents competition. In: Proceedings of the 18th International Conference on Intelligent Virtual Agents. pp. 23–28. ACM (2018)
19. Pommeranz, A., Visser, W., Broekens, J., Wiggers, P., Hindriks, K., Jonker, C.M.: Duo meta-model for knowledge elicitation and bidding support in nss. In: Proc. of 11th Group Decision and Negotiation Conference. pp. 120–123 (2010)
20. Raiffa, H.: The art and science of negotiation: How to resolve conflicts and get the best out of bargaining. Harvard University Press, Cambridge, MA (1982)
21. Thompson, L.: The Mind and heart of the negotiator. Prentice Hall Press, Upper Saddle River, NJ, USA, 3rd edn. (2000)

A Preference Profiles

Weights of Issues	**Employer's Preferences**	**Employee Preferences**
Salary	0.2	0.24
Fte	0.165	0.32
Work from home	0.245	0.18
Lease car	0.08	0.06
Permanent Contract	0.21	0.16
Career development Opportunity	0.10	0.04

Table 5. Importance of Each Issue for Employer and Employee

Issue		**Employer's Preferences**	**Employee Preferences**
Salary	Value	Evaluation Value	Evaluation Value
	2000	1.0	0.0
	2500	0.75	0.25
	3000	0.5	0.50
	3500	0.25	0.75
	4000	0.0	1.0
Fte	Value	Evaluation Value	Evaluation Value
	0.6	0.36	0.25
	0.8	0.5	0.50
	1.0	0.64	0.75
Work from home	Value	Evaluation Value	Evaluation Value
	0 day	1.0	0.33
	1 day	0.5	0.50
	2 days	0.0	0.67
Lease car	Value	Evaluation Value	Evaluation Value
	Yes	0.0	1.0
	No	1.0	0.0
Permanent Contract	Value	Evaluation Value	Evaluation Value
	Yes	0.24	1.0
	No	0.76	0.0
Career Opportunity	Value	Evaluation Value	Evaluation Value
	Low	1.0	0.0
	Medium	0.5	0.5
	High	0.0	1.0

Table 6. Evaluation Values of Employer and Employee