

How to Recognize and Explain Bidding Strategies in Negotiation Support Systems

Koeman, Vincent J.; Hindriks, Koen; Gratch, Jonathan; Jonker, Catholijn M.

DOI 10.1007/978-981-16-0471-3_3

Publication date 2021 **Document Version** Final published version

Published in Recent Advances in Agent-based Negotiation - Formal Models and Human Aspects

Citation (APA) Koeman, V. J., Hindriks, K., Gratch, J., & Jonker, C. M. (2021). How to Recognize and Explain Bidding Strategies in Negotiation Support Systems. In R. Aydoğan, T. Ito, A. Moustafa, T. Otsuka, & M. Zhang (Eds.), *Recent Advances in Agent-based Negotiation - Formal Models and Human Aspects* (pp. 35-53). (Studies in Computational Intelligence; Vol. 958). Springer. https://doi.org/10.1007/978-981-16-0471-3_3

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

https://www.openaccess.nl/en/you-share-we-take-care

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

How to Recognize and Explain Bidding Strategies in Negotiation Support Systems



Vincent J. Koeman, Koen Hindriks, Jonathan Gratch, and Catholijn M. Jonker

Abstract Effective use of negotiation support systems depends on the systems capability of explaining itself to the user. This paper introduces the notion of an explanation matrix and an aberration detection mechanism for bidding strategies. The aberration detection is a mechanism that detects if one of the negotiating parties deviates from their expected behaviour, i.e. when a bid falls outside the range of expected behaviour for a specific strategy. The explanation matrix is used when to explain which aberrations to the user. The idea is that the user, when understanding the aberration, can take effective action to deal with the aberration. We implemented our aberration detection and our explanation mechanisms in the Pocket Negotiator (PN). We evaluated our work experimentally in a task in which participants are asked to identify their opponent's bidding strategy, under different explanation conditions. As the number of correct guesses increases with explanations, indirectly, these experiments show the effectiveness of our aberration detection mechanism. Our experiments with over 100 participants show that suggesting consistent strategies is more effective than explaining why observed behaviour is inconsistent. An extended abstract of this article can be found in [15].

V. J. Koeman · K. Hindriks

K. Hindriks e-mail: k.v.hindriks@vu.nl

J. Gratch USC Institute for Creative Technologies, 12015 Waterfront Drive, Playa Vista, CA 90094-2536, USA e-mail: gratch@ict.usc.edu

C. M. Jonker (⊠) Interactive Intelligence Group, Delft University of Technology, van Mourik Broekmanweg 6, Delft, The Netherlands e-mail: c.m.jonker@tudelft.nl

LIACS, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands

Vrije Universiteit Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands e-mail: vj.koeman@quicknet.nl

[©] The Author(s), under exclusive license to Springer Nature Singapore Pte Ltd. 2021 R. Aydoğan et al. (eds.), *Recent Advances in Agent-based Negotiation*, Studies in Computational Intelligence 958, https://doi.org/10.1007/978-981-16-0471-3_3

Keywords Negotiation support systems · Recognizing · Explanation · Bidding strategies

1 Introduction

Negotiation support systems aim to assist human negotiators in their complex decision-making processes aimed at reaching an agreement to exchange goods or services. One such system, the Pocket Negotiator (PN) [10], states its goal as 'to enhance the negotiation skills and performance of the user ... through synergy between the human negotiator and the Pocket Negotiator'. The PN supports the major activities of a negotiation: modelling the interests of the user and the opponent, bidding, and closing. Techniques currently used to provide this support include preference elicitation methods [3, 20], visualization of the negotiation space [14], and multi-criteria optimization techniques for advising what to bid and when to accept the opponent's offer [2].

This paper is the first in a line of research to develop a full fledged explanation framework for negotiation support systems. We decided to start with the explanation of what happens during the bidding phase of a negotiation, as the effectiveness of the bidding strategies determines, to a large extent, the utility of the negotiation outcome.

During the bidding phase, the support currently provided by the PN consists of an interface, see Fig. 1 with a range of options and tools. The snapshot is taken at a moment when the user has just received a bid from the opponent bid suggestions, PN provides an intuitive bid analysis in the form of the horizontal red bars in the right upper corner of Fig. 1. That same bid is also presented as a dot in the visualization of the bid space and its Pareto Optimal Frontier.

Expert negotiators use this interface to quickly create bids by either clicking on the points projected on the Pareto Optimal Frontier or by asking for a bid suggestion. Bid suggestions are generated by a bidding support agent. The user can pick any of a number of typical bidding strategies provided by PN. Finally, the visualization provides an overview of the bids made by the user and by the other party. Note that the visualization of the bid space is based on an estimation of the preference profile of the opponent, and the current view of the negotiator of his/her own preferences. If the estimation is wrong, then so is visualization. Furthermore, the bid suggestions, and the advice of when to accept and what to accept of the support agent depend on that estimation.

To get the most out of interface in Fig. 1, the human user and the bidding support agent have to be able to collaborate at a high knowledge level about the ongoings in the negotiation. The team needs to make sophisticated analysis of the bidding by both parties, creating an understanding of why a player makes this particular bid now. Although the reader is referred to the state of the art in human negotiation theories for a thorough discussion of these questions, see, e.g. [17, 21, 22], we provide an example here. Suppose that one of the players, say the opponent, bids below the Pareto Optimal Frontier, then we need to know why. Several reasons come to mind:

TUDelft Pocket Negotiator Log in Help	My bid salary 4000	Good for me Good for other	
 Introduction Prepare & Explore Strategic info My interests My preferences Other's interests Other's preferences 	fte 1.0 work from home 0 kease car yes permanent contract yes career development opportunities how	I De boord de contraction de la contraction de l	
e Bidding	Enter my bid Suggest bid	0 utility other side 1 S End negotiations	

Fig. 1 An example of a bidding phase in the pocket negotiator

- 1. The opponent doesn't realize his bid is not Pareto Optimal bids. This might be the case if the opponent is human, as humans in multi-issue negotiations often find this difficult.
- 2. The opponent's preferences might differ from what we estimated.
- 3. It might be a tactic to play a bit unpredictable.
- 4. All of the above might hold at the same time.

The different cases ask for different actions on our side, and thus we need to identify which case holds.

Similarly, users might deviate intentionally or unintentionally from their chosen bidding strategy. If intentionally, our negotiation supports agent should know about this, so that it can match its advise and support activities to that strategy. If it happens unintentionally, alerting the user might be the best support to give. We wrote 'seems to deviate', as it might also be the case that the preferences of a user change or are for some other reason different from the preferences entered in the negotiation support system. Again it is important to discover this as quickly as possible, and make sure that the team of human negotiator and negotiation support system have a shared understanding of what is going on.

Now we have come to the core of the problem: a negotiation support system can only discuss these matters with the user, if it can explain to the user what we wrote here. Furthermore, the need to discuss this can only be established if the agent is capable of detecting and analysing these and other strange behaviours that we decided to call aberrations. This motivates the need for *aberration detection* and *explanation mechanisms* that we introduce in this paper, and for which we present experimental results.

In Sect. 2, we discuss the state-of-the-art literature relevant for this work. Section 3 discusses the most characteristic and typical negotiation strategies used in real life. We use these strategies to further focus our research. The concepts already discussed informally in Sect. 3 are formalized and extended in Sect. 4 to form the basis for the analytical framework that forms the core of our aberration detection and explanation mechanisms as presented in Sect. 5 and 6. The experimental setup for the evaluation of our mechanisms is presented in Sect. 7. The experimental results are presented in Sect. 8. Conclusions can be found in Sect. 9.

2 Related Work

Explanations are currently employed in many sub-fields of artificial intelligence, such as justifying autonomous agent behaviour, debugging of machine learning models, explaining medical decision-making, and explaining predictions of classifiers [18]. Reference [8] identifies, however, that allowing users of negotiation support systems to 'trust the system through co-participation, transparency, and proper representation' is still an open challenge. For negotiation agents representing humans specifically, the authors identify that a user's trust and willingness to relinquish control is conditional on a sufficient understanding of the agent's reasoning and consequences of its actions.

Reference [24] focuses on explaining the preferences of a user and his or her opponent in the Pocket Negotiator. The authors propose a mechanism to analyse discrepancies between the system's mental model and the user's (assumed) mental model. However, aside from addressing a different sub-topic within negotiation support than we do, generating the content of the explanations and evaluating their effectiveness are also not addressed in this work.

Reference [19] states that 'artificial agents should be equipped with explanation and argumentation capabilities in order to be able to convince their users of the validity of their recommendations'. Reference [23] identifies seven possible aims of such explanations: transparency, scrutability, trust, effectiveness, persuasiveness, efficiency, and satisfaction. The authors also consider these aim metrics for good qualifications, of which trade-offs are inevitable. The goal of an explanation should thus be carefully considered. Reference [18] argues that explanations in AI should be contrastive, selective, non-probabilistic, and social.

Although most research on 'opponent modelling' in (automated) negotiation focuses on determining the preferences of the opponent [7], in this work, we focus explicitly on determining the (bidding) strategy that an opponent uses. Reference [7] identifies two main approaches: regression analysis and time-series forecasting. Specific implementations are, however, either overly simplistic (e.g. classifying an opponent as 'positive' when its average concession rate exceeds some pre-set amount [16]) or opaque (e.g. using techniques like neural networks and Markov chains). In this work, we aim to devise an approach that balances the level of sophistication with the degree of explainability, focusing on increasing a (novice) human negotiator's understanding of the opponent's strategy rather than determining that strategy as good as possible.

Using a negotiation support system as a training tool for novice negotiators, as [13] do for example, shares similarities with our aim of providing insight into bidding strategies of opponents in those systems, as information about (digital) negotiations is to be conveyed to a novice user in both situations. Current work in the field of training is, however, mainly focused on evaluating the (actions of the) participant itself, e.g. focusing on factors such as making efficient concessions and avoiding early commitment. The explanation mechanism as developed in this paper for opponent strategy recognition could be relevant for negotiation training, but we do not explicitly examine that aspect here.

3 Typical Bidding Strategies

As we study negotiation support systems for human negotiators, the number of rounds in a negotiation is low, with the highest numbers typically found in the markets of Northern Africa, where people enjoy haggling, and thus the process takes much longer than in USA where the number of rounds of bidding is typically no more than 3.

In that light, the essence of human negotiation strategies can be captured by the following four typical strategies as identified in [5]: **Hardheaded** ('Tough negotiator that does not yield easily', i.e. makes mostly silent moves or small concessions), **Conceder** ('Nice negotiator that tends to move towards you', i.e. generally makes concessions), **Tit-for-Tat** ('Somewhat mirrors the moves you are making', i.e. responds with the same type of moves) and **Random** ('Does not follow any of the other strategies'.

Examples of how these bidding strategies are in action are depicted in Fig. 2. Note that in that picture, the user (playing the 'me' role) is playing a strategy that allows her to differentiate between these four typical strategies, and, in particular, between the Conceder and the Tit-for-Tat strategy.

Note that the literature on automated negotiating agents, see, e.g. [4, 6], is full of all kinds of sophisticated bidding strategies. However, the core of most of these strategies is formed by (combinations of) two commonly used negotiation tactics: *time-dependent tactics* and *behaviour-dependent tactics*, in which some aspects of randomness is used to prevent the strategy from becoming too predictable. The Conceder and Hardheaded strategy fall under the time-dependent tactics, and Tit-for-Tat is a behaviour-dependent tactic.

It is quite a challenge to recognize the essence of someone's negotiation strategy during a negotiation of only a few rounds. To be able to develop a mechanism to do so that works independent of the domain of negotiation and independent of the opponent one is playing, we need a way to abstract away from the exact details of bids and offers. The following section presents an abstract framework to do so.



Fig. 2 Typical bidding strategies

4 Bids, Utilities and Moves

This section presents the notation and definitions for bids, utilities and moves as used in the remainder of this article.

 $N = \{H, O\}$ is the set of negotiators, where *H* denotes the human participant, and *O* the opponent. Variable $a \in N$ ranges over the negotiators ('agents').

B denotes the bid space for the negotiation and $B_i^a \in B$ denotes the bid made by agent *a* in round *i*. Let $u_a : B \mapsto [0, 1]$ denote the utility function of agent *a* (i.e. *a*'s preferences), then $m_a \subseteq B = \{b \in B | \forall b' \in B : u_a(b') \leq u_a(b)\}$ is the set of bids that have maximum utility for agent *a*. These are the so-called maximum utility bids for the agent.

 $U_N = [0, 1]^{|N|}$ denotes the multi-dimensional utility space (i.e. all possible bids in a domain) over the negotiators in set N, where |N| denotes the number of elements in N. The bids made by the negotiators are mapped to U_N according to the utility functions of negotiators by function $\upsilon : B \mapsto U_N$ defined by $\upsilon(b) = \langle u_H(b), u_O(b) \rangle$. A bid sequence $\beta_{i,i}^a = (B_i^a, \dots, B_i^a)$ is the sequence of bids made by an agent $a \in N$ from round *i* up to and including round *j*, where $i \le j$, in the negotiation.¹ A move μ is a pair (b, b') of two sequential bids $b, b' \in B$ made by the same agent. Any negotiating party can make offers that from the perspective of an agent *a* can be seen as concessions, selfish moves or silent moves. For any two agents $a, a' \in N$ and any move $\mu = (b, b') \in B \times B$, we define the following:

- [move size:] $\sigma_a(\mu) = u_a(b') u_a(b)$, is the size of the move (i.e. difference in utility) according to *a*.
- [silent moves:] $sil_a^{\delta}(a', \mu)$ if $|\sigma_a(\mu)| \le \delta$, which means that agent *a* considers the move with a size less than δ made by agent *a'* to be a silent move.
- [concession moves:] We differentiate between
 - $conc_a^{\delta}(a', \mu)$ if $a \neq a' \land \sigma_a(\mu) > \delta$, which means that from the perspective of *a* agent *a'* conceded at least δ . The generic case $conc_a$ refers to $conc_a^0$. • $conc_a^{\delta}(a, \mu)$ if $\sigma_a(\mu) < 0 \land |\sigma_a(\mu)| \ge \delta$, which means that *a* conceded at least δ .
- [selfish moves:] $self_a^{\delta}(a, \mu)$ if $\sigma_a(\mu) > \delta$, which means that *a* thinks he made a selfish move of at least δ size.
- [move types:] the parametrized characterizing relations defined above also define sets M_a of move types according to agent $a: M_a = \{sil_a^{\delta_1}, conc_a^{\delta_2}, self_a^{\delta_3}\}$, for given δ parameters. Thus, it is the set of all move types that satisfy the corresponding predicates.

These notions are inspired by the Dynamics Analysis of Negotiation Strategies (DANS) framework of [11], which we simplified by modelling unfortunate moves as selfish moves, and fortunate/nice moves as concessions.

The next section illustrates the effectiveness of our abstract framework by introducing an optimal strategy detection algorithm that is based on this framework.

4.1 Optimal Bidding Strategy Recognition

The following strategy correctly determines the opponent's strategy in three to four rounds, unless the opponent would play the random strategy as that could theoretically behave consistently with a different strategy over multiple moves:

Round 1: Randomly select B_1^H from m_H .

Round 2:

Bidding: Randomly select B_2^H from $\{b \in B \mid \tau_H^O((B_1^H, b)) = conc_H^\delta\}$, where δ corresponds to a moderate concession, e.g. $\delta = .1$.

Analysis: After *O* made a bid, compare the first moves of both players to form a first hypothesis Hyp.

Let $t^{O} = \tau_{H}^{O}(\beta_{1,2}^{O}).$

¹For simplicity, we disregard the possibility of using information from a previous encounter with the same opponent here.

If $\sigma_H(\beta_{1,2}^O) \leq \sigma_H(\beta_{1,2}^H) \wedge t^O \in \{conc_H, sil_H\}$, then Hyp := {HH, R}. If $\sigma_H(\beta_{1,2}^O) \geq \sigma_H(\beta_{1,2}^H)$ then Hyp := {CC, R, TT}.

Round 3: The bid of *H* depends on the analysis of round 2.

Bidding in case Hyp = {*HH*, *R*} Randomly select B_3^H from { $b \in B | \tau_H^H ((B_2^H, b)) = conc_H^{\delta_l}$ } where δ_l is the boundary of a large concession, e.g. > 0.2 Bidding in case Hyp = {*CC*, *TT*, *R*} Randomly select B_3^H from { $b \in B | \tau_H^H ((B_2^H, b)) = self_H^{\delta_m}$ } where δ_m is the boundary of a moderately selfish move. Analysis in case Hyp = {*HH*, *R*}: Let $t^O = \tau_H^O(\beta_{2,3}^O)$. if $t^O \in \{self_H, conc_H\}$ and $\sigma_H(\beta_{2,3}^O) \le \sigma_H(\beta_{2,3}^H)$ }, then Conclude that *O* plays Hardheaded, and stop else Conclude that *O* plays Random.

Analysis in case Hyp = {CC, TT, R}: Let $t^{O} = \tau_{H}^{O}(\beta_{2,3}^{O})$. If $t^{O} = self_{H}$, then O does not play Conceder, so we update: Hyp:= {TT, R}. If $t^{O} = \{sil_{H}, conc_{H}\}$, then O does not play Tit-for-Tat, so we update: Hyp:= {CC, R}.

Round 4: Only needed if round 3 ended without conclusions

H's bid: Randomly select B_4^H from $\{b \in B | \tau_H^H((B_3^H, b)) = sil_H\}$. Conclusions: Let $t^O = \tau_H^O(\beta_{2,3}^O)$.

Analysis in case Hyp = {TT, R}: if $t^O = sil_H$, then this does not fit with R and we conclude that O is playing Tit-for-Tat else we conclude that O is playing R.

Analysis in case Hyp = {CC, R}: If $t^O \in {sil_H, conc_H}$, then we conclude that *O* plays Conceder.

else we conclude that O plays Random.

5 Expectations and Aberrations

As our aim is to pro-actively discuss bids with respect to a user's expectation ('guess') of the bidding strategy of the opponent, a mechanism is needed that can detect when a bid deviates from that strategy. The mechanism should be sensitive to the user's estimation of the opponent bidding strategy, which we refer to as *assumption* in the remainder of this paper. A deviation can only be detected if also an expectation can be formulated on the types of move that a negotiator would play if he or she were to play a certain strategy.

Let *S* be a set of bidding strategies. We define an expectation function $\rho : S \times \mathbb{N} \times N \mapsto \Pi^N(\mathcal{P}(M))$ to be a function that given a strategy $s \in S$, a finite number of rounds $r \in \mathbb{N}$, and a negotiator $a \in N$ and produces a sequence of length *r* of

expected move types from M_a corresponding to *s*. Strategy descriptions should be specific enough to derive the δ parameters of the move types, and the behaviour over the rounds and in relation to possible deadlines and/or discount factors. For each of the four typical strategies [5], **Hardheaded** ('Tough negotiator that does not yield easily', i.e. makes mostly silent moves or small concessions), **Conceder** ('Nice negotiator that tends to move towards you', i.e. generally makes concessions), **Tit-for-Tat** ('Somewhat mirrors the moves you are making', i.e. responds with the same type of moves) and **Random** ('Does not follow any of the other strategies', i.e. makes concessions or selfish or silent moves randomly), we give an example for four rounds of negotiation in which the role of the δ parameters is ignored. In each example, the human user is the first to bid in a round.

If the human user *H* estimates the opponent *O* to play a Hardheaded strategy (denoted $HH \in S$) for four rounds, then $\rho(HH, 4, H) = (\{sil_H\}, \{sil_H\}, \{sil_H\})$. Similarly, if the strategy is estimated to be a Conceder strategy *CC*, then $\rho(CC, 4, H) = (\{conc_H\}, \{conc_H\}, \{conc_H\})$. A Random strategy *R* would yield a set of all possible move types per move: $\rho(R, 4, H) = (M_H, M_H, M_H)$. Note that the definition of ρ function for a Tit-for-Tat strategy (denoted *TT*) can only be determined if the bidding strategy of the human user is also given, or if a move type sequence for the same rounds of the human user is provided. Therefore, in case of *TT*, the function is called by $\rho(\langle TT, (conc_H, self_H, conc_H) \rangle$, $4, H) = (\{conc_H\}, \{self_H\}, \{conc_H\})$.

In order to detect aberrations, we need to compare the move types of the actually made moves with the expected move types. For this, we define a set of functions $\tau_a^{a'}$ for all $a, a' \in N$ over bid sequences as follows:

$$\forall \mu \in B \times B : \tau_a^{a'}(\mu) = \begin{cases} sil_a & \text{if } sil_a(a', \mu) \\ conc_a & \text{if } conc_a(a', \mu) \\ self_a & \text{if } self_a(a', \mu) \end{cases}$$

 $\forall a, a' \in N, \forall i, j \in \mathbb{N}, \forall \beta_{i,j}^{a'}$

$$\tau_a^{a'}(\beta_{i,j}^{a'}) = (\tau_a^{a'}((B_i^{a'}, B_{i+1}^{a'})), \dots, \tau_a^{a'}((B_{j-1}^{a'}, B_j^{a'})))$$

Aberration detection is now as simple as checking for each element in $\tau_a^{a'}(\beta^{a'})$ if it occurs in the corresponding element of $\rho(s, r, a)$. By setting the δ parameters appropriately, minor deviations can be ignored.

6 Generating Explanations

Now that we have a method to indicate a party's bid as deviating from the user's assumption of that party's strategy and a classification of the deviation in terms of a

Our μ^{-1}	Expected μ	Actual μ	Explanation (of aberration)
Silent	Silent	Concession	A tit-for-tat player would typically not respond with a conceding move to your inaction
		Selfish	A tit-for-tat player would typically not respond with a selfish move to your inaction
Concession	Concession (equal)	Silent	A tit-for-tat player would typically not respond with inaction to your conceding move
		Concession (smaller)	A tit-for-tat player would typically not respond with a conceding move that is much smaller than your concession
		Concession (larger)	A tit-for-tat player would typically not respond with a conceding move that is much larger than your concession
		Selfish	A tit-for-tat player would typically not respond with a selfish move to your conceding move
Selfish	Selfish (equal)	Silent	A tit-for-tat player would typically not respond with inaction to your selfish move
		Concession	A tit-for-tat player would typically not respond with a conceding move to your selfish move
		Selfish (smaller)	A tit-for-tat player would typically not respond with a selfish move that is much smaller than your selfish move
		Selfish (larger)	A tit-for-tat player would typically not respond with a selfish move that is much larger than your selfish move

 Table 1
 The aberration explanation matrix for our tit-for-tat expectation function (according to template 1)

direction and size according to (a simplification of) the DANS framework, we need to convey this information to the user. To this end, we propose the use of *aberration explanation matrices*, providing a visualization of the expectation function as well as an explanation for all combinations (i.e. aberrations) of the expected move type(s) and size(s) and the actual move type(s) and size(s) of the opponent.

As an example, we provide the aberration explanation matrix for the expectation function $\rho(\langle TT, \mu^{-1} \rangle, 2, H)$ in Table 1, which provides explanations for aberrations from an expectation of the tit-for-tat strategy. The matrix is set up according to the following template (Template 1): **'An expected strategy player would typically not respond with an** *actual* μ **to your** μ^{-1} , where *expected strategy* and *actual* are parameters to be instantiated. For simplicity, we leave move size information out.

Note that we use μ^{-1} to signify the last two bids of our user, i.e. defining a decrease of our own utility with x as a concession towards the opponent of size r = -x; any bids before those last two are not used. Moreover, we use 'smaller' and 'larger' here as a difference between the expected r (which is equal to the user's own r in μ^{-1}) and the actual r that is larger than 10% ($\delta = .1$).

For each supported negotiation strategy, an explanation matrix should be provided, establishing a design from which the implementation can be constructed. As discussed in Sect. 7, the results from two pilot studies encouraged us to design an additional explanation template. The idea of the second template is to suggest to the user which strategies would be consistent with the observed behaviour, instead of only pointing out the behaviour is not consistent with the user's current guess, as is done in Template 1. The alternative explanation template that we used is Template 2: **'Responding with an** *actual* μ **to your** μ^{-1} **is more consistent with** *consistent* strategies', where *actual* and *consistent* are parameters to be instantiated.

7 Evaluation

This section describes our evaluation of the aberration detection mechanism and explanation matrix we introduced.

If we would try to introduce our mechanisms at once for all aspects of the bidding phase, the experiments would have to cover too many variables at once for a meaningful evaluation. In direct relation to that, the number of participants would make the experiment infeasible. Finally, if we just let participants negotiate then we cannot control how often aberrations would occur, or whether they would occur at all.

Therefore, we designed the experiment in such a way that greatly reduced the number of variables we would test for and in a manner that gives us control of the aberrations that would occur in the experiment.

We decided to test the participants' understanding of the typical bidding strategies discussed in Sect. 3. In a between-subject setup, participants negotiated against automated opponents. The bidding strategy used by the automated opponents (agents) varied over the well-known bidding strategies. The participants were asked to identify the bidding strategy of the opponent. We controlled the variation over the bidding strategies, as well as whether or not the participant was supported by our explanation mechanism. We evaluated the effectiveness of this mechanism in improving a participants' understanding of the opponent's bidding negotiation strategy. We hypothesized that our explanation mechanism improves a PN user's understanding of a negotiation, and specifically, of the strategy that the other party uses. By some pilot experiments we found that this, more than expected, depends on the contents of an explanation (of an aberration), suggesting consistent strategies is more effective than explaining why observed behaviour is inconsistent for example.

Therefore, we finally evaluated our hypothesis that our explanation mechanism based on aberrations increases a user's understanding of the opponent's strategy through controlled between-subjects experiments, in which one group did not receive such explanations, a second group received explanations of why a chosen strategy seemed less likely to fit and the third group received explanations about which strategies would be consistent with the behaviour of the opponent. All participants were tasked with negotiating against a (computer-controlled) opponent that employed one of the four defined strategies, with the goal to find out which strategy this opponent is playing.

7.1 Preparation

In the experiments, each participant first received short definitions of the four possible negotiation strategies, and a brief training in the use of the PN itself. The goal of determining the opponent's strategy without regarding the result of the negotiation itself was made clear. All negotiations were performed in the multi-issue Jobs domain (see Fig. 3), which was selected due to being easily understandable for novice users while still providing enough complexity and thus flexibility and variation in the negotiations. The issues and values in this domain could be explored by the user in the PN; all issue weights and valuations were fixed for both parties, i.e. all preferences are fully known from the start and never change. Each participant was asked to perform at least four negotiation sessions. In the first four negotiations, each participant played against each possible opponent at least once, in a random order. Participants were not informed about the fact that each opponent would only be encountered once. In all sessions, the participant's experiment condition did not change.

Based on the optimal bidding strategy presented in Sect. 4.1, which requires three to four rounds of bidding, we allowed each participant sufficient room with at most ten rounds per negotiation. Our evaluation results show that this is indeed sufficient.

7.2 Conditions

During a negotiation, the opponent would never accept a bid (i.e. the opponent never ended the negotiation); only the participant could end the negotiation when he or she was convinced of having identified the strategy of the opponent successfully (which happened automatically after the ten bid limit as well). This was known to the participants. The participant's assumption about the opponent's strategy was requested after each move of the opponent, as illustrated in Fig. 3. As the participant always had to start the negotiation with an opening bid, the first move of the opponent was already a response to the participant's first move. Thus, with the participant always making the opening bid, the first assumption about the opponent's strategy is requested after four bids (i.e. a move from both parties). If the opponent would start the negotiation, a participant would have only three bids in total to base his or her first estimate on, which we considered too much of a guess.

Participants were not informed of the correctness of their assumptions of the opponent's strategy at any point during the experiment. Moreover, the order in which the four strategies (i.e. assumption options) were displayed was randomized in each negotiation session. In the explanation conditions, the request for selecting the strategy the user thinks the opponent is employing was potentially accompanied by an



The other side has made a new off	er.
salary	2000
fte	1.0
work from home	0
lease car	no
permanent contract	no
career development opportunities	low
Good for other	
I think the opponent is playing this	strategy:
Conceder Random Hardhe	eaded . Tit-For-Tat

explanation as detailed in this paper. Note that such an explanation was always based on the participant's previous selection of the opponent's strategy, as it would otherwise be too easy to just 'try all buttons' and see how the system responds.

7.3 Metrics

Each bid, and each selection of an assumption, was logged. Moreover, after each negotiation, participants were asked to rate on a five-point Likert scale (i) how sure he or she was about the determination of the opponent's strategy and (ii) how well he or she was assisted by the system in making this determination. Before starting the first negotiation, participants had to rate their prior knowledge on negotiations (on a scale of 1–10) and indicate what kind of moves they would expect from each of the fours strategies (with percentages). We did this in order to measure the participant's understanding of the four negotiation strategies, and posed the same questions in a post-questionnaire. For hardheaded, we counted the answer as correct when silent moves got the largest portion, along with a non-zero portion for concessions. For tit-for-tat, each portion had to be at least 20%. For random, each portion had to be at least 30%. Finally, for conceding, concession moves had to have the largest portion. In the post-questionnaire, we also asked how difficult the user found the task.

8 Results

This section describes and discusses the results of our experiments. Two pilots were held with relatively small groups, after which a large-scale online experiment was performed.

8.1 Pilot 1

To determine the suitability of the experimental setup and our software for the goals of our evaluation, we performed an exploratory pilot study with 11 participants, all male post-graduates in the department of the authors. Compared to the final setup as described above, in this first setting only a post-questionnaire was held, in which the questions about 'What kinds of moves would you expect ...' were not posed. Furthermore, we included the question 'You had at most 10 bidding rounds to identify the strategy of each opponent. Was this sufficient?'. Finally, the participant's existing knowledge on negotiation was requested on a Likert scale (instead of on a scale of 1 to 10 as in the final setup). The results of this pilot are summarized in Fig. 4.

On average, each participant negotiated five times. No technical problems were encountered. In about 80% of the negotiations, the final answer on which strategy the opponent was playing was correct; 6 out of the 11 participants even achieved a 100% score on this in the pilot. These high scores are also apparent in the questionnaire results (see Fig. 4), as the participants were very sure about their answers ($\mu = 4.4$) and did not find the task very difficult ($\mu = 2.0$). The condition (i.e. receiving explanations or not) did not have any significant effect, which we believe was both due to the small sample size (only five people received explanations) and the low difficulty of the task for this highly educated group.

8.2 Pilot 2

Following the inconclusive results from the initial pilot, an additional pilot was held with a mixed-gender group of 39 third-year bachelor's students following a minor on negotiation. For this pilot, the opponents were tweaked in order to slightly increase the difficulty of the task, and as aforementioned, the single post-questionnaire was split into pre- and post-questionnaires, to which questions were added in order to measure the participant's understanding of the negotiation strategy. In addition, as it was clear the limit of ten bids was more than sufficient from the initial pilot, both from the related questionnaire question ($\mu = 4.3$) and the fact that on average only five bids were made per negotiation (in about 2.5 min), the question related to this fact was removed. In the second pilot, on average, seven bids were made per negotiation (in about 2.5 min as well).



Fig. 4 Results from the first pilot (N = 11)



Fig. 5 Results from the second pilot (N = 30)

As illustrated in Fig. 5, the results from the second pilot were in some sense the opposite of the results of the initial pilot. The 30 students that completed the task correctly identified only 39% of the opponents. As there are just four options to pick from, it can safely be concluded that the participants performed very poorly, indicated by the participants themselves as well through being less sure ($\mu = 3.7$) and feeling the task was more difficult ($\mu = 3.2$). Just like in the initial pilot, no significant results based on the condition were found. Due to this fact and space constraints, further results from the pilots will not be discussed here.

8.3 Full Experiment

Based on the two pilots, we introduced the 'new' explanation strategy in which strategies that would be consistent with an aberration are identified. Our main reason for doing this is that the original explanations that detail the aberrations only gave participants more knowledge about the strategy they had currently guessed, while this new form would also communicate information about one or more other strategies. In addition, in order to gain more participants from more varied backgrounds, we decided to perform a large-scale online experiment. Therefore, instead of face-to-face training as given in both pilots, this part was digitalized.²

To gain a sufficient number of participants, we made use of the Amazon Mechanical Turk [1], 'a marketplace for work that requires human intelligence'. In order to ensure high-quality work from participants that we could not have direct interaction with, a number of measures were taken (that are common practice [12]):

- Only participants who performed at least 1000 tasks with at least a 99% acceptance rate were allowed in.
- Only participants from English-speaking countries were allowed to participate.
- In the pre-questionnaire, besides questions ensuring informed consent, questions were added to verify that the participant understood the training. Participants that did not answer these questions correctly were prevented from continuing in the experiment.
- A unique code was generated upon completion; participants submitting incorrect codes were rejected.

Out of 198 'turkers' that started the task, 84 completed the experiment successfully.³ 31% of participants was female.

The main results of the experiment are shown in Fig. 6. Participants correctly identified the strategy of 44% of their opponents, using 6.7 bids on average (in about 2 minutes). Independent sample T-tests were used to identify differences between participants that received any form of explanation, the 'old' inconsistent-behaviour explanations, the 'new' suggesting consistent strategies explanations or no explanations at all. Participants receiving any form of explanation on average had a 23.2% ($\pm 11.4\%$) better score against opponents playing a random strategy (t(79) = 2.029, p = 0.046) than participants that received no explanations at all. Moreover, such participants had a 13.5% ($\pm 6.4\%$) better score for correctly specifying the hardheaded strategy in the post-questionnaire (t(79) = 2.098, p = 0.039)as well. Participants receiving the 'new' form of explanation on average had a 15.3% ($\pm 5.7\%$) better score against any opponent (t(79) = 2.691, p = 0.009). As no significant difference was found for the 'old' form of explanation, we conclude that suggesting consistent strategies is more effective than explaining why observed behaviour is inconsistent. Participants receiving the 'new' form of explanation on average also had a 46.2% ($\pm 10.9\%$) better score against opponents playing a

²The training for our experiment can be found at *anonymized*.

³These numbers fall within the expected range for MTurk experiments of this type [9].



Fig. 6 Results from the final experiment (N=81)

random strategy (t(79) = 4.253, p = 0). Such a difference was not found for the other opponents, perhaps because there is more overlap in their behaviours. This is especially true for the hardheaded and the conceding opponent (mainly varying their concession rate), but also for the tit-for-tat opponent if the participant mainly performs concessions him or herself, a sub-optimal strategy. In addition, participants receiving **the 'new' form of explanation** on average had a 24.8% (±11.8%) **better score in their second negotiation** (t(77) = 2.103, p = 0.039) and a 26.4% (±11.6%) **better score in their third negotiation** (t(78) = 2.269, p = 0.026) than the other participants. Interestingly, such differences were not found for the first and fourth negotiations, suggesting a relatively steep learning curve.

No further significant mean differences were found based on the explanation conditions. However, further correlation analysis suggests that **participants that on average did more bids in a negotiation identified the opponent correctly more often** (r = 0.233, p = 0.037), but were also less sure of their answer each time (r = -0.388, p = 0). Besides the 'new' form of explanations as aforementioned, no other factor directly correlates with the amount of correct answers of a participant. Interestingly, just like in both pilots, participants that indicated they were more knowledgeable about negotiations at the start did not perform significantly better. However, such participants seemed to perform less bids (r = -0.327, p = 0.003), to be more sure about their answers (r = 0.297, p = 0.007), and to feel more assisted (r = 0.304, p = 0.006) regardless. Being more sure ($\mu = 3.9$) and feeling more assisted ($\mu = 3.7$) positively correlates in general (r = 0.361, p = 0.001). Finally, no significant impact of either the order of opponents in the four negotiations or the ordering of the strategies in the interface in each negotiation, which were both randomized, was found.

9 Conclusion

If we can automatically detect when the user or the opponent *seems to deviate* from a strategy, or that our opponent model might be wrong, or that the user or the opponent might have changed his preferences, or might just simply have made a mistake, this opens the possibility for pro-actively discussing these strategies with the user.

The technology we introduce in this paper has been developed with the aim of supporting human negotiators in gaining insight into the bidding strategy of the opponent and into their own bidding behaviour. The core technology we developed consists of two aspects: *aberration detection*, and the notion of an *explanation matrix*. The aberration detection mechanism identifies when a bid falls outside the range of expected behaviour for a specific strategy. The explanation matrix is used to decide when to provide what explanations. We evaluated our work experimentally in a task in which participants are asked to identify their opponent's strategy in the Pocket Negotiator. On a technical note, our explanations, as these could quickly be implemented in our explanation matrices. As the number of correct guesses increases with explanations, indirectly, these experiments show the effectiveness of our aberration detection mechanism. Our experiments show that suggesting consistent strategies is more effective than explaining why observed behaviour is inconsistent.

Future Work

Our evaluations used a single negotiation domain and four negotiation strategies. Although we believe the domain and strategies are representative, the effects of using more complex domains and/or strategies can be examined.

Finally, note that our work is applicable to evaluating the bids a user makes him or herself as well, e.g. for confirming that a user's bids comply with that user's intended strategy (as set in the system) and providing an explanation when this is not the case (before a bid is actually made).

References

- 1. Amazon. Mechanical Turk (2018). https://www.mturk.com/. Accessed 16 Nov 2018
- 2. Baarslag, T.: What to bid and when to stop. Ph.D. thesis, Delft University of Technology (2014)
- Baarslag, T., Gerding, E.H.: Optimal incremental preference elicitation during negotiation. In: Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15, pp. 3–9. AAAI Press (2015)
- Baarslag, T., Hindriks, K.V., Jonker, C.M., Kraus, S., Lin, R.: The first automated negotiating agents competition (ANAC 2010). In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (eds.) New Trends in Agent-Based Complex Automated Negotiations. Studies in Computational Intelligence, vol. 383, pp. 113–135. Springer, Berlin (2012)
- Baarslag, T., Fujita, K., Gerding, E.H., Hindriks, K., Ito, T., Jennings, N.R., Jonker, C., Kraus, S., Lin, R., Robu, V., Williams, C.R.: Evaluating practical negotiating agents: results and analysis of the 2011 international competition. Artif. Intell. 198, 73–103 (2013)

- Baarslag, T., Aydoğan, R., Hindriks, K.V., Fuijita, K., Ito, T., Jonker, C.M.: The automated negotiating agents competition, 2010-2015. AI Mag. 36(4), 115–118 (2015)
- Baarslag, T., Hendrikx, M.J.C., Hindriks, K.V., Jonker, C.M.: Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. Auton. Agents Multi-Agent Syst. 30(5), 849–898 (2016)
- 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 26th International Joint Conference on Artificial Intelligence, IJCAI'17, pp. 4684–4690. AAAI Press (2017)
- 9. Goodman, J.K., Cryder, C.E., Cheema, A.: Data collection in a flat world: the strengths and weaknesses of Mechanical Turk samples. Behav. Decis. Mak. **26**(3), 213–224 (2012)
- Hindriks, K.V., Jonker, C.M.: Creating human-machine synergy in negotiation support systems: towards the pocket negotiator. In: Proceedings of the 1st International Working Conference on Human Factors and Computational Models in Negotiation, HuCom'08, New York, NY, USA, pp. 47–54. ACM (2009)
- Hindriks, K.V., Jonker, C.M., Tykhonov, D.: Let's DANS! an analytic framework of negotiation dynamics and strategies. Web Intell. Agent Syst. 9(4), 319–335 (2011)
- Ipeirotis, P.G., Provost, F., Wang, J.: Quality management on Amazon Mechanical Turk. In: Proceedings of the ACM SIGKDD Workshop on Human Computation, HCOMP'10, New York, NY, USA, pp. 64–67. ACM (2010)
- 13. Johnson, E., Gratch, J., DeVault, D.: Towards an autonomous agent that provides automated feedback on students' negotiation skills. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS'17, Richland, SC, USA, pp. 410– 418. International Foundation for Autonomous Agents and Multiagent Systems (2017)
- Jonker, C.M., Hindriks, K.V., Wiggers, P., Broekens, J.: Negotiating agents. AI Mag. 33(3), 79–91 (2012)
- 15. Koeman, V., Hindriks, K.V., Gratch, J., Jonker, C.: Recognising and explaining bidding strategies in negotiation support systems, extended abstract. In: Agmon, N., Taylor, M.E., Elkind, E., Veloso, M. (eds.) Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2019, pp. 2063–3064. International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org) (2019)
- 16. Lee, W.-P.: Towards agent-based decision making in the electronic marketplace: interactive recommendation and automated negotiation. Expert Syst. Appl. **27**(4), 665–679 (2004)
- Lewicki, R.J., Saunders, D.M., Barry, B., Minton, J.W.: Essentials of Negotiation. McGraw-Hill, Boston (2003)
- Miller, T.: Explanation in artificial intelligence: insights from the social sciences. Artif. Intell. 267, 1–38 (2019)
- Moulin, B., Irandoust, H., Bélanger, M., Desbordes, G.: Explanation and argumentation capabilities: towards the creation of more persuasive agents. Artif. Intell. Rev. 17(3), 169–222 (2002)
- Pommeranz, A., Broekens, J., Wiggers, P., Brinkman, W.-P., Jonker, C.M.: Designing interfaces for explicit preference elicitation: a user-centered investigation of preference representation and elicitation process. User Model. User-Adapt. Interact. 22(4), 357–397 (2012)
- 21. Raiffa, H.: The Art and Science of Negotiation: How to Resolve Conflicts and Get the Best Out of Bargaining. Harvard University Press, Cambridge (1982)
- Raiffa, H., Richardson, J., Metcalfe, D.: Negotiation Analysis: The Science and Art of Collaborative Decision Making. Harvard University Press, Cambridge (2003)
- Tintarev, N., Masthoff, J.: Explaining recommendations: design and evaluation. In: Ricci, F., Rokach, L., Shapira, B. (eds.) Recommender Systems Handbook, pp. 353–382. Springer, Boston (2015)
- 24. Van De Kieft, I., Jonker, C.M., Van Riemsdijk, M.B.: Explaining negotiation: obtaining a shared mental model of preferences. In: Proceedings of the 24th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems Conference on Modern Approaches in Applied Intelligence Volume Part II, IEA/AIE'11, pp. 120–129. Springer, Berlin (2011)