Recognising and Explaining Bidding Strategies in Negotiation Support Systems

Extended Abstract

Vincent J. Koeman
Delft University of Technology
Delft, The Netherlands
v.j.koeman@tudelft.nl

Koen V. Hindriks
Delft University of Technology
Delft, The Netherlands
k.v.hindriks@tudelft.nl

Jonathan Gratch
USC Institute for Creative Technologies
Playa Vista, CA, United States
gratch@ict.usc.edu

Catholijn M. Jonker
Delft University of Technology
Delft, The Netherlands
c.m.jonker@tudelft.nl

ABSTRACT

To improve a negotiator's ability to recognise bidding strategies, we pro-actively provide explanations that are based on the opponent's bids and the negotiator's guesses about the opponent's strategy. We introduce an aberration detection mechanism for recognising strategies 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 environment of a negotiation support system, namely the Pocket Negotiator (PN). We implemented our explanation mechanism in the PN and experimented with different explanation matrices. 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.

KEYWORDS

strategy recognition; aberration detection; explanation matrix; bidding strategies; negotiation support system

ACM Reference Format:

Vincent J. Koeman, Koen V. Hindriks, Jonathan Gratch, and Catholijn M. Jonker. 2019. Recognising and Explaining Bidding Strategies in Negotiation Support Systems. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

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 is the Pocket Negotiator (PN) [5]. Our focus is on supporting (novice) negotiators in the bidding phase of the PN through experiential learning. The state of the art in research on bidding strategies focuses on automated negotiating agents, i.e., agents that negotiate on the user's behalf. The bidding strategies developed for these agents form the core of the bidding advise that the PN provides to its users. The support provided by the PN consists of bid suggestions and a visualisation of the bid space and its Pareto Optimal Frontier. An expert negotiator can use this interface to quickly create bids that are in line with his or her bidding strategy to the opponent. Similarly, the visualisation gives the negotiator an overview of the bids made by himself and by the other party. For expert negotiators this might be enough to estimate the bidding strategy of the opponent, but is this also enough for non-professionals?

The technology we introduce has been developed with the aim of supporting human negotiators in gaining insight into the bidding strategy of the opponent. The core technology we developed consists of two aspects: *aberration detection*, and the notion of an *explanation matrix*. If we can automatically detect when the user or the opponent *seems to deviate* from a strategy, this opens the possibility for pro-actively discussing these strategies with the user. The user might deviate intentionally or unintentionally. 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 PN. In such a case, it is important to discover this as quickly as possible. Quickly detecting a deviation in the opponent's behaviour is just as important for the negotiation. Similar reasons can be the cause for the deviation: our opponent model might be wrong, the opponent might have changed his preferences, the opponent might have changed strategy, or might just simply have made a mistake. Finding the real cause of the aberration is beyond the scope of this work, however, a mechanism for explanation is essential for all further steps.

We thus introduce the aberration detection mechanism and the notion of an explanation matrix, and we test these in controlled human-machine experiments in which we test the participants' understanding of bidding strategies. In a between-subject set-up, participants negotiated against automated opponents. The bidding strategy used by the automated opponents (agents) varied over well-known bidding strategies. The participants were asked to identify the bidding strategy of the opponent. We controlled the variation

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

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 hypothesised that our explanation mechanism improves a PN user's understanding of a negotiation, and specifically, of the strategy that the other party uses. 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.

2 RELATED WORK

Explanations are employed in many sub-fields of artificial intelligence [7]. Baarslag et al. [3] identify, however, that allowing users of negotiation support systems to "trust the system through coparticipation, 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.

Most research on 'opponent modelling' in (automated) negotiation focuses on determining the preferences of the opponent [2]. We instead focus on determining the (bidding) strategy that an opponent uses. We aim for 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 Johnson et al. [6] 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. Our explanation mechanism for opponent strategy recognition could be directly relevant to negotiation training.

3 CONTRIBUTIONS

As our aim is to pro-actively discuss bids with respect to a user's expectation ('guess') of the bidding strategy of the opponent, we propose a mechanism that can detect when a bid deviates from that strategy. The mechanism is sensitive to the user's estimation of the opponent bidding strategy. A deviation can thus only be detected if an expectation is formulated on the types of move that a negotiator would play if he or she were to play a certain strategy. Due to space constraints, we cannot go into details on this mechanism here.

Based on our aberration detection method, we convey the resulting information to the user. To this end, we use *aberration explanation matrices*, providing 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. The following template was initially used for each explanation: **"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. μ represents the last *move* of the opponent, i.e. the difference between the last two bids of the opponent. μ^{-1} signifies the same for our own user. For each supported negotiation strategy, an explanation matrix should be provided, establishing a design from which the implementation can be constructed.

However, the results from two pilot studies encouraged us to design explanations according to a different template. The idea 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. The explanation template we thus eventually used is: **"Responding with an *actual* μ to your μ^{-1} is more consistent with *consistent strategies*."**, where *actual* and *consistent* are parameters to be instantiated.

4 EVALUATION AND CONCLUSIONS

We 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, whilst others did (upon aberrations). All participants were tasked with negotiating against a (computer-controlled) opponent that employed one of four common negotiation strategies, in order to find out which strategy this opponent is playing. Each participant was trained on the use the PN itself and the various negotiation strategies. 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, which was selected due to being easily understandable for novice users whilst 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 four negotiations in the PN, in which the participant's experiment condition did not change. The participant's assumption about the opponent's strategy was requested after each move of the opponent, in one condition always accompanied by an explanation.

For our experiment, we made use of the Amazon Mechanical Turk [1]. Out of the 198 'turkers' that started our task, 84 completed the experiment[1]. 31% of participants was female. Participants correctly identified the strategy of 44% of their opponents, using 6.7 bids on average (in about two minutes). A t-test shows that **participants receiving explanations correctly identified 15.3% (±5.7%) more opponents on average** ($t(84) = 2.691$, $p = .009$).

We introduced an aberration detection mechanism for recognising strategies 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. 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.

[1] These numbers fall within the expected range for MTurk experiments of this type [4].

REFERENCES

[1] Amazon. 2018. Mechanical Turk. https://www.mturk.com/. (2018). Accessed: 2018-11-16.

[2] Tim Baarslag, Mark J. C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems* 30, 5 (Sept. 2016), 849–898.

[3] Tim Baarslag, Michael Kaisers, Enrico H. Gerding, Catholijn M. Jonker, and Jonathan Gratch. 2017. 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)*. AAAI Press, 4684–4690.

[4] Joseph K. Goodman, Cynthia E. Cryder, and Amar Cheema. 2012. Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples. *Behavioral Decision Making* 26, 3 (April 2012), 213–224.

[5] Koen V. Hindriks and Catholijn M. Jonker. 2009. 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)*. ACM, New York, NY, USA, 47–54.

[6] Emmanuel Johnson, Jonathan Gratch, and David DeVault. 2017. 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)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, USA, 410–418.

[7] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (Feb. 2019), 1–38.