



How truthful explanation of a conceder negotiation agent affects human trust, compared to a neutral one

Pavel Germanov¹

Supervisors: Myrthe Tielman¹, Carolina Jorge¹

¹EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 25, 2023

Name of the student: Pavel Germanov
Final project course: CSE3000 Research Project
Thesis committee: Myrthe Tielman, Carolina Jorge, Rafa Bidarra

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

Trust in negotiation agents plays a crucial role in their adoption and utilization. However, there is not enough research on what factors influence it. This paper aims to investigate how different explanations of a negotiation agent's strategy affect human trust and decision-making. Specifically, it compares the effects of a truthful explanation and a neutral explanation on trust in a conceder negotiation agent. An experiment was conducted involving human participants interacting with a negotiation agent in an imaginary negotiation scenario. The participants were divided into two groups, with one group receiving a truthful explanation of the agent's strategy, and the other receiving a neutral explanation. Trust levels were assessed through a questionnaire and the results were analyzed using an independent samples *t* test. No significant difference was found between the two groups, but further research needs to be done before conclusions are drawn.

1 Introduction

The field of negotiation agents keeps on expanding and there is an increasing amount of research being done on it [2; 16]. They have a wide variety of uses but in their essence, they serve as artificial intelligence (AI) tools that assist with achieving the best outcome for their human counterparts in a negotiation scenario [14]. They are able to assist with the exploration of the negotiation space and also to use a preset strategy and preferences in order to suggest the next optimal bid in a negotiation sequence. Then, unless it is a fully automatic system, the human can choose whether to make use of the agent's suggestion or act independently [2]. This decision can be influenced by different factors, but one of the most significant ones is the human's trust in the agent [23]. Understanding what can affect this trust allows for designing and presenting negotiation agents in such a way that people would be more willing and confident about using them [17]. With this kind of technology becoming more prevalent in different industries this is an important aspect that must be considered by the agents' developers in order to provide for a smoother adoption process [21].

The topic of human trust in AI systems has been analyzed in-depth in many different studies with [9] providing a detailed review of the current research, while in [24] a similar analysis can be found specifically about trust in AI-assisted decision-making. Nonetheless, particularly the topic of trust in negotiation agents is less researched than in some other areas of AI. This means that more analysis is needed in order to conclude what affects human trust in negotiation agents.

Researchers agree that the way an AI is explained to its users affects people's trust in it [27]. Providing an interpretable description of how it works increases transparency which leads to wider adoption of the AI system [17]. This paper investigates how different explanations of a negotiation agent's strategy affect human trust in it. More specifically it focuses on:

How truthful explanation of a conceder negotiation agent affects human trust, compared to a neutral one.

To analyze this, an experiment was designed and conducted. It consisted of human participants being asked to use a negotiation agent in an imaginary negotiation scenario, and then to fill out a questionnaire that reflected their level of trust in the agent. To make a comparison between the two different explanation styles, half of the subjects were provided with a detailed and accurate description of the strategy (truthful explanation), and the other half with a minimalistic one that does not reveal any meaningful details (neutral explanation). For both groups, the negotiation agent followed a "conceder" strategy that made the agent to suggest bids that were progressively more favorable to the other party over time.

The rest of the paper is structured as follows: In Chapter 2, literature related to concepts used throughout the study is reviewed. Chapter 3 focuses on the methodology of the experiment, its results are reported in Chapter 4, and Chapter 5 summarizes the ethical aspects of the research. The data is analyzed and discussed in Chapter 6, Chapter 7 concludes the main part of the paper, and Chapter 8 is dedicated to acknowledgments.

2 Literature review

2.1 Trust

Trust is a vague concept and there are many different interpretations that can also depend on the context in which they are being used [15]. It is desirable to have an explicit definition of trust used throughout the entire study because there are many different interpretations that can also depend on the context in which they are being used [5]. Therefore, for this research, the following definition of trust was used:

A sentiment resulting from knowledge, beliefs, emotions, and other elements derived from lived or transmitted experience, which generates positive or negative expectations concerning the reactions of a system and the interaction with it [4, p. 2].

This definition was chosen because it suits this paper's context well. It describes trust in the context of human-system interaction and also in [4] it is used for a similar research - analyzing the effect of a cruise control's explanation on human trust.

2.2 Explainable artificial intelligence

With artificial intelligence occupying an increasing number of roles requiring trust, the conditions that influence it must be analyzed [22]. One of these aspects is the influence of AI interpretability on trust. Explainable AI (XAI) is the main field of study that focuses on this [6], and the consensus is that developers should be providing AI users with a transparent and comprehensible description of its behavior in order to increase human trust in a system [27].

Nonetheless, the influence that XAI has on trust may depend on different factors. For example, in [3] and [10] it is shown that the explanation type significantly affects trust in autonomous vehicles and medical decision systems, while in

[19] other factors were concluded to be more important for automatic classification systems.

3 Methodology

As mentioned in the introduction, the experiment in this study followed a between-subjects design where half of the subjects were presented with a truthful explanation of the negotiation agent and the other half with a neutral one that does not reveal any specifics. Each participant was asked to complete a simulated negotiation against a computer opponent while using the conceder agent as an assistant. The following subsections discuss different aspects of the experiment in more detail.

3.1 Hypothesis

The objective of this paper is to compare the trust effect of a truthful versus a neutral explanation of a conceded negotiation agent. Based on this goal, the null hypothesis that is used is:

People trust a conceder negotiation agent with a neutral strategy explanation equally as much as an agent with a truthful one.

3.2 Participants

In total 30 participants were recruited for the study, resulting in 15 participants receiving the "neutral" explanation and 15 the "truthful" one. All of them were university students with a technical background between the ages of 18 and 24 with the participation being voluntary and with no compensation.

3.3 Materials

Pocket Negotiator

The negotiation simulation was done using the web tool Pocket Negotiator [13] which allows for setting up different domains and then performing negotiations using them. It provides options for configuring negotiation issues and preference profiles, native support for negotiation agents, and simulating negotiations against an automatic (computer) opponent [1; 13].

Conceder negotiation agent strategy

There exists a wide range of strategies employed by negotiation agents [26], but this paper is concerned only with the conceder one. It is described as a time-dependent tactic that quickly goes to the agent's reservation value [7], and since Pocket Negotiator provides an assistant agent with such a strategy, it was directly used for the purpose of this experiment. The technical implementation of the strategy in Pocket Negotiator corresponds to the one used in the Genius environment [11].

Negotiation domain

The domain used in the experiment was a theoretical pizza preparation scenario with different cooking tasks that had to be allocated between the participant and the computer opponent. All but one of the negotiation issues resembled parts of the preparation process and the participants had to bid on which tasks should be assigned to which party. The final issue was about the amount of pizza each party should receive at the end, with the options being 33%, 50%, and 66%. The full domain description can be found in Appendix A.

Agent explanations

As the independent variable of the study, agent explanations had to be picked carefully to reach a reliable conclusion. To mitigate any side effects, the two explanations followed a similar style while also following their original goal - one truthfully explaining the 'conceder' approach and the other neutrally explaining the agent without mentioning the strategy. To achieve this, both explanations started with the same neutral introduction of the negotiation agent:

"You shall be assisted by a negotiation agent in the following negotiation session. In the process of evaluating, the agent shall take your preferences into consideration, and based on that it shall evaluate bids to and from the opposing party."

Then, the truthful explanation was continued with a short description of the conceder strategy discussed in Section 3.3:

"The agent would propose bids that will be progressively more favorable to the other party over time until the end of the negotiation. The goal is to find a mutually satisfactory outcome, even if it means accepting less favorable terms than initially desired."

While the neutral explanation was finished with a single universal sentence:

"The goal of the agent is to maximize the overall result attainable for both parties."

Because the participants could be students who might have not had prior experience with similar systems, technical terms were avoided. Instead, the focus was on making the explanations easy to interpret for everyone [6].

Measurements

The comparison between the two explanations was done by quantifying the participants' trust in the negotiation agent for both scenarios. Accordingly, a two-part questionnaire was used: one part to gather background information about the participants, and another to measure their self-reported trust after they had used the agent.

The first part consisted of general demographic questions (e.g. age group, gender), some questions regarding negotiations (e.g. prior experience, theoretical knowledge), and a question about their tendency to trust AI to make accurate and reliable decisions. The goal of these questions was to make it easier to expose any possible confounding variables that could influence the results.

The second part was taken from the validated trust scale recommended in [12]. It consists of eight questions that cover multiple aspects of trust in artificial intelligence and uses a 5-point Likert scale ranging from *Strongly disagree* to *Strongly agree* (see Appendix B). The questions reflect on subjective topics like confidence, predictability, likability, and more, with the end goal being to measure the participants' perceived trust in the negotiation agent. The individual trust scores were calculated by taking the sum of the questionnaire answers, and then the data sets of the two groups were used to either support or reject the null hypothesis.

3.4 Procedure

The participation process of the experiment involved the following steps:

1. Completing a consent form
2. Answering the first part of the questionnaire
3. Receiving a tutorial on using Pocket Negotiator
4. Reading the provided negotiation scenario and the negotiation agent strategy explanation
5. Negotiating within Pocket Negotiator using the negotiation agent
6. Answering the second part of the questionnaire

The participants took part either online or in person under the guidance of a researcher following the aforementioned method.

4 Results

4.1 Participants

To ensure that they did not influence the results, the even distribution of demographics, prior experiences, and trust in AI between the "neutral" and the "truthful" group of the study had to be verified. All of the participants reported being in the 18-24 age group and that they had grown up in Europe so no further analysis is needed for these factors. In the following subsections, the remaining distributions are examined for their significance.

Gender

The participants could report their gender with the options being male, female, non-binary, or other. Only male and female responses were submitted, with their distribution between the two groups visualized in Figure 1. Since the data is not normally distributed, a Mann-Whitney U test ($\alpha = .05$) [18] was used to show that there is no significant difference between the gender distribution of the two groups, $U = 105$, $p = .772$.




Figure 1: Gender distribution per group

Field of study

The participants were also asked to submit their field of study with the options being computer science, mathematics, electrical engineering, or other. Only the first three options were selected and the distribution can be viewed in Figure 2. Because the "truthful" data set consists only of computer science students, the mathematics students were combined with the electrical engineering students to become a larger non-computer science group (see Figure 3). This way a more effective statistical test could be performed by only analyzing the number of non-computer science students in each group. Since the data is not normally distributed, a Mann-Whitney U test ($\alpha = .05$) was used and no significant difference was found between the study distribution of the two groups, $U = 82.5$, $p = .222$.

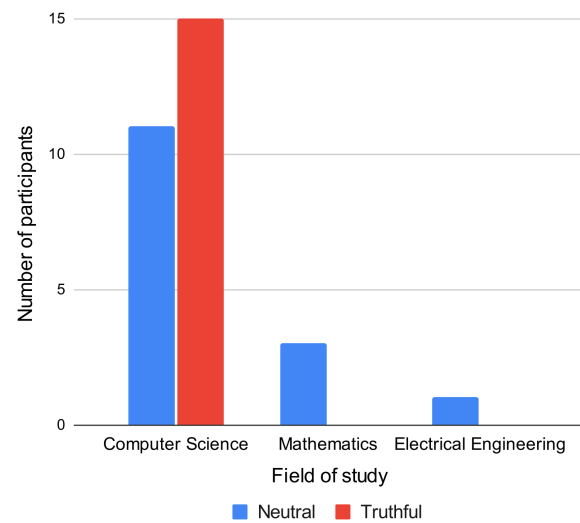


Figure 2: Study distribution per group

Theoretical negotiation knowledge

In both groups, some participants responded that they have theoretical knowledge in negotiations (see Figure 4). The data is not normally distributed, so a Mann-Whitney U test ($\alpha = .05$) was used to show that there is no significant difference between the number of people with negotiation knowledge in the two groups, $U = 67.5$, $p = .064$.

Experience with Pocket Negotiator

As can be seen in Figure 5 most of the participants reported that they have never seen or used the Pocket Negotiator tool before. Due to the data not having a normal distribution, a Mann-Whitney U test ($\alpha = .05$) was used on it and it shows no significant difference between the two groups, $U = 90$, $p = .362$.

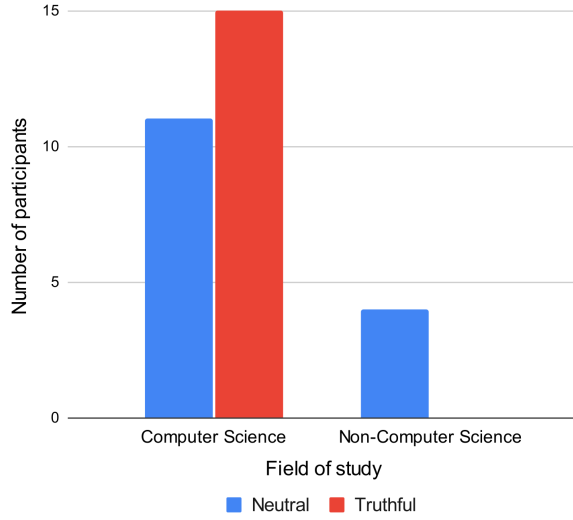


Figure 3: Computer science students distribution per group

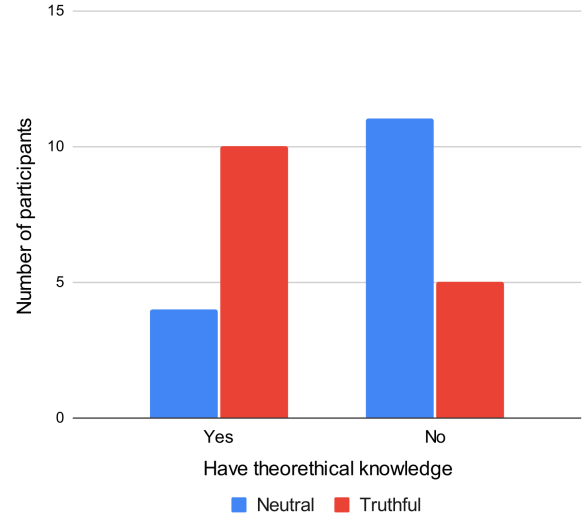


Figure 4: Participants with theoretical negotiation knowledge per group

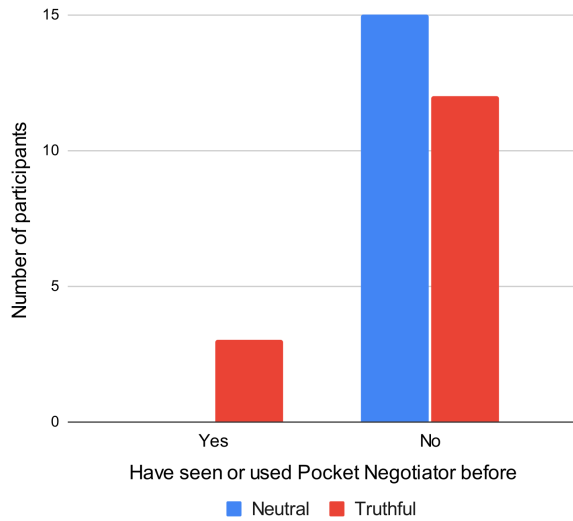


Figure 5: Participants that have seen or used Pocket Negotiator before per group

Propensity to trust AI

The participants were also asked about their propensity to trust AI on a 5-point Likert scale. The group with a neutral explanation had a score mean of 3.2 and a standard deviation of 1.15, while the group with a truthful explanation, had a mean of 3 and a standard deviation of .76 (see Figure 6). The data set is not normally distributed, so again a Mann-Whitney U test ($\alpha = .05$) was used. The results show that there is no significant difference between the propensity to trust AI between the two groups, $U = 101$, $p = .646$.

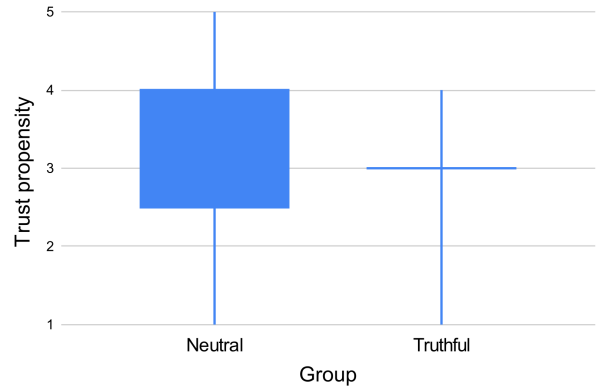


Figure 6: Distribution per group of propensity to trust AI

4.2 Trust analysis

Analysis of the experiment results was used to show whether the research hypothesis holds, and since the questionnaire used to measure the trust level of the participants was unchanged, there was no further need for its validation. To decide whether a parametric or a non-parametric statistical test needed to be done, a Shapiro-Wilk test ($\alpha = .05$) [20] was done first to check each group's data for normality. The results showed the data not having a significant departure from normality, $W_{neutral} = .88$, $p_{neutral} = .054$, and $W_{truthful} = .95$, $p_{truthful} = .558$. Therefore a parametric test with assumptions for normal distribution could be used. In the end, because of the presence of two separate categorical groups, a two-tailed independent samples t -test ($\alpha = .05$) [8] was used.

To quantify the questionnaire answers, each participant's trust score was calculated by taking the sum of their answers with the option values ranging from 1 (strongly disagree) to 5 (strongly agree). This resulted in the first (neutral explana-

tion) group of 15 participants having a median of 22.53, and a standard deviation of 6.12, while the second (truthful explanation) group of 15 participants having a median of 24.86 and a standard deviation of 4.93. The distribution between the two groups can be seen in Figure 7 with a higher trust score interpreted as higher trust in the negotiation agent.




Figure 7: Distribution per group of the trust scores

The newly calculated trust values were used to perform the t -test ($\alpha = .05$). Its results show that there is no significant difference between the trust of the group provided with a "neutral" explanation and the group with a "truthful" one, $t(28) = 1.15$, $p = .26$.

5 Responsible Research

5.1 HREC

Due to the human participation in the experiment, approval was requested from the Human Research Ethics Committee (HREC) of TU Delft before the study was started. Because the experiment does not involve vulnerable groups, health risks, third parties, etc., it is considered as Minimal Risk, and approval was granted. The only factor that is considered to pose a moderate risk to the participants is the collection and management of their data, which is further elaborated upon in the next section.

5.2 Privacy and data

Since the study uses participants' answers to a questionnaire, their privacy and anonymity must be ensured. For this purpose, no personal information that could be used to identify the subjects was collected, and only GDPR-compliant tools were used for the data retrieval and storage. The participants were also informed that the collected anonymous data might be stored in long-term storage for future research use.

5.3 Consent form

To ensure that the participants were informed about the experiment's terms, they were asked to read and fill in a consent form before proceeding with the experiment. The form was based on the template provided by HREC and contained general information about the experiment, the privacy policy,

how data is handled, and the contact details of the responsible researchers. In the end, all of the participants had no issues with the terms and agreed in written form.

5.4 Reproducibility

To ensure that the method of this study is easy to reproduce it is described in detail in Chapter 3.3. In addition, the materials needed for the experiment (i.e. questionnaire and domain) can be found in the appendix of the paper.

6 Discussion

The interpretation of the results resolves around the main research topic:

How truthful explanation of a conceder negotiation agent affects human trust, compared to a neutral one.

To check whether the null hypothesis behind this question holds, in Chapter 4 an independent samples t -test was done on the collected data. What was found is that there is no significant statistical difference in trust between the group provided with a neutral explanation of the conceder negotiation agent and the one provided with a truthful one. Also, various tests were done to confirm that the outcome was not influenced by the participants' distribution between the two groups. Therefore, there is not enough proof against the null hypothesis and it can not be rejected.

One of the possible explanations of the outcome is that people's trust in negotiation agents is influenced relatively more by other factors (e.g. the results of the final outcome) like the case in [19]. Further research is needed to determine whether the end results of a negotiation affect human trust in a negotiation agent.

Another possibility is that the sample size for the two groups is not large enough to reflect the differences in trust. In Figure 7 it can be seen that there is a moderate increase in the trust scores for the group provided with a truthful explanation. Therefore, if a larger sample size follows the same trend, a significant difference might be detected [25]. Due to time constraints, more participants could not be recruited for this research, but a complementary future study, with a larger sample size, may reject the null hypothesis.

Finally, the trust results might be affected by the fact that only students with technical backgrounds were recruited. This potential limitation raises concerns about the generalizability of the study since by exclusively recruiting students with similar education, the findings may not accurately reflect the behaviors of the wider population. To ensure the broader applicability of the study's results, future research could consider expanding the participant pool to include individuals with different educational backgrounds and occupations.

7 Conclusion

This research investigated how a truthful explanation affects human trust in a conceder negotiation agent compared to a neutral one. To analyze this, a between-subject experiment was designed such that half of the participants had to use a conceder negotiation agent, while provided with a truthful

explanation of its strategy, whereas the other half had only a neutral (uninformative) description of it. The main collected data was the participants' subjective trust in the agent after they were done using it in a simulated negotiation scenario.

Analyzing the results showed no significant difference between the trust scores of the two groups. Accordingly, it cannot be concluded whether the two explanations influence people's trust in conceder negotiation agents differently. There are various possible explanations for the outcome, some of them being the presence of more influential trust factors and the technical background of the participants. Also, despite the fact that no significant difference was found, it is notable that the group, provided with a truthful explanation, showed relatively higher trust scores. Further research, with a larger and more diverse set of participants, might reveal the significance of an agent's explanation.

In conclusion, this research provides insights into the presentation of a negotiation agent, introduces a method how to measure the influence its description has on trust, and gives a better understanding of the effects of truthful and neutral agent explanations on people's trust.

8 Acknowledgements

I want to express my gratitude to Kostadin Penchev for providing the data set of the "neutral" group, and also to him, Deniz Tan Hasdemir, Justin Luu, and Kevin van der Werff for the collaboration on the methodology part of this study. Furthermore, I want to thank Carolina Jorge and Myrthe Tielman for their guidance and feedback throughout the entire research.

References

- [1] Reyhan Aydoğan and Catholijn M Jonker. Bidding support by the pocket negotiator improves negotiation outcomes. In *Recent Advances in Agent-Based Negotiation: Applications and Competition Challenges*, pages 52–83. Springer, 2023.
- [2] Tim Baarslag, Michael Kaisers, Enrico Gerding, Catholijn M Jonker, and Jonathan Gratch. When will negotiation agents be able to represent us? The challenges and opportunities for autonomous negotiators. *International Joint Conferences on Artificial Intelligence*, 2017.
- [3] Adrian Bussone, Simone Stumpf, and Dymrna O'Sullivan. The role of explanations on trust and reliance in clinical decision support systems. In *2015 international conference on healthcare informatics*, pages 160–169. IEEE, 2015.
- [4] Béatrice Cahour and Jean-François Forzy. Does projection into use improve trust and exploration? an example with a cruise control system. *Safety science*, 47(9):1260–1270, 2009.
- [5] Jin-Hee Cho, Kevin Chan, and Sibel Adali. A survey on trust modeling. *ACM Computing Surveys (CSUR)*, 48(2):1–40, 2015.
- [6] Filip Karlo Došilović, Mario Brčić, and Nikica Hlupić. Explainable artificial intelligence: A survey. In *2018 41st International convention on information and communication technology, electronics and microelectronics (MIPRO)*, pages 0210–0215. IEEE, 2018.
- [7] Peyman Faratin, Carles Sierra, and Nick R Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [8] Banda Gerald. A brief review of independent, dependent and one sample t-test. *International journal of applied mathematics and theoretical physics*, 4(2):50–54, 2018.
- [9] Ella Glikson and Anita Williams Woolley. Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2):627–660, 2020.
- [10] Taehyun Ha, Sangyeon Kim, Donghak Seo, and Sangwon Lee. Effects of explanation types and perceived risk on trust in autonomous vehicles. *Transportation research part F: traffic psychology and behaviour*, 73:271–280, 2020.
- [11] Koen Hindriks, Catholijn M Jonker, Sarit Kraus, Raz Lin, and Dmytro Tykhonov. Genius: negotiation environment for heterogeneous agents. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 1397–1398, 2009.
- [12] Robert R Hoffman, Shane T Mueller, Gary Klein, and Jordan Litman. Metrics for explainable ai: Challenges and prospects. *arXiv preprint arXiv:1812.04608*, 2018.
- [13] Catholijn M Jonker, Reyhan Aydoğan, Tim Baarslag, Joost Broekens, Christian A Detweiler, Koen V Hindriks, Alina Huldgren, and Wouter Pasman. An introduction to the pocket negotiator: a general purpose negotiation support system. In *Multi-Agent Systems and Agreement Technologies: 14th European Conference, EUMAS 2016, and 4th International Conference, AT 2016, Valencia, Spain, December 15-16, 2016, Revised Selected Papers*, pages 13–27. Springer, 2017.
- [14] Catholijn M Jonker, Koen V Hindriks, Pascal Wiggers, and Joost Broekens. Negotiating agents. *AI Magazine*, 33(3):79–79, 2012.
- [15] John D Lee and Katrina A See. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80, 2004.
- [16] Fernando Lopes, Michael Wooldridge, and Augusto Q Novais. Negotiation among autonomous computational agents: principles, analysis and challenges. *Artificial Intelligence Review*, 29:1–44, 2008.
- [17] Aniek F Markus, Jan A Kors, and Peter R Rijnbeek. The role of explainability in creating trustworthy artificial intelligence for health care: a comprehensive survey of the terminology, design choices, and evaluation strategies. *Journal of Biomedical Informatics*, 113:103655, 2021.

- [18] Patrick E McKnight and Julius Najab. Mann-whitney u test. *The Corsini encyclopedia of psychology*, pages 1–1, 2010.
- [19] Andrea Papenmeier, Gwenn Englebienne, and Christin Seifert. How model accuracy and explanation fidelity influence user trust. *arXiv preprint arXiv:1907.12652*, 2019.
- [20] Samuel Sanford Shapiro and Martin B Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611, 1965.
- [21] Ofir Turel and Yufei Yuan. You can’t shake hands with clenched fists: potential effects of trust assessments on the adoption of e-negotiation services. *Group Decision and Negotiation*, 17:141–155, 2008.
- [22] Daniel Ullman and Bertram F Malle. What does it mean to trust a robot? steps toward a multidimensional measure of trust. In *Companion of the 2018 acm/ieee international conference on human-robot interaction*, pages 263–264, 2018.
- [23] Alicia Nahmad Vazquez. Evaluating team fluency in human-industrial robot collaborative design tasks. In *Computer-Aided Architectural Design. Design Imperatives: The Future is Now: 19th International Conference, CAAD Futures 2021, Los Angeles, CA, USA, July 16–18, 2021, Selected Papers*, pages 378–402. Springer, 2022.
- [24] Oleksandra Vereschak, Gilles Bailly, and Baptiste Caramiaux. How to evaluate trust in ai-assisted decision making? a survey of empirical methodologies. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2):1–39, 2021.
- [25] Elise Whitley and Jonathan Ball. Statistics review 4: sample size calculations. *Critical care*, 6:1–7, 2002.
- [26] Colin Richard Williams. *Practical strategies for agent-based negotiation in complex environments*. PhD thesis, University of Southampton, 2012.
- [27] Feiyu Xu, Hans Uszkoreit, Yangzhou Du, Wei Fan, Dongyan Zhao, and Jun Zhu. Explainable ai: A brief survey on history, research areas, approaches and challenges. In *Natural Language Processing and Chinese Computing: 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9–14, 2019, Proceedings, Part II* 8, pages 563–574. Springer, 2019.

A Negotiation domain

A.1 Scenario description

To introduce the negotiation scenario the following description is provided to the experiment participant:

You and your best friend are hungry and want to make something and share it for lunch. You both like pizza and there is pizza dough already in the fridge. You need to divide the tasks on making the pizza and discuss how much of the pizza each of you would eat.

A.2 Negotiation issues

The negotiation domain consists of the following issues:

- Received pizza share
- Wash ingredients
- Chop ingredients
- Roll out dough
- Assemble pizza topping
- Wash used utensils

The possible values for the “Received pizza share” issue are 33%, 50%, and 66%. For the rest of the issues, the options are “me”, “you”, and “both”.

B Trust questionnaire

The questionnaire used to determine trust is the one recommended in [12]. It is a 5-point Likert scale with the following questions:

1. I am confident in the [tool]. I feel that it works well.
2. The outputs of the [tool] are very predictable.
3. The tool is very reliable. I can count on it to be correct all the time.
4. I feel safe that when I rely on the [tool] I will get the right answers.
5. The [tool] is efficient in that it works very quickly.
6. I am wary of the [tool].
7. The [tool] can perform the task better than a novice human user.
8. I like using the system for decision making.

For the study’s experiment “[tool]” was replaced with “negotiation assistant” and the answer to question 6 is reverse-scored for the analysis (i.e. “Agree” is treated as “Disagree” and vice versa).