Opponent Modeling for Virtual Human Negotiators

Zahra Nazari(✉), Gale M. Lucas, and Jonathan Gratch

Institute for Creative Technologies, University of Southern California, Los Angeles, USA
{zahra, lucas, gratch}@ict.usc.edu

Abstract. Negotiation is a challenging domain for virtual human research. One aspect of this problem, known as *opponent modeling,* is discovering what the other party wants from the negotiation. Research in automated negotiation has yielded a number opponent modeling techniques but we show that these methods do not easily transfer to human-agent settings. We propose a more effective heuristic for inferring preferences both from a negotiator's pattern of offers and verbal statements about their preferences. This method has the added advantage that it can detect negotiators that lie about their preferences. We discuss several ways the method can enhance the capabilities of a virtual human negotiator.

1 Introduction

Negotiation is an important and challenging domain for virtual human research. Several efforts have advanced the design of conversational agents that negotiate with human users, both as a means for teaching negotiation skills [1] and as a means to advance the social intelligence of virtual agents [2, 3]. Negotiation engages a wide range of individual and interpersonal skills that are not only crucial for people, but essential for machines that socially engage with humans. In this paper, we present a novel algorithm for advancing one of these skills – opponent modeling.

Understanding what the other party wants is key to successful negotiation [4]. Unfortunately, these must often be inferred from words and actions, and many negotiation contexts, parties may feel they should withhold or misrepresent their true preferences. Opponent modeling is a focus of research in the multi-agent community and various models are proposed. However, models developed in the context of agent-agent negotiations may be inappropriate in human negotiations. In this paper, we illustrate that standard opponent modeling approaches do not perform well on human data and we introduce a heuristic that performs far better.

Opponent modeling is part of a larger effort of ours to create a virtual human platform that teaches negotiation skills (see Fig. 1). The agent allows students to engage with different virtual human role-players across a variety of negotiation problems via natural language and nonverbal expressions of emotion. Opponent modeling plays a crucial role for this agent to (a) understand what the student wants, (b) understand how the agent's own behavior is influencing the student's beliefs, and (c) a metric for quantifying the information being exchanged between parties. For example,

W.-P. Brinkman et al. (Eds.): IVA 2015, LNAI 9238, pp. 39–49, 2015.
DOI: 10.1007/978-3-319-21996-7_4

Fig. 1. The conflict resolution agent and a human user [5].

a student might think they are accurately communicating their preferences but their actual words and deeds might communicate something quite different to the other party, and the difference between the true and inferred models can serve to quantify this discrepancy.

The rest of this paper is organized as follows. Section 2, explains the negotiation setting considered for this paper and an overview of the currently existing models. In Sect. 3, we use a large corpus of human negotiations to address the differences between human-involved and multi-agent settings, and propose alternative models that address these differences. In Sect. 4, results from the existing and proposed models are shown. Finally, Sect. 5 provides a summary and discusses the future work.

2 Preliminaries/Background

We adopt a standard formulation of a bilateral negotiation. Parties must find agreement over set of independent issues. Each issue consists of a set of discrete levels (levels might correspond to attributes of an agreement, such as salary and benefits of a job offer). A negotiation succeeds if both players agree on a level for each issue, at which point each party receives a utility (unknown to the other party). Here, we adopt the conventional assumption of a linear additive utility function in the range of [0, 1]. The value of ω, is the weighted (w_i) sum of assigned levels for each issue (l_i):

$$u(\omega) = \sum_{i=1}^{n} w_i \cdot l_i \tag{1}$$

The weight w_i represents the preference the opponent holds for issue i. The set of all possible deals are known as outcome space (Ω) that is known for both parties and negotiators try to reach an outcome that maximizes their own utility.

An important concept in negotiation is *Pareto efficiency* which is a measure of the quality of a negotiated agreement. A deal is inefficient if it is possible to improve the other party's position without harming one's own. Rational agents should only offer Pareto efficient deals as this maximizes joint gains and increases the chance of a beneficial agreement for oneself. Unfortunately, one needs perfect information about both parties' preference to calculate the efficiency of a deal. Having a good estimate of the opponent's preferences helps the negotiators to make offers that are closer to Pareto optimal and are more likely to be accepted by their opponent.

Several algorithms have been proposed to estimate the opponent' utility function (see [4] for a recent review). These models only use the pattern of offers for their estimations and fall into two main types: Bayesian models and frequency models. Bayesian models generate a set of candidate preference profiles first, and then use certain assumptions about the opponent concessions to update their models. One of the main assumptions in this set of models is that the opponent starts with asking for maximum possible utility and then gradually concedes towards lower utilities. Among Frequency models, N.A.S.H Frequency [4] learns the issue weights based on how often the best value for each issue is offered. Hardheaded [6] learns the issue weights based on how often the level of an issue changes, assuming when a party changes an asked level for an issue frequently, they must assign lower utility for that issue.

Various metrics have been proposed to assess the quality of an opponent modeling approach. In this paper, we use a standard accuracy measure in negotiation contexts called rank distance of the deals [7]. This metric compares the utility of all possible deals in the outcome space (Ω), given the estimated (u'_{op}) and the actual weights (u_{op}), and calculates the average number of conflicts in how deals are ranked using the estimated vs. actual utility function:

$$d_r\left(u_{op}, u'_{op}\right) = \frac{1}{|\Omega|^2} \sum_{\omega \in \Omega, \omega' \in \Omega} c_{<u, <u'}(\omega, \omega') \tag{2}$$

The function c, the *conflict indicator,* takes any pair of deals (ω and ω') and returns 1 if the ranking between the deals changes when calculated by the actual vs. estimated weights; otherwise it returns 0. An opponent modeling approach is considered to be more accurate if it produces a *smaller* rank distance than another approach.

One limitation of rank distance is that it minimized the average number of errors that can arise from the estimated weights but not the severity of these errors. An alternative strategy, known as *minmax regret,* is to minimize the worst-case negative consequences that arise from using the estimated weights. Minmax regret is considered a more conservative and robust procedure for selecting amongst models. To capture this intuition in the opponent modeling context, we propose the following measure.

Max-regret finds the maximum absolute difference between the utility of a deal given actual weights $(u_{op}(\omega))$ and its utility given estimated weights $\left(u'_{op}(\omega)\right)$ over all possible deals (Ω):

$$d_M\left(u_{op}, u'_{op}\right) = \max_{\omega \in \Omega}\left|u_{op}(\omega) - u'_{op}(\omega)\right| \tag{3}$$

As we assume utility has been normalized to the range [0…1], a max-regret of 0.5 indicates that the estimate may be off by half of its maximum possible value. For example, if our opponent could receive $1000 for his best deal, with a max-regret of 0.5 we estimate his payout at only $500, leading us to think we're getting a great deal when we could have extracted greater concessions. We will use both rank distance and max-regret to assess the value of different opponent modeling techniques.

3 Challenges with Modeling Human Negotiation Preferences

Opponent modeling techniques do a reasonable job of inferring the preferences of automated negotiation agents [4], but they may fail to capture the true preferences of human negotiators. A review of the literature on human negotiations emphasizes several differences between how people and automated negotiation agents behave:

1. Automated agents tend to start with the offer that is best to them, and concede monotonically from that point. Human negotiators (especially novice negotiators) tend to start much closer to what they perceive to be a fair offer [8, 9]. As fairness depends on an estimate of the Pareto frontier, which may be incorrect, a 'fair' offer may differ considerably from what is truly fair, adding noise into any algorithm that focuses strictly on offers.
2. Automated agents impose strict mechanisms to structure the negotiation process. Most algorithms assume that parties take turns exchanging complete offers. In contrast, humans are more flexible and usually make partial offers (i.e., offers that specify levels on only a subset of the issues [8].
3. Automated agents can be tediously patient and often exchange thousands of offers before reaching a deal, whereas humans may only exchange a few offers before concluding a negotiation [10].
4. Most importantly, people communicate via language whereas most agents only communicate via a pattern of offers. Language allows people to distinguish between preferences and deals (e.g., I'd prefer the best stereo possible for this car but can only afford $10,000). Of course, people often lie in negotiations, so there may be differences between what people say and do.

In this section, we describe a large corpus of face-to-face negotiation data, illustrate that humans indeed violate the assumptions of current opponent modeling techniques, and propose some simple heuristics to infer the preferences of human negotiators.

3.1 Corpus

Participants: 113 same-sex dyads (38 female dyads, 75 male dyads) were recruited from craigslist for participation in this study. Participants were paid for completing the study, and based on the outcome they achieved in the negotiation, were given additional entries into a lottery for a cash prize. They reported a mean age of 26.83 (SD = 12.77), and 41.6 % racially identified as African-American, 32.3 % as Caucasian, 8.8 % as Hispanic, 8.0 % as other, 6.2 % as Asian, 3.1 % as Native American/Hawaiian.

Design: Each dyad engaged in a negotiation task, in which they role-played characters who are negotiating over six antique items (three crates of LP records, two art deco lamps, and one art deco painting), which have varied levels of value to them. They are told that their task is to decide how to divide up these six items with another participant, and if they fail to reach agreement, that they will receive the number of lottery tickets that they would have received for one of their highest value items.

Table 1. Participant's preferences for different issues across task and side.

Task	Side	Records	Lamps	Painting
Distributive	A	High (30)	Moderate (15)	Low (5)
	B	High (30)	Moderate (15)	None (0)
Integrative	A	High (20)	Moderate (10)	Low (5)
	B	Moderate (10)	High (30)	None (0)

Dyads were randomly assigned to either a distributive or integrative task: the items either have the same level of value to each participant, or different levels of value to each partner (Table 1). Specifically, in the distributive task, for both participants, the records are the highest value items, and the lamps have moderate value to them, whereas in the integrative task, for participant A, the records are the highest value items and the lamps have moderate value, but for participant B, the lamps have the highest value and the records have moderate value. In both conditions, the painting is of little value to participant A, and of no value at all to participant B.

Of the original 226 participants, 8 were removed for non-compliance with procedures (e.g., obvious intoxication, beginning the negotiation before reading all instructions), 5 dyads were excluded for failing to reach agreement in the negotiation, and an additional 28 participants were excluded from analyses because, after the study, they failed to accurately report the preferences to which they had been assigned.

All dialogues were manually transcribed and annotated with several different dialog acts. Two dialog acts are relevant for the current paper (the complete annotation scheme is described here [11]). *Offers* correspond to statements about a division of items. For example, the statement "How about if I take two records and you take both lamps?" is a partial offer. It is also partial in two distinct senses as (a) fails to mention the painting, and (b) it requests two records while leaving the third record unspecified.

Preference-assertions correspond to statements about an individual issue such as "I like records" or "I don't like paintings." The full annotation scheme makes several distinctions in preferences (e.g., I-like-best, i-like, i-might like, etc.), however, for the purpose of this paper, we collapse these into two broad classes (I-like-ITEM, and I-don't-like-ITEM) as more subtle distinctions would be problematic for automatic language recognition. Thus, to summarize, preference-assertions indicate an issue (e.g., records) and a sentiment (positive or negative) expressed toward that issue. If multiple issues are stated in a single utterance ("I like the records and the lamps"), these are represented as two preference assertions.

Preference-assertions are annotated for their veracity. False statements do not necessarily mean the participant is lying (they could have misspoken), however it was clear that many participants attempted to misrepresent their true preferences.

3.2 Testing the Assumptions of Standard Opponent Models

In this section, we report some statistics from the corpus explained above as empirical evidence illustrating the differences between human and agent negotiators.

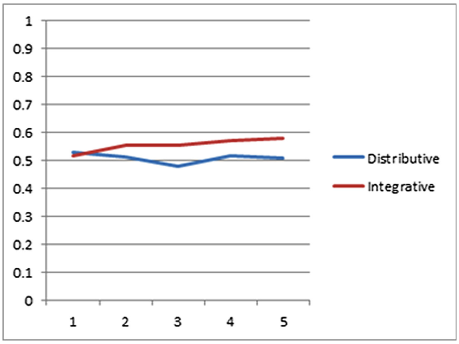

Fig. 2. Concession rates: y-axis is relative amount of total value claimed by each offer by round.

1. Participants did not start high and concede over time, as assumed by many opponent modeling methods. Figure 2 shows the average strength of offers by round and by the type of negotiation. A value of 1.0 would indicate that the participant made the high-utility offer they could obtain. Instead, they only asked for about half of this on the first round. Rather than conceding over time, participants, if anything, show the opposite trend. This can be explained as participants in the integrative condition discover a win-win solution is possible (i.e., they revise their demands upward as they form a more accurate model of their opponent).
2. The complete offer protocol adopted in automated agent frameworks was indeed violated. Of the 522 offered in the dataset, 370 (70 %) were partial.
3. In contrast to the thousands of offers exchanged by automated agents, human participants exchanged only 5.8 offers on average per negotiation.
4. People talked extensively about their preferences. People made 9.9 statements about their preferences on average in each negotiation (as compared to 5.8 explicit offers). Therefore, about twice as much as the information relevant for preference standard opponent modeling methods discards modeling. Most of this information was truthful, although 18 % of participants misstated their preferences, which emphasizes a challenging in using human data.

The observed differences between agent and human negotiators clearly emphasize the need for models that are better suited to the characteristics of human negotiators. We approached this problem by proposing three straightforward heuristic models: the *Issue-ratio heuristic* uses the sequence of offers to estimate the preference profile, the *Issue-sentiment heuristic* derives an estimate from explicit preference assertions, and the *Offer/Sentiment heuristic* simply uses the mean of these two models to integrate information in both offers and language. Note these are the simplest models we could imagine and future work could improve on these using machine-learning methods.

Issue-Ratio Heuristic: Rather than looking at concession rates, this heuristic examines each issue separately and assumes (1) if an issue is important, the participant will offer a greater percentage of the possible value to themselves, and (2) if an issue is important, they will include it more often in their partial offers. We realize these two intuitions in the following metric. If an issue (i) is discussed in an offer (k), it is assumed to have two parts: which level of that issue was claimed for self (l_k) and which level was assigned to the opponent (l'_k). The heuristic estimates each issue weight (w_i) for a participant by comparing the average level claimed for self ($\bar{l}_k$) to the average level offered to the opponent ($\bar{l'}_k$) across all offers made by that participant:

$$w_i = \frac{\bar{l}_k}{\bar{l'}_k} \tag{4}$$

Issue-Sentiment Heuristic: Participants made a large number of explicit statements about their preferences (e.g., "I like the records the most"). Although the trustworthiness of these expressions remains unknown, they could be considered as a valuable source of information for our preference models. We propose a very simple heuristic based merely on counting the number of times a preference is expressed towards an issue. More precisely, every time a positive preference is asserted towards an item ("I like the painting"), we add one to a weight associated with the issue. Every time a negative preference is expressed ("I don't really care for the painting"), we subtract one to the weight associated with the issue. All weights are normalized to compute a set of weights that are comparable to the weights derived from offers:

$$w_i = |P_i| - |N_i| \tag{5}$$

P_i is the set of all positive statements a participant asserted about issue i, and N_i is the list of all negative statements that participant made about issue i.

Offer/Sentiment Heuristic: As negotiators reveal preferences both through their offers and through explicit preference statements, a heuristic that incorporates both sources of information might perform best. A simple way to combine these two factors is to average the weights that arise from these two estimates. Our Offer/Like heuristic accomplishes just this. If no information is given by one source (e.g., the participant makes offers but does not make any preference assertions), we use just one estimator. If both estimators produce valid weights, we simply take the mean of the two estimates.

Offer/Sentiment Discrepancy: Finally, the fact that we have two sources of preference information (through offers and through explicit preference statements) it becomes possible to examine the discrepancy between these two estimates. This might have potential for lie detection. For example, if a negotiator implies one thing through his words but makes offers inconsistent with these preferences; we would have reason to be suspicious. Alternatively, this might suggest the negotiator is confused. Either way, such discrepancies are important to note and could be valuable to a pedagogical negation agent. We define Offer/Like discrepancy as the rank distance between the weights estimated via offers and the weights estimated via preference statements.

4 Evaluation

Here, we contrast the performance of our heuristics against the state-of-the art in opponent modeling techniques. We compare the performance of all heuristics on the dataset discussed above and report their accuracy in terms of rank-distance and max-regret. We also considered if Offer/Like discrepancy could be used for lie detection.

4.1 Heuristic Accuracy

We chose HardHeaded and N.A.S.H. Frequency as representatives of the state-of-the-art in opponent modeling. Hardheaded won the 2011 automated negotiation agents' competition and had one of the highest accuracies in modeling opponents (see [4]). N.A.S.H. Frequency did not perform as well in practice but represents the standard Bayesian view on how to model opponents.

Hardheaded is designed to solve a more general problem than was faced by our human participants, so we also created a *HardHeaded_Modified* to allow a more direct comparison. In the human negotiation task, participants were given knowledge about the ranking of levels within an issue (i.e., 2 lamps had more value than 1 lamp which had more value than no lamps). HardHeaded estimates this ranking from the pattern of offers. Thus, to create a fair comparison, we created a version of the algorithm with these parameters were fixed to their true values. In this way, both humans and the algorithms started the negotiation with the same knowledge about the task.

Each model was given the sequence of offers and preference assertions that were produced by each participant in the corpus (preference assertions were only used by the Issue-Sentiment model and the hybrid Offer/Sentiment model). For each participant, we calculate the rank distance of the deals, and Max-regret between the estimated weights produced by the model and the true weights provided to the participant. We also include a model that produces random weights as a point of comparison. Results are shown in Fig. 3. There were significant differences between models (Rank distance: $F(4, 176) = 85.242$, $p < .001$), Max-regret: $F(4, 176) = 44.148$, $p < .001$). As predicted, the existing models did not fare well on human data and their performance was close to random. Issue-sentiment and issue-ratio heuristics, however, made significantly better estimations in terms of both rank distance of the deals, and Max-regret measures. The composite Offer/Sentiment heuristic made the best estimation for the participants' preferences. These differences are significant ((Offer/Sentiment)/issue-ratio, Rank distance: $t(1, 179) = 5.006$, $p < .001$), Max-regret: $t(1, 179) = 3.195$, $p = .002$) ((Offer/Sentiment)/Issue-sentiment, Rank distance: $t(1, 179) = 7.656$, $p < .001$), Max-regret: $t(1, 179) = 7.21$, $p < .001$).

4.2 Lie Detection

We next tested if the Offer/Sentiment heuristic could serve as the basis of a lie detector. We compared the discrepancy of the weights estimated by the Issue-Ratio heuristics

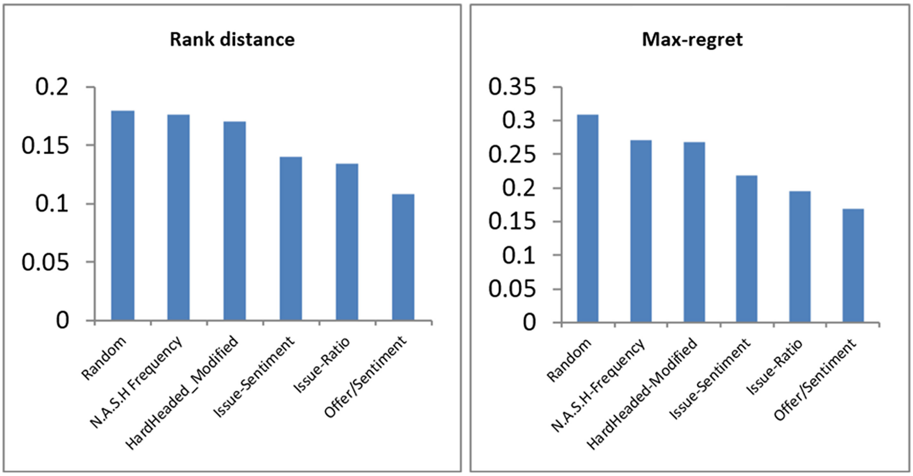

Fig. 3. Performance of various models based on two measures, rank distance of the deals, and max-regret [Note that the smaller our measures are the better that prediction is]

with the weights estimated by the Offer-Sentiment heuristic. This discrepancy was significantly positively correlated with the number of false preference assertions made by each participant ($r = 0.174$, $p = 0.03$), which confirms our hypothesis that agents can use the discrepancy between language channel and the pattern of offers to detect suspicious information. This is especially promising as both heuristics are quite simplistic and we expect even stronger correlations with more improved estimators.

5 Conclusion

Opponent preference modeling is an important skill for a virtual agent who intends to have efficient negotiations with humans or engage with them as a trainer. Despite the wide attention this area of research has received in agent-agent frameworks, negligible work was done for human involved settings. Using a corpus of human-human negotiations, we addressed the main differences between human and agent negotiators. Human negotiators violate many of the assumptions of automated negotiation agents, rendering these existing models useless for our purpose. Human participants made frequent use of partial offers, whereas agents typically assume complete offers; humans started with fair offers, whereas agents assume initial offers will be tough. Most importantly, humans use language to communicate their preferences whereas agents only focus on the pattern of offers. As a result of these differences, state-of-the-art opponent modeling methods perform close to chance on human data.

We proposed a new set of simple methods to estimate opponent preferences in human negotiations. Our first heuristic used this pattern of behavior to capture people's preferences and showed a significantly better performance than other existing models. Our second heuristic only used participants' preference assertions and performed as good as the offer heuristic. We used the average between offer based and the assertion based weight estimations as our third model performed the best of all. More interesting, the Offer/Sentiment heuristic allowed a simple form of lie detection.

These heuristics have several potential uses for a virtual human negotiator. By having a more accurate estimate of the human's preferences, the virtual agent can make more efficient offers. The agent can also use the heuristic on its own offers to estimate how its language or offers might influence the human's belief about the agent's own preferences. This can guide dialog strategies and even personality differences. For example, a cooperative agent could select speech acts that make its own preferences transparent, whereas a Machiavellian agent might misrepresent its own preferences in order to gain strategic advantage. Similarly, discrepancies between the human's words and offers could serve to detect when the human is attempting to exploit the agent. Finally, these heuristics can serve as metrics for quantifying the information people are exchanging in negotiations. This could have potential use in psychological research on emotion to give insight into the negotiation process. It could also have value for pedagogical agents for assessing the performance of student negotiators.

Future work will proceed on several fronts. First, we hope to improve upon these heuristics using machine-learning approaches. The current heuristics are very simple and we expect substantial improvements are possible, both in how to estimate weights from issues and language separately, but also how to best combine these two sources of information. Second, we plan to use active probing to improve our estimates. In the current paper, we passively observed a sequence of assertions and offers, but in a real negotiation, negotiators can select offers or ask about preferences in order to reduce uncertainty about the opponent's preferences [12]. Finally, we will incorporate these methods into a virtual human negotiator being developed at our lab [5].

Acknowledgement. This research was supported by the US Army. The content does not necessarily reflect the position or the policy of any Government, and no official endorsement should be inferred.

References

1. Kim, J.M., Hill, R.W., Technologies, C., Lane, H.C., Forbell, E., Core, M., Marsella, S., Pynadath, D., Hill Jr, R.W., Durlach, P.J., Hart, J.: BiLAT: A game-based environment for practicing negotiation in a cultural context. Int. J. Artif. Intell. Educ. **19**, 289–308 (2009)
2. Traum, D.R., Marsella, S.C., Gratch, J., Lee, J., Hartholt, A.: Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In: Prendinger, H., Lester, J.C., Ishizuka, M. (eds.) IVA 2008. LNCS (LNAI), vol. 5208, pp. 117–130. Springer, Heidelberg (2008)
3. Rosenfeld, A., Segal-halevi, E., Drein, O., Kraus, S., Zuckerman, I.: NegoChat: a chat-based negotiation agent. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, International Foundation of Autonomous Agents and Multi-agent Systems (2014)
4. Baarslag, T., Hendrikx, M., Hindriks, K., Jonker, C.: Predicting the performance of opponent models in automated negotiation, vol. 2, pp. 59–66. In: Proceedings of 2013 IEEE/WIC/ACM International Conferences on Intelligent Agent Technologies (IAT 2013) (2013)

5. Gratch, J., Devault, D., Lucas, G., Marsella, S.: Negotiation as a challenge problem for virtual humans. Paper presented at the 15th international conference on intelligent virtual agents, Delft (2015)
6. van Krimpen, T., Looije, D., Hajizadeh, S.: HardHeaded. In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) Complex Automated Negotiations: Theories, Models, and Software Competitions. SCI, vol. 435, pp. 225–230. Springer, Heidelberg (2012)
7. Hindriks, K.V., Tykhonov, D.: Towards a quality assessment method for learning preference profiles in negotiation. In: Ketter, W., La Poutré, H., Sadeh, N., Shehory, O., Walsh, W. (eds.) AMEC 2008. LNBIP, vol. 44, pp. 46–59. Springer, Heidelberg (2010)
8. Kelley, H.: A classroom study of dilemmas in interpersonal negotiations. Strategic Interaction and Conflict, pp. 49–73. Institute of International Studies, University of California, Berkeley (1966)
9. Fehr, E., Schmidt, K.M.: The economics of fairness, reciprocity and altruism - experimental evidence and new theories. Handb. Econ. Giving Altruism Reciprocity **1**(6), 615–691 (2006)
10. Lin, R., Kraus, S.: Can automated agents proficiently negotiate with humans? Commun. ACM **53**, 78 (2010)
11. Devault, D., Mell, J., Gratch, J.: Toward natural turn-taking in a virtual human negotiation agent. In: Paper Presented at the AAAI Spring Symposium on Turn-taking and Coordination in Human-Machine Interaction. Stanford (2015)
12. Boutilier, C., Sandholm, T., Shields, R.: Eliciting bid taker non-price preferences in (combinatorial) auctions, pp. 204–211. In: Proceedings of Nineteenth National Conference on Artificial Intelligence (2004)