A Survey of Opponent Modeling Techniques in Automated Negotiation

(JAAMAS Extended Abstract)

Tim Baarslag
University of Southampton
United Kingdom
T.Baarslag@soton.ac.uk

Mark J.C. Hendrikx
Delft University of Technology
The Netherlands
M.J.C.Hendrikx@tudelft.nl

Koen V. Hindriks
Delft University of Technology
The Netherlands
K.V.Hindriks@tudelft.nl

Catholijn M. Jonker
Delft University of Technology
The Netherlands
C.M.Jonker@tudelft.nl

1. INTRODUCTION

Negotiation is a process in which parties interact to settle a mutual concern to improve their status quo. Traditionally, negotiation is a necessary, but time-consuming and expensive activity. Therefore, in the last two decades, there has been a growing interest in the *automation* of negotiation.

One of the key challenges for a successful negotiation is that usually only limited information is available about the opponent. Although sharing private information can result in mutual gains, negotiators often avoid this to prevent exploitation. This problem can be partially overcome by deriving information from the opponent's actions. Exploiting this information to learn aspects of the opponent is called *opponent modeling*. Creating an accurate opponent model is a key factor in improving the quality of the outcome and can further increase the benefits of automated negotiation.

Despite the advantages of opponent modeling and two decades of research, there is no recent study that provides an overview of the field. Therefore, in order to stimulate the development of efficient future opponent models, and to outline a research agenda, we provide an overview of existing opponent models in bilateral negotiation [2]. As our main contributions, we classify opponent models using a comprehensive taxonomy and provide recommendations on how to select the best model depending on the negotiation setting.

2. LEARNING ABOUT THE OPPONENT

A bilateral negotiation may be viewed as a two-player game of incomplete information. An opponent model is then simply an abstracted description of a player (and/or their behavior) during the game. In negotiation, opponent modeling often revolves around three concerns: what does the opponent *want*; what will the opponent *do*; and what *type* of player is the opponent. These considerations are often highly related; for example, to understand the opponent's behavior, an agent first needs to know what the opponent desires.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

An important aspect in which opponent models differ is their time of creation. *Offline* opponent models are constructed before the negotiation using historical data. *Online* opponent models are constructed during the negotiation by interpreting the exchange of offers. A major challenge for online models are real-time deadlines, which impose a restriction on the time available for maintaining the model.

Even though there are large differences between opponent models, there is a common set of motivations for using them:

1. **Minimize negotiation cost.** It costs time and resources to negotiate and as a consequence, early agreements are often preferred. Learning can aid in identifying promising bids that lead to swift agreements.
2. **Adapt to the opponent.** Agents can adapt their behavior according to their opponent; e.g., by estimating the deadline or reservation value in an attempt to press for an outcome the opponent will ultimately settle for.
3. **Reach win-win agreements.** In a cooperative environment, agents aim for a fair result. An estimate of the opponent's preference profile can aid in identifying mutually beneficial outcomes.

We found that existing work can fulfill these goals by learning a combination of *four* opponent attributes listed below.

Acceptance Strategy

The decision to accept an offer is made by the *acceptance strategy* of a negotiating agent. Upon acceptance of an offer, the negotiation ends in agreement; otherwise, the agents continue exchanging offers. Learning an opponent's acceptance strategy is potentially of great value as it can help to find the deal with the highest utility that is still acceptable for the opponent.

In negotiations about a single, quantitative issue (such as the price of a service), where the opponent's have opposing preferences that are publicly known, *estimating the reservation value* is sufficient to determine all acceptable bids. The reservation value can be learned by extrapolating the opponent's concessions, for example by applying Bayesian learning. Alternatively, a model may assume that the opponent uses a particular decision function of which the unknown variables can be estimated using non-linear regression.

In multi-issue negotiations, we can *estimate the probability of acceptance* for every possible outcome. By keeping track of what offers were offered and accepted in previous negotiations, an agent can estimate the probability that a bid will be accepted. As it is unlikely that such an estimate can be derived for all possible bids, regression methods can be applied to determine the acceptance probability for the entire outcome space.

Deadline

The *deadline* of a negotiation refers to the time before which an agreement must be reached which is better than the best alternative for each party. Each agent can have their own private deadline, but it is also common for the negotiation deadline to be shared. The deadline may be specified as a maximum number of rounds, or alternatively as a real-time cutoff point.

Private deadlines are important to learn about, as an agent is likely to concede strongly near the deadline to avoid non-agreement. Because of the strong connection with the reservation value, most of the procedures to learn the reservation value can also be applied here.

Preference Profile

The *preference profile* of an agent represents the private valuation of possible negotiation outcomes. Learning the preference profile assists in locating mutually beneficial outcomes and recognizing potential for meaningful concessions.

Four approaches have been used so far to estimate the opponent's preference information:

1. **Importance of the issues.** It is often easier to estimate the *weight* of all issues under negotiation, rather than the preference over all outcomes. The idea is to analyze the opponent's concessions, assuming that stronger concessions are made on issues that are valued less.
2. **Classify the negotiation trace.** Given the opponent's negotiation actions, we can determine which opponent type is most likely, and subsequently apply a classification algorithm to categorize preferences of the opponent.
3. **Aggregate negotiation data.** When offline data is available, we can derive the opponent's preference profile from a large database of negotiation traces from similar – but not identical – opponents.
4. **Importance of outcomes.** A popular technique is the frequency analysis heuristic. The main idea is that preferred values of an issue are offered relatively more often in a negotiation trace. For the issue weights, the opposite holds: an issue is likely unimportant if its value changes often.

Bidding Strategy

The *negotiation strategy* determines an agent's offer in any given negotiation state. Strategies may range from simple time-dependent concessions to complex decision-making that depends on the opponent's behavior. Learning the opponent's bidding strategy allows an agent to anticipate and manipulate the opponent's behavior.

An agent is generally unaware of the opponent's exact strategy, but it might have knowledge about the type of strategy used. If such knowledge can be captured in a closed-form formula with unknown variables, *regression analysis* can be applied.

When limited knowledge is available about the opponent's type, *time series forecasting* is an alternative, which is a technique to extrapolate a set of observations that is sequentially ordered in time. In negotiation, the time series typically consists of the utilities of offers received. Learning the opponent's bidding strategy then boils down to creating a forecast of future offers, using a set of statistical techniques and smoothing methods.

3. DISCUSSION AND CONCLUSION

Constructing an opponent model may be viewed as a classification problem, where the type of the opponent needs to be determined from a range of possibilities. In combination with *opponent classification*, there are two other opponent model categories: models that learn *what the opponent wants*, in terms of its reservation value, deadline, and preferences, and secondly, models that learn *what the opponent will do*, in terms of the bidding and acceptance strategy.

A large number of opponent modeling techniques exist in literature. For each learned opponent attribute, we have identified key approaches that are currently being applied. A natural question that arises with regard to agent design is: which of the depicted modeling techniques is *best*? To date, answering this question is unfeasible, as most authors evaluate opponent models in their own specific setting using different quality measures. A main direction for future work is therefore to outline a procedure for finding the best opponent modeling technique given the particular circumstances.

A fair procedure should consist of two components: a balanced experimental setup to evaluate the model under different conditions, together with an unbiased set of *quality measures* [1]. *Performance measures* are the most commonly used quality measure; they determine the performance gain when a negotiation strategy is supplemented by a learning method, for example the increase in average utility. *Accuracy measures* quantify how closely the model resembles reality, e.g. by calculating the correlation between the actual and estimated utilities of all outcomes. Both approaches augment each other, and in [2] we discuss for each of the four opponent attributes that are learned which measures to apply. In addition, for each of the three reasons for using an opponent model (c.f. Section 2) we recommend a set of corresponding performance measures.

Consistently applying the right negotiation benchmarks would greatly improve comparability of results in negotiation literature, and would provide insight in possible improvements to existing models and how they can be combined to augment each other. Our work serves as a starting point towards finding the best opponent modeling techniques, and thereby the transition from theory to practice.

REFERENCES

[1] T. Baarslag, M. J. Hendrikx, K. V. Hindriks, and C. M. Jonker. Predicting the performance of opponent models in automated negotiation. In *IEEE/WIC/ACM*, volume 2, Nov 2013.

[2] T. Baarslag, M. J. Hendrikx, K. V. Hindriks, and C. M. Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, pages 1–50, 2015.