

Designing an Automated Negotiator: Learning What to Bid and When to Stop (Doctoral Consortium)

Tim Baarslag
Interactive Intelligence Group
Delft University of Technology
Mekelweg 4, Delft, The Netherlands
T.Baarslag@tudelft.nl

ABSTRACT

In this paper, the ongoing research of the author on designing an automated negotiator is described. One of the key challenges of designing a successful negotiation agent is that usually only limited information is available about the other party. Therefore, we need to combine various learning techniques to decide what offers to make, and when to accept. Our goal is to investigate techniques for developing a versatile automated negotiator that can effectively conduct negotiations in an incomplete information setting.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*intelligent agents, multi-agent systems*

General Terms

Algorithms, Experimentation, Performance, Theory

Keywords

Negotiation; Artificial Intelligence; Machine Learning

1. INTRODUCTION

Negotiation is an important activity in human society, and is studied by various disciplines, ranging from economics and game theory [15], to electronic commerce [9], social psychology, and artificial intelligence [10, 12].

Traditionally, negotiation is a necessary, but also time-consuming and expensive activity. Therefore, in the last decades there has been a large interest in the *automation* of negotiation [7, 9, 10], for example in the setting of e-commerce [14]. This interest is fueled by the promise of automated agents eventually being able to negotiate on behalf of human negotiators.

One of the key challenges for a successful negotiation is that usually only limited information is available about the other party. Despite the fact that sharing private information can result in mutual gains, negotiators are often unwilling to share this information to avoid exploitation. This

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

problem can be partially overcome by *learning* from the offers that are exchanged during the negotiation. This can be done in many ways, ranging from strategy prediction (“*what will the opponent do?*”) to preference estimation (“*what does the opponent want?*”).

What is learnt about the opponent can then be used to improve the decision of what offers to make, and what offers to accept. For example, in order to send out bids that are beneficial to the opponent (thereby increasing the chances of reaching a deal), it is necessary to estimate the opponent’s preferences. Conversely, when the agent is presented with an offer by the opponent, strategy prediction may help in deciding what to do next. The agent has to make a choice between accepting the offer that is currently on the table, or rejecting it and continuing the negotiation; both options involve an inherent risk, and predicting the opponent’s future offers is essential for making the right decision.

2. APPROACH

Every year, automated negotiation agents are improving in various ways, and there is now a large body of negotiation strategies available, all with their unique strengths and weaknesses (for an exposition, see [1, 7]). For example, some agents are able to predict the opponent’s preferences very well, while others focus more on having a sophisticated bidding strategy. Naturally, we would like to learn from previous approaches, and improve upon them, in order to create a versatile and generic negotiating agent.

However, most agents are designed to function in very different environments, so there is no straightforward way to compare different approaches, let alone combine them. Therefore, our approach has been as follows: first of all, we created a generic negotiation environment called GENIUS [13], which can fully support a diversity of different negotiation protocols, scenarios, and agents. Second, we amended the GENIUS repository with various existing agents [8], scenarios [11], and protocols [15] from literature. Additionally, we organized a yearly international negotiation competition (ANAC) [1, 6] to harvest even more strategies and scenarios, and to learn new, improved approaches to effective agent design.

With this in place, we were able to pinpoint additional structure in most agent designs. We identified three main components of a general negotiation strategy; namely a *bidding strategy*, possibly an *opponent model*, and an *acceptance strategy* (BOA) [4]. These BOA components enable

us to do two things: first, they allow us to study the behavior and performance of individual components; and second, they make it possible to systematically explore the space of possible negotiation strategies by recombining different components.

We started with the first part: seeking out the best of each BOA component. We mainly focused on the logical starting points, namely the opponent modeling and acceptance mechanism components.

For opponent modeling, there was no recent overview of the field available, so we conducted a survey of currently existing opponent modeling techniques (to appear, see Section 3). In tandem with surveying the state of the art, we studied the performance of a variety of different opponent models, and we concluded that simple heuristics usually outperform more sophisticated methods such as Bayesian learning [2]. We believe this is due to the fact that usually, the simple methods assume less about the opponent, which makes them more robust in everyday negotiation settings.

For the acceptance mechanisms, we studied and classified current approaches in [5]. We found that designing a good acceptance mechanism amounts to solving the acceptance dilemma: accepting bad to mediocre offers yields more agreements of relatively low utility, while accepting only the best offers produces less agreements, but of higher utility. Most of the current acceptance mechanisms are on either side of the extreme, and this insight led us to believe that more consideration has to be given to exactly what is the right time to accept. Consequently, we have adopted a more principled approach in [3], where we calculate the optimal time to accept when estimates of the opponent's future behavior is available.

3. FUTURE WORK

The next steps in our research revolve around a number of themes. First, we believe there is still room for improvement in both opponent preference estimation and strategy prediction methods, and we think online genetic algorithms can be promising in this regard. Also, learning across multiple negotiation sessions (against different opponents, and across different scenarios) has not yet received much attention, but ANAC 2013 will incorporate such learning mechanisms, which we hope will stimulate research in this area.

Second, we believe there is still a lot of research to be conducted in terms of the bidding strategy. Research in automated negotiation has mostly focused on learning techniques, but it is still rather unclear what to do with the information once it is learned.

Lastly, after analyzing all components separately, we have begun working on putting the pieces back together again. We are basically interested in three questions:

- Does combining the best of each component in the end create the best strategy?
- Which component is the most important with respect to the agent's performance?
- What aspects of the opponent should we learn, and what measures can we use to predict a good outcome?

We hope in the end, this will shed some light on where researchers should direct their effort and time when designing a negotiating agent.

4. REFERENCES

- [1] Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence Journal*, 2012.
- [2] Tim Baarslag, Mark Hendriks, Koen Hindriks, and Catholijn Jonker. Measuring the performance of online opponent models in automated bilateral negotiation. In Michael Thielscher and Dongmo Zhang, editors, *AI 2012: Advances in Artificial Intelligence*, volume 7691 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2012.
- [3] Tim Baarslag and Koen Hindriks. Accepting optimally in automated negotiation with incomplete information. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '13, 2013.
- [4] Tim Baarslag, Koen Hindriks, Mark Hendriks, Alex Dirkzwager, and Catholijn Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In *Proceedings of The Fifth International Workshop on Agent-based Complex Automated Negotiations (ACAN 2012)*, 2012.
- [5] Tim Baarslag, Koen Hindriks, and Catholijn Jonker. Acceptance conditions in automated negotiation. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 95–111. Springer Berlin / Heidelberg, 2013.
- [6] Tim Baarslag, Koen Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The first automated negotiating agents competition (ANAC 2010). In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, pages 113–135. Berlin, Heidelberg, 2012. Springer-Verlag.
- [7] C. Beam and A. Segev. Automated negotiations: A survey of the state of the art. *Wirtschaftsinformatik*, 39(3):263–268, 1997.
- [8] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [9] R.H. Guttman, A.G. Moukas, and P. Maes. Agent-mediated electronic commerce: a survey. *The Knowledge Engineering Review*, 13(02):147–159, 1998.
- [10] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, M.J. Wooldridge, and C. Sierra. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215, 2001.
- [11] Gregory E. Kersten and Grant Zhang. Mining inspire data for the determinants of successful internet negotiations. *InterNeg Research Papers INR 04/01 Central European Journal of Operational Research*, 2003.
- [12] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94(1-2):79–97, 1997.
- [13] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 2012.
- [14] Alessio R. Lomuscio, Michael Wooldridge, and Nicholas R. Jennings. A classification scheme for negotiation in electronic commerce. *Group Decision and Negotiation*, 12:31–56, 2003.
- [15] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, 50:97–109, 1982.