Automated Negotiating Agents Competition (ANAC)

Catholijn M. Jonker, Reyhan Aydoğan, Tim Baarslag,
Katsuhide Fujita, Takayuki Ito, Koen Hindriks

Interactive Intelligence Group, Delft University of Technology, The Netherlands
Computer Science, Özyeğin University, Istanbul,Turkey
Centrum Wiskunde & Informatica (CWI), The Netherlands
Department of Computer Science and Engineering, Nagoya Institute of Technology, Japan
Institute of Engineering, Tokyo University of Agriculture and Technology, Tokyo, Japan

Abstract

The annual International Automated Negotiating Agents Competition (ANAC) is used by the automated negotiation research community to benchmark and evaluate its work and to challenge itself. The benchmark problems and evaluation results and the protocols and strategies developed are available to the wider research community.

Motivation and Aims

The negotiations studied are classified into bilateral and multilateral. Automated bilateral negotiations were already studied extensively before the start of the ANAC competition. However, work before ANAC presented individual solutions, claiming improvements over other approaches on the basis of performing well in rather specific example domains. The construction of GENIUS[1] (Lin et al. 2014) was done with aim of addressing this problem and immediately opened the possibility of organizing ANAC. This brought the research community together and led to significant improvements on the agents for automated bilateral negotiations on linear additive domains. In a few years the improvements were getting smaller and the community realized that it was time to tackle new challenges.The automated negotiating agents competition has the following aims:

1. to provide an incentive for the development of effective and efficient negotiation protocols and strategies for bidding, accepting and opponent modeling for different negotiation scenarios

2. to collect and develop a benchmark of negotiation scenarios, protocols and strategies

3. to develop a common set of tools and criteria for the evaluation and exploration of new protocols and new strategies against benchmark scenarios, protocols and strategies

4. to set the research agenda for automated negotiation.

Originally, the competition focused on the area of bilateral multi-issue closed negotiation. Over the years the competition has addressed various topics: varying the number of negotiators, the complexity of the negotiation domains (additive linear versus non-linear), repeated negotiations with the

[1]http://ii.tudelft.nl/genius

same set of opponents, negotiations in special domains, and negotiating against humans. The topics of upcoming challenges is determined by the research community. The meeting during which the results are presented is used to gather opinions, decisions on the topics are finalized through emails and polls.

The next sections discuss the benchmark characteristics of negotiation scenarios, an overview of which challenges were posed in what year (Table 1), a short description of GENIUS, and some future challenges. The references section contains publications of the results of competitions, descriptions of finalist agents, and descriptions of the platforms used.

Benchmark Characteristics

Negotiation scenarios describe the protocol, the domain of negotiation, and the preference profiles. Protocols describe what information can be exchanged and the timing aspects. For the sake of research the scenarios also define what information the agents can maintain over negotiation sessions.

Domain Complexity In all competitions (except the Diplomacy track) we vary the size of the domain: small domains of some dozens of outcomes, thousands, and hundreds of thousands (or more) possible outcomes. In some competitions we use relatively simple structured domains for which the preference profiles can be modelled using additive linear utility functions, but in others we use rectilinear hypercubes to model domains having non-linear interdependencies between issues.

Information Sharing Closed negotiation, when opponents do not reveal their preferences to each other, is an important class of real-life negotiations. As the game-theoretic approaches cannot be directly applied to design efficient negotiating agents due to the lack of information about opponent, instead, heuristic approaches are used to design negotiating agents. However, when humans are at the negotiating table, they typically prefer to share more than just bids. Negotiators can share some information of what issues are important to them, can indicate that they like one offer better than another, and so on. Per competition the type of information that is shared is indicated. So far, we never chose for completely open negotiations, i.e., where

Table 1: ANAC 2010-2017 Overview

Criterion \ Competition	2017-A	2017-B	2017-C	2016	2015	2014	2013	2012	2011	2010
Nr of Players	3	2	7	3	3	2	2	2	2	2
Against Human Player		✓								
Domain(s)	linear	linear	Diplomacy	linear	linear	non-linear	linear	linear	linear	linear
Discount Factor	✓			✓	✓	✓	✓	✓	✓	
Shared deadline	✓	✓	✓	✓	✓	✓	✓	✓	✓	per agent
Reservation Value	✓			✓	✓	✓	✓	✓		
Learning	✓						✓			
Partial Offers		✓								
Emotions		✓								
Framework	GENIUS	IAGO	BANDANA	GENIUS	GENIUS	GENIUS	GENIUS	GENIUS	GENIUS	GENIUS

the negotiators share their full preference profiles. In 2013, we increased the complexity by adding to each preference profile a private reservation value. That means that even if a lot of the preference profile would correctly estimated by the opponent, the opponent would still be uncertain about when the agent might walk away from the negotiation, because of outside options represented by the private reservation value.

Number of Negotiating Parties The negotiations studied are classified into bilateral and multilateral negotiations. The early years of ANAC led to significant improvements in the strategies for automated bilateral negotiations on linear additive domains. In 2015 the community saw that the improvements were getting smaller and realized that it was time to tackle new challenges. Increasing the number of negotiating parties inspired new innovations for the protocols and required changes in the strategies for bidding, accepting and opponent modeling.

Time Frames and Discount Factors In all competitions we use a deadline. The reasons for doing so are both pragmatic and to make the competition more interesting from a theoretical perspective. Without a deadline, the negotiation might go on forever, especially without any discount factors. Also, with unlimited time an agent may simply try a large number of proposals to learn the opponents preferences. In addition, as opposed to having a fixed number of rounds, the competition runs in real time. As it is unknown how long it takes an opponent to compute a counter offer, it introduces uncertainty about the number of negotiation rounds,

In ANAC 2010 each agent had three minutes to deliberate. To be effective, agents need to keep track of their own time and the time the opponent has left. From ANAC2011 onward, the agents share a time window of three minutes.

As of 2011 discount factors are frequently part of the scenarios. Discount factors reduce the utility of deals with the progression of time. Adding discount factors provides an incentive to the agents to reach deals faster.

Learning from past negotiations Human negotiators learn during a negotiation session, but also from past negotiations. Negotiating often in the same domain, leads to better estimations of preference profiles of different opponents, but also to insights in the negotiation strategies of opponents. In 2013

agents were allowed to save information from ongoing negotiations, and to load information from past negotiations. By analyzing past negotiation sessions, agents can estimate the opponents utility function based on exchanged bids. They can also analyse under which conditions the opponent concedes (e.g., in response to the bidding behaviour of the other, or in response to the progression of time). The agent can adapt its strategies to best negotiate in this domain, against this opponent.

Tournaments in GENIUS

Negotiating agents designed using heuristic approaches need extensive evaluation, typically through simulations and empirical analysis, since it is usually impossible to predict precisely how the system and the constituent agents will behave in a wide variety of circumstances. To facilitate this research the GENIUS system was introduced and is continuously further developed. Use it to run tournaments, access our repository of protocols, domains, preferences, agents and the BOA-framework, or to let humans play against your agents.

Overall Impact and Lessons Learned

The state of the art in generic automated bilateral negotiating agents is hard to beat. We found that 1) Tough agents perform better, 2) Opponent models are less important than thought, 3) Simple opponent modeling techniques perform best. More in (Baarslag et al. 2015).

Future Competitions and Challenges

We challenge interested readers to join our community and motivate us to address challenges that you might pose to us. To entice you, here are some of our future plans.

Note that humans typically don't fully know their preference profiles when they start negotiating. This can be modeled by changing the issues, value ranges per issue, and preference profiles during the negotiations and requiring that the agents adequately adapt their behavior. Another challenge is to develop agents that pick the best negotiation strategies for a given type of domain and opponent.

Acknowledgments

To all contributors and co-organizers over the years.

References

ANAC. 2010. Anac 2010. http://mmi.tudelft.nl/negotiation/index.php/ANAC_2010.

ANAC. 2011. Anac 2011. http://www.itolab.nitech.ac.jp/ANAC2011/.

ANAC. 2012. Anac 2012. http://anac2012.ecs.soton.ac.uk.

ANAC. 2013. Anac 2013. http://www.itolab.nitech.ac.jp/ANAC2013/.

ANAC. 2014. Anac 2014. http://www.itolab.nitech.ac.jp/ANAC2014/.

ANAC. 2015. Anac 2015. http://web.tuat.ac.jp/~katfuji/ANAC2015/.

ANAC. 2016. Anac 2016. http://web.tuat.ac.jp/~katfuji/ANAC2016/.

ANAC. 2017. Anac 2017. http://web.tuat.ac.jp/~katfuji/ANAC2017/.

Aydoğan, R.; Baarslag, T.; Jonker, C. M.; Fujita, K.; Ito, T.; Hadfi, R.; and Hayakawa, K. 2016. A baseline for nonlinear bilateral negotiations: The full results of the agents competing in ANAC 2014.

Baarslag, T.; Hindriks, K.; Jonker, C. M.; Kraus, S.; and Lin, R. 2012. The first automated negotiating agents competition (ANAC 2010). In Ito, T.; Zhang, M.; Robu, V.; Fatima, S.; and Matsuo, T., eds., *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, 113–135. Berlin, Heidelberg: Springer-Verlag.

Baarslag, T.; Fujita, K.; Gerding, E. H.; Hindriks, K.; Ito, T.; Jennings, N. R.; Jonker, C.; Kraus, S.; Lin, R.; Robu, V.; and Williams, C. R. 2013. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence* 198(0):73 – 103.

Baarslag, T.; Aydoğan, R.; Hindriks, K. V.; Fuijita, K.; Ito, T.; and Jonker, C. M. 2015. The automated negotiating agents competition, 2010-2015. *AI Magazine* 36(4):115–118.

Bandana. 2017. Bandana framework. http://www.iiia.csic.es/~davedejonge/bandana/.

Calhamer, A. B. 2017. Diplomacy rules. https://www.wizards.com/avalonhill/rules/diplomacy.pdf.

Gal, K., and Ilany, L. 2015. *The Fourth Automated Negotiation Competition*. Tokyo: Springer Japan. 129–136.

IAGO. 2017. Iago framework. http://people.ict.usc.edu/~mell/IAGO/.

Lin, R.; Kraus, S.; Baarslag, T.; Tykhonov, D.; Hindriks, K.; and Jonker, C. M. 2014. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence* 30(1):48–70.

Mell, J., and J. Gratch, J. 2016. Iago: Interactive arbitration guide online. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems International Foundation for Autonomous Agents and Multiagent Systems*.

Sierra, C. 2017. Introduction to diplomacy. https://www.youtube.com/watch?v=z40JP-PJ1vI&feature=youtu.be.

Williams, C.; Robu, V.; Gerding, E.; and Jennings, N. 2014. An overview of the results and insights from the third automated negotiating agents competition (ANAC 2012). In Marsa-Maestre, I.; Lopez-Carmona, M.; Ito, T.; Zhang, M.; Bai, Q.; and Fujita, K., eds., *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*. Springer, Japan. 151–162.