The Sixth Automated Negotiating Agents Competition (ANAC 2015)

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

Abstract In May 2015, we organized the Sixth International Automated Negotiating Agents Competition (ANAC 2015) in conjunction with AAMAS 2015. ANAC is an international competition that challenges researchers to develop a successful automated negotiator for scenarios where there is incomplete information about the opponent. One of the goals of this competition is to help steer the research in the area of multi-issue negotiations, and to encourage the design of generic negotiating agents that are able to operate in a variety of scenarios. 24 teams from 9 different institutes competed in ANAC 2015. This chapter describes the participating agents and the setup of the tournament, including the different negotiation scenarios that were used in the competition. We report on the results of the qualifying and final round of the tournament.

K. Fujita (✉)
Faculty of Engineering, Tokyo University of Agriculture and Technology, Tokyo, Japan
e-mail: katfuji@cc.tuat.ac.jp

R. Aydoğan
Computer Science Department, Özyeğin University, Istanbul, Turkey
e-mail: reyhan.aydogan@gmail.com

T. Baarslag
Agents, Interaction and Complexity group, University of Southampton, Southampton, UK
e-mail: tb1m13@ecs.soton.ac.uk

T. Ito
Techno-Business Administration (MTBA), Nagoya Institute of Technology, Aichi, Japan
e-mail: ito.takayuki@nitech.ac.jp

K. Hindriks · C. Jonker
Man Machine Interaction Group, Delft University of Technology, Delft, The Netherlands
e-mail: K.V.Hindriks@tudelft.nl

C. Jonker
e-mail: C.M.Jonker@tudelft.nl

1   Introduction

Success in developing an automated agent with negotiation capabilities has great advantages and implications. In order to help focus research on proficiently negotiating automated agents, we have organized the automated negotiating agents competition (ANAC). The results of the different implementations are difficult to compare, as various setups are used for experiments in ad hoc negotiation environments [6]. An additional goal of ANAC is to build a community in which work on negotiating agents can be compared by standardized negotiation benchmarks to evaluate the performance of both new and existing agents. Recently, the analysis of ANAC becomes important fields of automated negotiations in multi-agent systems [1].

In designing proficient negotiating agents, standard game-theoretic approaches cannot be directly applied. Game theory models assume complete information settings and perfect rationality [8, 9]. However, human behavior is diverse and cannot be captured by a monolithic model. Humans tend to make mistakes, and they are affected by cognitive, social and cultural factors [7]. A means of overcoming these limitations is to use heuristic approaches to design negotiating agents. When negotiating agents are designed using a heuristic method, we need an extensive evaluation, typically through simulations and empirical analysis.

We employ an environment that allows us to evaluate agents in a negotiation competition: GENIUS [6], a **G**eneral **E**nvironment for **N**egotiation with **I**ntelligent multi-purpose **U**sage **S**imulation. GENIUS helps facilitating the *design* and *evaluation* of automated negotiators' strategies. It allows easy development and integration of existing negotiating agents, and can be used to simulate individual negotiation sessions, as well as tournaments between negotiating agents in various negotiation scenarios. The design of general automated agents that can negotiate proficiently is a challenging task, as the designer must consider different possible environments and constraints. GENIUS can assist in this task, by allowing the specification of different negotiation domains and preference profiles by means of a graphical user interface. It can be used to train human negotiators by means of negotiations against automated agents or other people. Furthermore, it can be used to teach the design of generic automated negotiating agents.

The First Automated Negotiating Agents Competition (ANAC 2010) was held in May 2010, with the finals being run during the AAMAS 2010 conference. Seven teams had participated and three domains were used. *AgentK* generated by the Nagoya Institute of Technology team won the ANAC 2010 [2]. The Second Automated Negotiating Agents Competition (ANAC 2011) was held in May 2011, with the AAMAS 2011 conference. 18 teams had participated and eight domains were used. The new feature of ANAC 2011 was the discount factor. *HardHeaded* generated by the Delft University of Technology won the ANAC 2011 [3]. The Third Automated Negotiating Agents Competition (ANAC 2012) was held in May 2012, with the AAMAS 2012 conference. 17 teams had participated and 24 domains were used. The new feature of ANAC 2012 was the reservation value. *CUHKAgent* generated by the Chinese University of Hong Kong won the ANAC 2012 [10]. The Forth

Automated Negotiating Agents Competition (ANAC 2013) was held in May 2013, with the AAMAS 2013 conference. 19 teams had participated and 24 domains were used. The new feature of ANAC 2013 was that agents can use the bidding history. *The Fawkes* generated by the Delft University of Technology won the ANAC 2013 [5]. The Fifth Automated Negotiating Agents Competition (ANAC 2014) was held in May 2014, with the AAMAS 2014 conference. 21 teams had participated and 12 domains were used. The new feature of ANAC 2014 was nonlinear utility functions. *AgentM* generated by Nagoya Institute of Technology won the ANAC 2014 [4].

ANAC organizers have been employing some of the new feature every year to develop the ANAC competition and the automated negotiations communities. The challenge of ANAC 2015 is to reach an agreement while negotiating with two opponents at the same time. In addition,the utility functions are linear again, as they were in ANAC 2010–2013. The multi-player protocol is a simple extension of the bilateral alternating offers protocol, called the Stacked Alternating Offers Protocol for Multi-Lateral Negotiation (SAOP).

The timeline of ANAC 2015 is mainly consisted by two parts: Qualifying Round and Final Round. First, the qualifying round was played in order to select the finalists from 24 agents by considering the individual utility and the nash product. In the qualifying round, 24 agents was divided into four groups (pools) randomly, and the best two agents of those pools proceed to the final in each category. After that, the final round was played among 8 agents in two categories, which won the qualifying round. The domains and preference profiles in the qualifying and final rounds were 10 domains generated by the organizers. The entire matches played among 8 agents in each category, and the ranking of ANAC 2015 is decided.

The remainder of this chapter is organized as follows. Section 2 provides an overview over the design choices for ANAC, including the model of negotiation, tournament platform and evaluation criteria. In Sect. 3, we present the setup of ANAC 2015 followed by Sect. 4 that layouts the results of competition. Finally, Sect. 5 outlines our conclusions and our plans for future competitions.

2 Setup of ANAC 2015

2.1 Negotiation Model

Given the goals outlined in the introduction, in this section we introduce the setup and negotiation protocol used in ANAC. The interaction between negotiating parties is regulated by a *negotiation protocol* that defines the rules of how and when proposals can be exchanged. The parties negotiate over a set of *issues*, and every issue has an associated range of alternatives or *values*. A negotiation *outcome* consists of a mapping of every issue to a value, and the set, Ω of all possible outcomes is called the negotiation *domain*. The domain is common knowledge to the negotiating parties and stays fixed during a single negotiation session. In addition to the domain, both

parties also have privately-known preferences described by their *preference profiles* over Ω. These preferences are modeled using a utility function U that maps a possible outcomes $\omega \in \Omega$ to a real-valued number in the range [0, 1]. In ANAC 2015, the utilities are *linearly additive*. That is, the overall utility consists of a weighted sum of the utility for each individual issue. While the domain (i.e. the set of outcomes) is common knowledge, the preference profile of each player is private information. This means that each player has only access to its own utility function, and does not know the preferences of its opponent.[1] Moreover, we use the term *scenario* to refer to the domain and the pair of preference profiles (for each agent) combined.

Finally, we supplement it with a deadline, reservation value and discount factors. The reasons for doing so are both pragmatic and to make the competition more interesting from a theoretical perspective. In addition, as opposed to having a fixed number of rounds, both the discount factor are measured in *real time*. In particular, it introduces yet another factor of uncertainty since it is now unclear how many negotiation rounds there will be, and how much time an opponent requires to compute a counter offer. In ANAC 2015, the discount factors and reservation value depend on the scenario, but the deadline is set to three minutes. The implementation of discount factors in ANAC 2015 is as follows:

A negotiation lasts a predefined time in seconds*(deadline)*. The timeline is normalized, i.e.: time $t \in [0, 1]$, where $t = 0$ represents the start of the negotiation and $t = 1$ represents the deadline. When agents can make agreements in the deadline, the individual utilities of each agent are the *reservation value*. Apart from a deadline, a scenario may also feature *discount factors*. Discount factors decrease the utility of the bids under negotiation as time passes. Let d in [0, 1] be the discount factor. Let t in [0, 1] be the current normalized time, as defined by the timeline. We compute the discounted utility U_D^t of an outcome ω from the undiscounted utility function U as follows:

$$U_D^t(\omega) = U(\omega) \cdot d^t \tag{1}$$

At $t = 1$, the original utility is multiplied by the discount factor. Furthermore, if $d = 1$, the utility is not affected by time, and such a scenario is considered to be undiscounted.

In the competition, we use the Stacked Alternating Offers Protocol for Multi-Lateral Negotiation (SAOP) as the new feature, in which the negotiating parties exchange offers in turns. All of the participants around the table get a turn per round; turns are taken clock-wise around the table. The first party starts the negotiation with an offer that is observed by all others immediately. Whenever an offer is made the next party in line can take the following actions:

---

[1]We note that, in the competition each agent plays *all* preference profiles, and therefore it would be possible in theory to learn the opponent's preferences. However, the rules explicitly disallow learning *between* negotiation sessions, and only *within* a negotiation session. This is done so that agents need to be designed to deal with unknown opponents.

- Make a counter offer (thus rejecting and overriding the previous offer)
- Accept the offer
- Walk away (e.g. ending the negotiation without any agreement)

This process is repeated in a turn taking clock-wise fashion until reaching an agreement or reaching the deadline. To reach an agreement, all parties should accept the offer. If at the deadline no agreement has been reached, the negotiation fails. The details of SAOP is written in the next chapter.

2.2   Running the Tournament

As a tournament platform to run and analyze the negotiations, we use the GENIUS environment (General Environment for Negotiation with Intelligent multi-purpose Usage Simulation) [6]. GENIUS is a research tool for automated multi-issue negotiation, that facilitates the design and evaluation of automated negotiators' strategies. It also provides an easily accessible framework to develop negotiating agents via a public API. This setup makes it straightforward to implement an agent and to focus on the development of strategies that work in a general environment.

GENIUS incorporates several mechanisms that aim to support the design of a general automated negotiator. The first mechanism is an analytical toolbox, which provides a variety of tools to analyse the performance of agents, the outcome of the negotiation and its dynamics. The second mechanism is a repository of domains and utility functions. Lastly, it also comprises repositories of automated negotiators. In addition, GENIUS enables the evaluation of different strategies used by automated agents that were designed using the tool. This is an important contribution as it allows researchers to empirically and *objectively* compare their agents with others in different domains and settings.

The timeline of ANAC 2015 consists of two phases: the qualifying round and the final round. The domains and preference profiles used during the competition are not known in advance and were designed by the organizers. An agent's success is measured using the evaluation metric in all negotiations of the tournament for which it is scheduled.

First, a qualifying round was played in order to select the finalists from the 24 agents that were submitted by the participating teams. Since there were too many agents, in the different domains, a whole tournament in the qualifying round is impossible. Therefore, 24 agents was divided to four groups (pools) randomly, and the best two agents in nash product and individual utility in each pool proceed to the final round. It took two weeks to finish the all pools of the qualifying round. In ANAC-2015, we didn't allow the updating agents between the qualifying round and the final round.

The final round was played among the the agents that achieved the best scores (individual utility and nash product) in each pool during qualifying. We prepared two categories in the final round of ANAC 2015: Individual utility category and nash

product categories. The domains and preference profiles are same as the qualifying round. The entire matches played among agents, and the final ranking of ANAC 2015 was decided. To reduce the effect of variation in the results, the final score calculates the average of the five trials.

3 Competition Domains and Agents

3.1 Scenario Descriptions

The ANAC is aimed towards modeling multi-issue negotiations in uncertain, open environments, in which agents do not know what the preference profile of the opponent is. The various characteristics of a negotiation scenario such as size, number of issues, opposition, discount factor and reservation value can have a great influence on the negotiation outcome. Therefore, we generated ten types of domains and profiles in the competition. Especially, in the qualifying round and final round, we used all 10 scenarios with different discount factors and reservation values and profiles. In other words, they have vary in terms of the number of issues, the number of possible proposals, the opposition of the preference profiles (see Table 1). The 3d negotiation space plotting in each domain are represented graphically in Fig. 1.

Table 1 The domains used in ANAC 2015

| ID	Number of issues	Size	Discount factor	Reservation value	Cooperativeness
1	1	5	None	0.5	Very competitive
2	1	5	None	0.5	A bit collaborative
3	2	25	0.2	None	Very competitive
4	2	25	None	0.5	Quite collaborative
5	4	320	0.5	None	Competitive
6	4	320	0.5	None	Collaborative
7	8	3^8	None	None	Competitive
8	8	3^8	None	None	Collaborative
9	16	2^{16}	0.4	0.7	Very collaborative
10	16	2^{16}	0.4	0.7	Very competitive
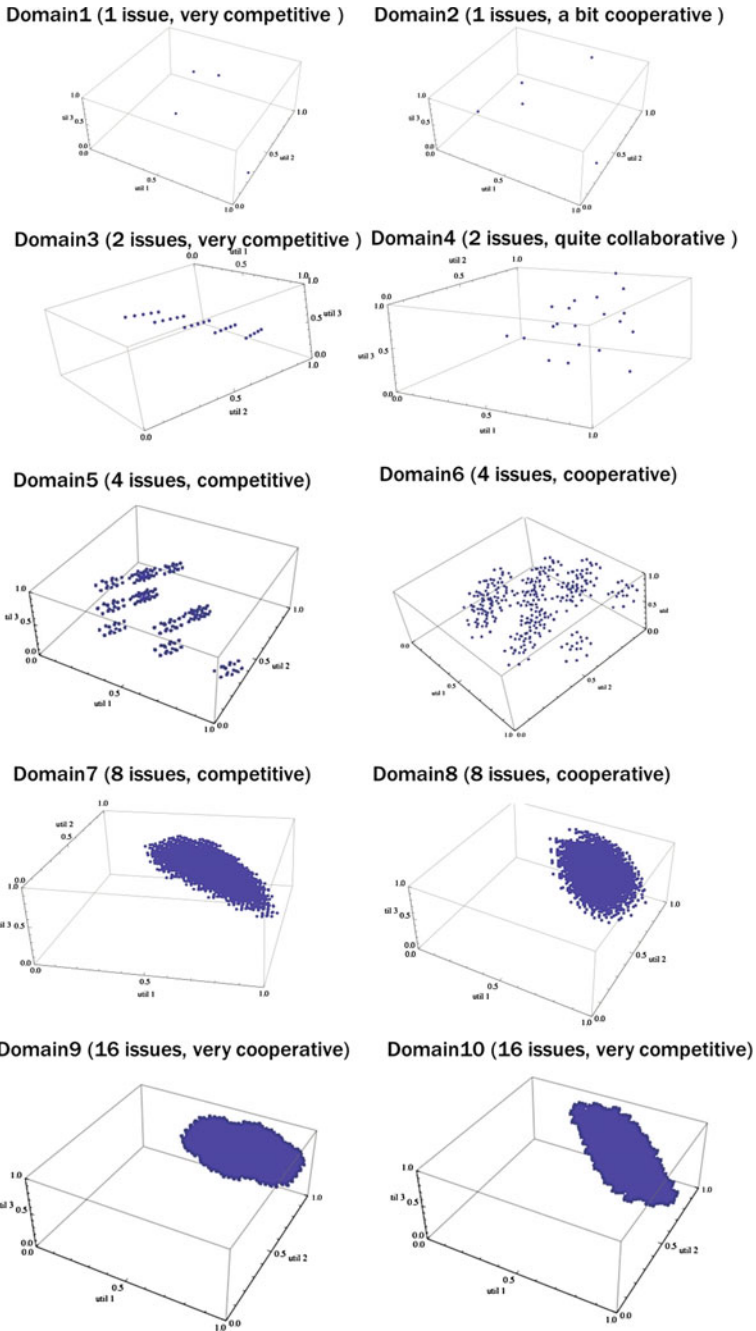
Fig. 1 3D negotiation space plotting in each domain

3.2   Agent Descriptions

ANAC 2015 had 24 agents, registered from 9 institutes from 7 countries: The Chinese University of Hong Kong, Hong Kong; Nanyang Technological University, Singapore; University of Isfahan, Iran; Nagoya Institute of Technology, Japan; Tokyo University of Agriculture and Technology, Japan; Delft University of Technology, Netherlands; Maastricht University, Netherlands; Norwegian University of Science and Technology, Norway; University of Tulsa, US. Table 2 shows the all participants in ANAC 2015.

Finalists are the winners of the qualifying round. In the rest of this book, we provide chapters of the individual strategies of the ANAC2015 finalists.

4   Competition Results

We describe the results of the qualifying and final rounds.

4.1   Qualifying Round

First, a qualifying round was played in order to select the finalists from the 24 agents that were submitted by the participating teams. 24 agents was divided to four groups (pools) randomly, and the best two agents in nash product and individual utility in each pool proceeded to the final round in each category. Each tournament wasn't repeated to prohibit the learning from the previous tournaments.

In order to complete such an extensive set of tournaments within a limited time frame, we used five high-spec computers, made available by Nagoya Institute of Technology and Tokyo University of Agriculture and Technology. Specifically, each of these machines contained an *Intel Core i7* CPU, at least 16GB of DDR3 memory, and a hard drive with at least 2TB of capacity.

Figures 2, 3, 4 and 5 show the results of each agent in the qualifying round (pool1, pool2, pool3 and pool4). The finalists are selected from all pools by considering the individual utilities and nash products. The individual utility means the average of utility of the individual agent in the tournaments. The nash products means the average of the product of utilities of three agents in the tournaments. As figures showing, the best two agents in each pool are selected by considering the individual utility and nash product. As a results, *agentBuyog* and *PokerFace* are selected as finalists from the pool1; *Atlas3* and *XianFaAgent* are selected as finalists from pool2; *ParsAgent* and *kawaii* are selected as finalists from pool3; *RandomDance* and *PhonexParty* are selected as finalists from pool4 in the individual category. Also, *agentBuyog* and *Mercury* are selected as finalists from the pool1; *Atlas3* and *AgentX* are selected as finalists from pool2; *CUHKAgent* and *Jonny Black* are selected as finalists from

Table 2 Team members and agent names in ANAC 2015

No.	Team members	Affliction	Agent name
1	Jeroen Peperkamp Vikko Smit	Delft University of Technology	Pokerface
2	Joao Almeida Hugo Zwaan Xiaoran Liu	Delft University of Technology	TUDMixedStrategyAgent
3	Dirk Schut Nikol Guljelmovic Jelle Munk	Delft University of Technology	Ai Caramba!
4	Shuang Zhou	Maastricht University	Mercury
5	Siqi Chen Jianye Hao Gerhard Weiss HF-Leung	Maastricht University	AresParty
6	Shinji Kakimoto	Tokyo University of Agriculture and Technology	RandomDance
7	Hiroyuki Shinohara	Tokyo University of Agriculture and Technology	AgentHP
8	Bhargav Sosale Swarup Satish Suyog Shivakumar Bo An	Nanyang Technological University	Agent Buyog
9	Neo Jun	Nanyang Technological University	AgentNeo
10	Chen Xian Fa Kelvin	Nanyang Technological University	XianFaAgent
11	Sengoku Akihisa	Nagoya Institute of Technology	SENGOKU
12	Ishida Kenta	Nagoya Institute of Technology	AgentW
13	Masayuki Hayashi	Nagoya Institute of Technology	Agent H - Hayashi
14	Bun Koku	Nagoya Institute of Technology	AgentX
15	Akiyuki Mori	Nagoya Institute of Technology	Atlas3
16	Takuma Inamoto	Nagoya Institute of Technology	Kawaii
17	Kazumasa Takahashi	Nagoya Institute of Technology	DragKnight
18	Zenefa Rahaman Kendall Hyatt Chad Crawford Sandip Sen	University of Tulsa	PNegotiator

Table 2 (continued)

No.	Team members	Affliction	Agent name
19	Nathaniel Beckemeyer Samuel Beckmann Abigail Sislo	University of Tulsa	MeanBot
20	Osman Yucel Jon Hoffman	University of Tulsa	Jonny Black
21	Lam Wing	The Chinese University of Hong Kong	PhoenixParty
22	Leung Hoi Tang Ng Chi Wing Ho-fung Leung	The Chinese University of Hong Kong	CUHKAgent2015
23	Zahra Khosravimehr Faria Nasiri Mofakham	University of Isfahan	ParsAgent
24	Lars Liahagen Haakon H. Rod	Norwegian University of Science and Technology (NTNU)	Forseti

Agent Name	Average (Individual)	Rank (Individual)	Average (Nash Product)	Rank (Nash)
agentBuyogV2	**0.597955067**	**1**	**0.275919234**	**2**
PokerFace	**0.594266467**	**2**	0.268609115	5
PNegotiator	0.591739600	3	0.268958032	4
DrageKnight	0.571351533	4	0.272096395	3
Mercury	0.550937867	5	**0.300345722**	**1**
SENGOKU	0.547276433	6	0.263140705	6

Fig. 2 Average scores of each agent in the qualifying round (pool1)

Agent Name	Average (Individual)	Rank (Individual)	Average (Nash Product)	Rank (Nash)
Atlas3	**0.680664517**	**1**	**0.309788773**	**1**
XianFaAgent	**0.633863800**	**2**	0.288605994	3
MeanBot	0.584072250	3	0.223485414	6
AgentX	0.571492517	4	**0.290723081**	**2**
AgentHP	0.535089883	5	0.254176388	5
TUDMixed	0.504430117	6	0.284713164	4

Fig. 3 Average scores of each agent in the qualifying round (pool2)

pool3; *RandomDance* and *AgentH* are selected as finalists from pool4 in the nash product category.

Agent Name	Average (Individual)	Rank (Individual)	Average (Nash Product)	Rank (Nash)
ParsAgent	**0.582228250**	**1**	0.226744176	6
kawaii	**0.575404450**	**2**	0.224477192	5
Group2	0.567122400	3	0.230310631	4
CUHKAgent2015	0.552638067	4	**0.268464466**	**1**
AgentW	0.518159433	5	0.237260891	3
JonnyBlack	0.491797117	6	**0.261849645**	**2**

Fig. 4 Average scores of each agent in the qualifying round (pool3)

Agent Name	Average (Individual)	Rank (Individual)	Average (Nash Product)	Rank (Nash)
RandomDance	**0.408558450**	**1**	**0.182175298**	**1**
PhoenixParty	**0.380885900**	**2**	0.136591820	6
AresParty	0.378801767	3	0.147792872	4
AgentNeo	0.356815667	4	0.141631385	5
AgentH	0.339627333	5	**0.163050174**	**2**
Forseti	0.258990217	6	0.151909061	3

Fig. 5 Average scores of each agent in the qualifying round (pool4)

4.2 Final Round

It is notable that *Atlas3* was the clear winner of the both categories (see Tables 3 and 4). However, the differences in utilities between many of the ranked strategies are small, so several of the agents were decided the ranking by a small margin. Finally, the first places in the individual utility and nash product categories were awarded to *Atlas3* ($450); The second place in the individual category was awarded to the ParsAgent ($150); The second place in the nash product was awarded awarded to *Mercury* ($175); The third place in the individual category was awarded to *RandomDance*

Table 3 Tournament results in the final round (Individual utility)

Rank	Agent	Score	Standard deviation
1	Atlas3	0.481042722	0.00156024
2	ParsAgent	0.470693979	0.003128712
3	RandomDance	0.46062548	0.003038258
4	kawaii	0.460129481	0.002715924
5	agentBuyog	0.458823101	0.003842303
6	PhoenixParty	0.442975836	0.005032221
7	XianFaAgent	0.353133027	0.001918821
8	PokerFace	0.344003999	0.001433044

Table 4 Tournament results in the final round (Nash product)

Rank	Agent	Score	Standard deviation
1	Atlas3	0.323992201	0.000405256
2	Mercury	0.321600864	0.001620108
3	JonnyBlack	0.313749314	0.001026152
4	AgentX	0.312427823	0.001393852
5	CUHKAgent	0.309464847	0.001726555
6	RandomDance	0.294950885	0.001088483
7	AgentH	0.292136808	0.001547118
8	agentBuyog	0.282378625	0.00236416

($100); The third place in the nash product was awarded awarded to *JonnyBlack* ($125).

5   Conclusion

This chapter describes the Sixth automated negotiating agents competition (ANAC2015). Based on the process, the submissions and the closing session of the competition we believe that our aim has been accomplished. Recall that we set out for this competition in order to steer the research in the area multi-issue closed negotiation. 24 teams have participated in the competition and we hope that many more will participate in the following competitions.

ANAC also has an impact on the development of GENIUS. We have released a new, public build of GENIUS[2] containing all relevant aspects of ANAC. In particular, this includes all domains, preference profiles and agents that were used in the competition. This will make the complete setup of ANAC available to the negotiation research community. Not only have we learnt from the strategy concepts introduced in ANAC, we have also gained understanding in the correct setup of a negotiation competition. The joint discussion with the teams gives great insights into the organizing side of the competition.

To summarize, the agents developed for ANAC will proceed the next step towards creating autonomous bargaining agents for real negotiation problems. We plan to organize the next ANAC in conjunction with the next AAMAS conference.

---

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

References

1. T. Baarslag, K. Fujita, E.H. Gerding, K.V. Hindriks, T. Ito, N.R. Jennings, C.M. Jonker, S. Kraus, R. Lin, V. Robu, C.R. Williams, Evaluating practical negotiating agents: results and analysis of the 2011 international competition. Artif. Intell. **198**, 73–103 (2013)
2. T. Baarslag, K.V. Hindriks, C.M. Jonker, S. Kraus, R. Lin, The first automated negotiating agents competition (ANAC 2010), in *New Trends in Agent-Based Complex Automated Negotiations* (Springer, Heidelberg, 2012), pp. 113–135
3. K. Fujita, T. Ito, T. Baarslag, K.V. Hindriks, C.M. Jonker, S. Kraus, R. Lin, The second automated negotiating agents competition (ANAC2011), in *Complex Automated Negotiations: Theories, Models, and Software Competitions* (Springer, Heidelberg, 2013), pp. 183–197
4. N. Fukuta, T. Ito, M. Zhang, K. Fujita, V. Robu, The fifth automated negotiation competition, in *Recent Advances in Agent-Based Complex Automated Negotiation* (Springer, Switzerland, 2016)
5. Y.K. Gal, L. Ilany, The fourth automated negotiation competition, in *Next Frontier in Agent-Based Complex Automated Negotiation* (Springer, Japan, 2015), pp. 129–136
6. R. Lin, S. Kraus, T. Baarslag, D. Tykhonov, K.V. Hindriks, C.M. Jonker, Genius: an integrated environment for supporting the design of generic automated negotiators. Comput. Intell. **30**(1), 48–70 (2014)
7. R.D. McKelvey, T.R. Palfrey, An experimental study of the centipede game. Econometrica **60**(4), 803–36 (1992)
8. M. Osborne, A. Rubinstein, in *Bargaining and Markets*, Economic Theory, Econometrics, and Mathematical Economics (Academic Press, Cambridge, 1990)
9. M.J. Osborne, A. Rubinstein, in *A Course in Game Theory*, MIT Press Books (MIT Press, Cambridge, 1994)
10. C.R. Williams, V. Robu, E.H. Gerding, N.R. Jennings, An overview of the results and insights from the third automated negotiating agents competition (ANAC2012), in *Novel Insights in Agent-based Complex Automated Negotiation* (Springer, Japan, 2014), pp. 151–162