

The Second Automated Negotiating Agents Competition (ANAC2011)

Katsuhide Fujita, Takayuki Ito, Tim Baarslag, Koen Hindriks, Catholijn Jonker, Sarit Kraus, and Raz Lin

Abstract. In May 2011, we organized the Second International Automated Negotiating Agents Competition (ANAC2011) in conjunction with AAMAS 2011. 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 bilateral multi-issue negotiations, and to encourage the design of generic negotiating agents that are able to operate in a variety of scenarios. Eighteen teams from seven different institutes competed in ANAC2011. 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.

1 Introduction

Negotiation is an important process to form alliances and to reach trade agreements. Research in the field of negotiation originates from various disciplines including

Katsuhide Fujita
School of Engineering, The University of Tokyo
e-mail: fujita@ipr-ctr.t.u-tokyo.ac.jp

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

Tim Baarslag, Koen Hindriks, Catholijn Jonker
Man Machine Interaction Group, Delft University of Technology
e-mail: {T.Baarslag, K.V.Hindriks, C.M.Jonker}@tudelft.nl

Sarit Kraus · Raz Lin
Computer Science Department, Bar-Ilan University
e-mail: {linraz, sarit}@cs.biu.ac.il

Sarit Kraus
Institute for Advanced Computer Studies, University of Maryland

economics, social science, game theory and artificial intelligence (e.g., [2, 10, 14]). Automated agents can be used side by side the human negotiator embarking on an important negotiation task. They can alleviate some of the efforts required of people during negotiations and also assist people that are less qualified in the negotiation process. There may even be situations in which automated negotiators can replace the human negotiators. Another possibility is for people to use these agents as a training tool, prior to actually performing the task. Thus, 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 first automated negotiating agents competition (ANAC). The principal goals of the ANAC competition are as follows:

- Encouraging the design of agents that can proficiently negotiate in a variety of circumstances,
- Objectively evaluating different bargaining strategies,
- Exploring different learning and adaptation strategies and opponent models, and
- Collecting state-of-the-art negotiating agents, negotiation domains, and preference profiles, and making them available and accessible for the negotiation research community.

A number of successful negotiation strategies already exist in literature [5, 6, 8, 9]. However, the results of the different implementations are difficult to compare, as various setups are used for experiments in ad hoc negotiation environments [12]. 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.

In designing proficient negotiating agents, standard game-theoretic approaches cannot be directly applied. Game theory models assume complete information settings and perfect rationality [15]. 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 [3, 4, 13]. 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 have recently introduced an environment that allowed us to evaluate agents in a negotiation competition such as ANAC: **GENIUS** [12], a **General Environment for Negotiation with Intelligent multi-purpose Usage Simulation**. **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 Automated Negotiating Agents Competition (ANAC) 2010 was held on May 12, 2010, with the finals being run during the AAMAS 2010 conference. Seven teams have participated in the first competition and *AgentK* generated by the Nagoya Institute of Technology team won the ANAC2010[1]. The tournament was ran on three different domains in ANAC2010 namely: Zimbabwe-England domain, IteX vs Cypress domain and Travel domain. The main differences between ANAC2010 in last year and ANAC2011 in this year are two points: Shared Timeline and Discount Factor. In ANAC 2010, the agents had three minutes each to deliberate. This means agents have to keep track of both their own time and the time the opponent has left. For ANAC2011, we have chosen a simpler protocol where both agents have a shared time window of three minutes. ANAC 2011 has domains that have discount factors. In ANAC 2010, almost every negotiation between the agents took the entire negotiation time of three minutes each to reach an agreement. Adding discount factors should provide more interesting negotiations with faster deals.

The timeline of ANAC2011 is mainly consisted by three parts: Qualifying Round, Updating Period and Final Round. The domains and preference profiles used during the competition are not known in advance and were designed by all participants. First, the qualifying round was played in order to select the best 8 agents from 18 agents. The entire pairwise matches played among 18 agents, and the best 8 agents of those tournaments proceed to the Finals. We set up the updating period for improving the finalists' agents for the final round. The detail results and all domains for the qualifying round are revealed to all finalists, and they tuned up their agents. Time period of updating period is for two weeks. Finally, the final round was played among 8 agents. The domains and preference profiles in the final were 8 domains submitted by all finalists for the final round. The entire pairwise matches played among 8 agents, and the ranking of ANAC2011 is decided.

The remainder of this paper 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 Section 3, we present the setup of ANAC2011 followed by Section 4 that layouts the results of competition. Finally, Section 5 outlines our conclusions and our plans for future competitions.

2 Set Up of ANAC

2.1 Negotiation Model

Given the goals outlined in the introduction, in this section we introduce the set-up and negotiation protocol used in ANAC. In this competition, we consider *bilateral* negotiations, i.e. negotiation between two parties. The interaction between negotiating parties is regulated by a *negotiation protocol* that defines the rules of how and

when proposals can be exchanged. In the competition, we use the alternating-offers protocol for bilateral negotiation as proposed in [16], in which the negotiating parties exchange offers in turns. The alternating-offers protocol conforms with our criterion to have simplicity of rules. Moreover, it is a protocol which is widely studied and used in literature, both in game-theoretic and heuristic settings of negotiation (a non-exhaustive list includes [7, 10, 11, 14, 15]).

Now, 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 ANAC2011, we focused on settings with a finite set of discrete values per issue.

In addition to the domain, both parties also have privately-known preferences described by their *preference profiles* over Ω . These preferences are modelled using a utility function U that maps a possible outcomes $\omega \in \Omega$ to a real-valued number in the range $[0, 1]$. In ANAC2011, 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 and discount factors. 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 opponent's preferences. 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. Also, this computational time will typically change depending on the size of the outcome space. In ANAC2011, the discount factors depend on the scenario, but the deadline is fixed and is set to three minutes, in which both agents shared this fixed time window². The implementation of discount factors in ANAC2011 is as follows. Let d in $[0, 1]$ be the discount factor. Let t in $[0, 1]$ be the current normalised time, as defined by the timeline. We compute the discounted utility U_D^t as follows:

¹ We note that, in the competition each agent plays *both* 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.

² In contrast, in ANAC 2010, the agents had three minutes *each* to deliberate. This means the agents had to keep track of both their own time and the time the opponent had left, otherwise they run the risk of the opponent walking away unexpectedly.

$$U_D^t(s_1, s_2) = U(s_1, s_2) \cdot d^t \quad (1)$$

If $d = 1$, the utility is not affected by time, and such a scenario is considered to be undiscounted, while if d is very small there is high pressure on the agents to reach an agreement. Note that, in the set-up used in ANAC2011 competition, discount factors are part of the preference profiles and are always *symmetric* (i.e. d always has the same value for both agents).

2.2 Running the Tournament

As a tournament platform to run and analyse the negotiations, we use the GENIUS environment (General Environment for Negotiation with Intelligent multi-purpose Usage Simulation) [12]. 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 ANAC2011 consists of three phases: the qualifying round, the updating period and the final round. The domains and preference profiles used during the competition are not known in advance and were designed by the participants themselves. 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 best 8 agents from the 18 agents that were submitted by the participating teams. Each participant also submitted a domain and pair of preference profiles for that domain. All 18 such domains were used in the qualifying rounds. For each of these domains, negotiations were carried out between all pairings of the 18 agents.

Since there were 18 agents, which each negotiate against 17 other agents, in 18 different domains, a single tournament in the qualifying round consists of $18 \times 17/2 \times 2 \times 18 = 5508$ negotiation sessions³. To reduce the effect of variation in the results, the tournament was repeated 3 times, leading to a total of 16,524 negotiation sessions, each with a time limit of three minutes. In order to complete such an extensive set of tournaments within a limited time frame, we used five high-spec

³ The combinations of 18 agents are $18 \times 17/2$, however, agents play each domain against each other twice by switching the roles.

computers, made available by Nagoya Institute of Technology. Specifically, each of these machines contained an *Intel Core i7* CPU, at least 4GB of DDR3 memory, and a hard drive with at least 500GB of capacity.

The best 8 agents of those tournaments proceed to the finals round. In the qualifying round, considering all possible pairwise matches among the submitted agents is fairer than randomly dividing agents into groups, because in this way, unfair grouping is avoided (e.g. it avoids the situation that some of the groups could be much more competitive than others). The results from the preliminary tournament matching all submitted agents was used for selecting the best 8 agents taking part in the final round.

Between the 3 rounds, we allowed a 2-week updating period, in which the 8 selected finalists were given the chance to improve their agents for the final round. The detailed results and all domains for the qualifying round were revealed to all finalists, and they could use this additional information to tune their agents. The updating period was set at two weeks.

The final round was played among the 8 agents that achieved the best average scores during qualifying. The domains and preference profiles in the final were 8 domains submitted by all finalists for the final round. The entire pairwise matches played among 8 agents, and the final ranking of ANAC2011 was decided. In the final, a single tournament consists of $8 \times 7/2 \times 2 \times 8 = 448$ negotiation sessions⁴. Again, each tournament was repeated three times.

3 Competition Domains and Agents

3.1 Scenario Descriptions

The ANAC is aimed towards modelling multi-issue negotiations in uncertain, open environments, in which agents do not know what the preference profile of the opponent is.

Table 1 The domains used in ANAC2011 Final Round

Domain Name	Number of issues	Size	Opposition
adg	5	25	Weak
NiceOrDie	1	3	Strong
Energy Domain	8	390,625	Strong
IS_BT_Acquisition	5	384	Medium
Grocery	5	1,600	Weak
Amsterdam_party	6	2,268	Weak
laptopdomain	3	27	Weak
CameraDoamin	6	3,600	Medium

⁴ The combinations of 8 agents are $8 \times 7/2$, however, agents play each domain against each other twice by switching the roles.

The various characteristics of a negotiation scenario such as size, number of issues, opposition, discount factor can have a great influence on the negotiation outcome. Therefore, in order to ensure a good spread of negotiation characteristics and fairness, and to reduce any possible bias on the part of the organisers, we gathered the domains and profiles from the participants in the competition. Specifically, in addition to submitting their agents, each participant submitted a scenario, consisting of both a domain and a pair of preference profiles. In the qualifying round, we used all 18 scenarios submitted by the participants. In the final round, eight scenarios submitted by the eight finalists were used. The final scenarios vary in terms of the number of issues, the number of possible proposals, the opposition of the preference profiles and the mean distance of all of the points in the outcome space to the Pareto frontier (see Table 1). The shape of the outcome space of each scenario is presented graphically in Figure 1.

The details of the scenarios are as follows:

Car

The *Car* domain represents a scenario in which a car dealer negotiates with a potential buyer. There are 6 negotiation issues, which represent the features of the car (such as CD player, extra speakers and air conditioning) and each issue takes one of 5 values (good, fairly good, standard, meagre, none), creating 15,625 possible agreements. The domain is almost symmetric and ensures that outcomes with very high utility for both parties can be achieved. Although the best bids of the domain are worth zero for the opponent, this domain is far from a zero sum game. Agents lean to make the agreement which is a 0.85 vs. 0.98 result. An agent can get close to the maximum possible utility (1.00), if it persuades its opponent to accept 0.85. The domain also allows agents to compromise to a fair division point (0.93, 0.93).

Nice Or Die

This domain is very different to the others, due to its very small size and competitiveness. In this domain, agents have to select between 3 possible agreement points: a fair division point, which is less efficient (in the sense that the sum of the agent's utilities is smaller) or one of two selfish points. The domain is symmetric and, naturally, there are only three possible outcomes. The fair division point has utility of (0.29, 0.29), while the other two selfish points have utilities of (1.00, 0.16) and (0.16, 1.00). In the selfish point, one agent can get a high utility despite the utility of the opponent is very low. If agents try to get high utilities, it is hard for them to reach agreements. However, if agents would like to make an agreement in this scenario, the social welfare is small (as, in the ANAC set-up, the agents cannot learn from previous interactions with an opponent).

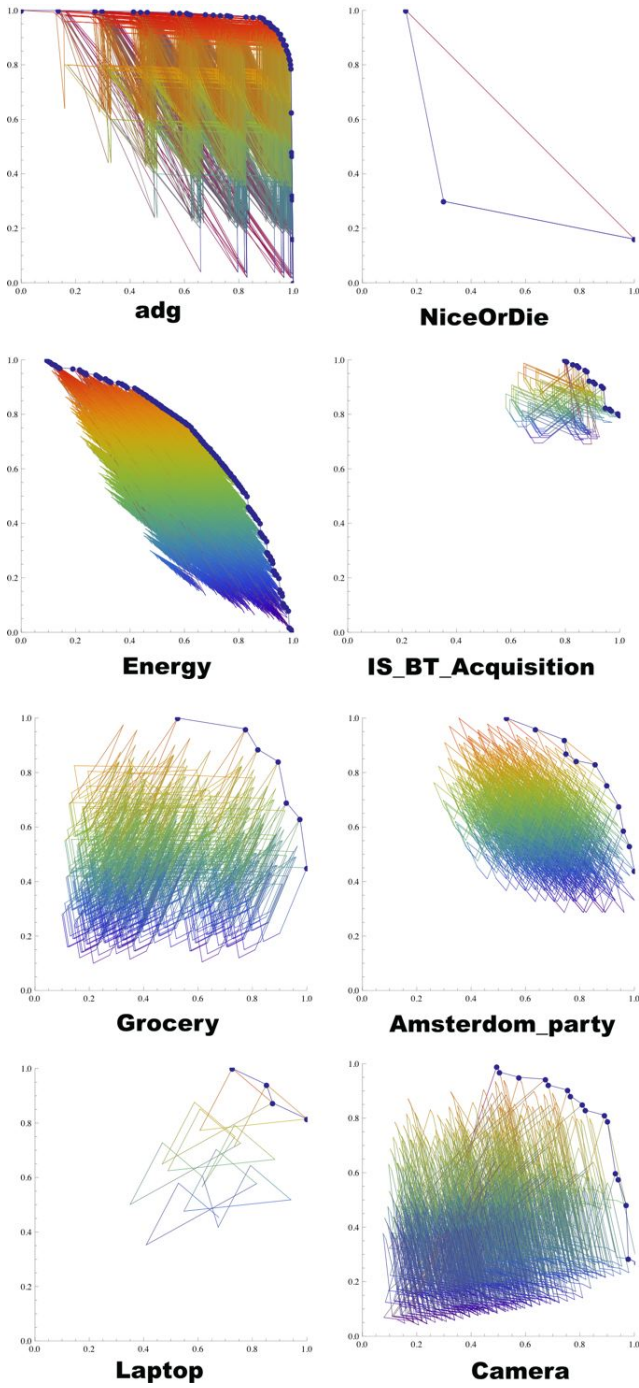


Fig. 1 Acceptance probability space

Energy

This domain considers the problem faced by many electricity companies to reduce electricity consumption during peak times, which requires costly resources to be available and puts a high pressure on local electricity grids. This domain models this application scenario as follows. One agent represents the electricity distribution company whilst the other represents a large consumer. The issues they are negotiating over represent how much the consumer is willing to reduce its consumption over a number of time slots for a day-ahead market (the 24 hours in a day are discretised into 3 hourly time slots). For each issue, there is a demand reduction level possible from zero up to a maximum possible (specifically, 100 kW).

In this domain, the distributor obtains utility by encouraging consumers to reduce their consumptions. Participants set their energy consumption (in kWh) for each of 8 time slots. In each slot, they can reduce their consumption by 0, 25, 50, 75 or 100 kWh. This domain is the largest in the ANAC11 competition (390,625 possible agreements) and has a highly competitive outcome space, therefore, reaching mutually beneficial agreements requires extensive exploration of the outcome space by the negotiating agents.

Company Acquisition

This domain represents a negotiation between two companies, in which the management of Intelligent Solutions Inc. (IS) wants to acquire the BI-Tech company. The negotiation includes five issues: the price that IS pays for BI Tech, the transfer of intellectual property, the stocks given to the BI-Tech founders, the terms of the employees' contracts and the legal liability of Intelligent Solutions Inc.

Each company wants to be the owner of the intellectual property. For IS, this issue is much more important. IS and BI-Tech have common interest that the BI-Tech co-founders would get jobs in IS. IS prefers to give BI-Tech only 2% of the stocks, while the BI-Tech co-founders want 5%. IS prefer private contracts, while firing workers is less desirable by them. BI-Tech prefers a 15% salary raise. For both sides this is not the most important issue in the negotiation. Each side prefers the least legal liability as possible.

The utility range is narrow and has high utility values such that all outcomes give both participants a utility of at least 0.5. The domain is relatively small, with 384 possible outcomes.

Grocery

The *Grocery* domain models a discussion in a local supermarket. The negotiation is between two people living together who have different tastes. The discussion is about five types of product: bread, fruit, snacks, spreads, and vegetables. Each category consists of four or five products, resulting in a medium sized domain with 1,600 possible outcomes. For their daily routine it is essential that a product of each

type is present in their final selection, however only one product can be selected for each type. Besides their difference in taste, they also differ in what category of product they find more important.

The profiles for agents Mary and Sam are modelled in such a way that a good outcome is achievable for both. Sam has a slight advantage, since he is easier to satisfy than Mary, and therefore is likely to have better outcomes. This scenario allows outcomes that are mutually beneficial, but the outcome space is scattered so agents must explore it considerably to find the jointly profitable ones.

Amsterdam Trip

The domain concerns the planning of a touristic trip to Amsterdam and includes issues representing the day and time of travel, the duration of the trip, the type of venues to be visited, the means of transportation and the souvenirs to buy. This domain is moderately large as the utility space has 3,024 possible bid configurations. The preference profiles specify a generous win-win scenario, since it would be unrealistic for two friends to make a trip to Amsterdam and to have it be a zero-sum game.

The size of the domain enables the agent to communicate their preferences (by means of generating bids), without having to concede far. Also the magnitude of the domain puts agents which use a random method of generating bids at a disadvantage, since the odds of randomly selecting a Pareto optimal bid in a large domain are small. So this domain will give an advantage to agents that make some attempt to learn the opponents preference profile, and those capable of rapidly choosing offers.

Laptop

This domain is a variation of the *Laptop* domain from the qualification rounds. Two parties, a seller and a buyer, are negotiating the specifications of a laptop. The domain has three issues: The laptop brand, the size of the hard disk, and the size of the external monitor. Each issue has only three options, making it a very small domain of 27 possible outcomes.

For example, in a negotiation about buying a laptop the buyer may prefer to have a middle-sized screen but the seller may prefer to sell laptops with small screens because s/he has more of those in stock. They could, however, agree on the brand of laptop that they want to buy/sell. An outcome of a negotiation reconciles such differences and results in a purchase.

Camera

This domain is another retail domain, which represents the negotiation between a buyer and a seller of a camera. It has six issues: makers, body, lens, tripods, bags, and accessories. The size of this domain is 3,600 outcomes. The seller gives priority

to the maker, and the buyer gives priority to the lens. The competitiveness of this negotiation domain is medium.

The range of the contract space is wide, which means the agents need to explore it to find the jointly profitable outcomes. While jointly profitable outcomes are possible (since the Pareto frontier is concave), no party has an undue advantage in this (since the Nash point is at an impartial position).

3.2 Agent Descriptions

ANAC2011 had 18 agents, registered from 7 universities: Bar Ilan University, Israel ($\times 5$); University of Southampton, United Kingdom ($\times 2$); Ben-Gurion University, Israel ($\times 4$); Delft University of Technology, The Netherlands ($\times 4$); Politehnica University of Bucharest, Romania; University of Alcalá, Spain and Nagoya Institute of Technology, Japan (one each).

Table 2 Team Members in the Final Round

No.	Team Members	Affiliation	Agent Name	Domain Name
1	Asaf Frieder Dror Sholomon Gal Miller	Bar Ilan University	ValueModelAgent	adg
5	Mai Ben Adar Nadav Sofy Avshalom Elimelech	Bar Ilan University	Gahboninho	NiceOrDie
6	Colin R. Williams Valentin Robu	University of Southampton	IAMhaggler2011	Energy Domain
8	Kadmila Fishel Maya Bercovitch, Ayelet Urieli Betty Sayag	Ben-Gurion University	BRAMAgent	IS_BT_Acquisition
12	Alex Dirkszwaiger Mark Hendrikx Julian de Ruiter	TU Delft	TheNegotiator	Grocery
13	Thijs van Krimpen Daphne Looije Siamak Hajizadeh	TU Delft	HardHeaded	Amsterdam_party
16	Tim Baarslag Koen Hindriks Catholijn Jonker	TU Delft	Nice Tit for Tat agent	laptopdomain
18	Shogo Kawaguchi Katsuhide Fujita Takayuki Ito	Nagoya Institute of Technology	AgentK2	CameraDomain

The final round in ANAC2011 had eight teams from four different universities, as listed in Table 2. They are the winners of the qualifying round. In the rest of the chapter in this book, we provide sections of the individual strategies of the ANAC2011 finalists based on descriptions of the strategies provided by the teams.

4 Competition Results

The ANAC11 competition consisted of two rounds: a qualifying round and a final round. We describe the results of these rounds in turn.

Table 3 Average scores of every strategy in the qualifying round

Rank	Agent Strategy	Mean Utility
1	Gahboninho	0.756
2	HardHeaded	0.708
3	ValueModelAgent	0.706
4	AgentK2	0.702
5	IAMhaggler2011	0.701
6	BRAMAgent	0.690
7	Nice Tit-For-Tat Agent	0.686
8	TheNegotiator	0.685
9	GYRL	0.678
10	WinnerAgent	0.671
11	Chameleon	0.664
12	SimpleAgentNew	0.648
13	LYYAgent	0.640
14	MrFriendly	0.631
15	AgentSmith	0.625
16	IAMcrazyHaggler	0.623
17	DNAgent	0.601
18	ShAgent	0.571

4.1 Qualifying Round

The qualifying round consisted of 18 agents that were submitted to the competition. For each pair of agents, under each preference profile, we ran a total of 3 negotiations. By averaging over all the scores achieved by each agent (against all opponents and using all preference profiles), eight finalists were selected based on their average scores. Formally, the average score $U_{\Omega}(p)$ of agent p in scenario Ω is given by:

$$U_{\Omega}(p) = \frac{\sum_{p' \in P, p' \neq p} U_{\Omega}(p, p') + U_{\Omega}(p', p)}{2 \cdot (|P| - 1)} \quad (2)$$

where P is the set of players and $U_{\Omega}(p, p')$ is the utility achieved by player p against player p' when player p is under the side A of Ω and player p' is under the side B of Ω . It is notable that *Gahboninho* was the clear winner of the qualifying round (see Table 3). However, the differences in utilities between many of the middle ranked strategies are rather small, so several of the agents which qualified for the final only did so by a small margin.

4.2 Final Round

For the final round, we matched each pair of finalist agents, under each preference profile, a total of 3 times. Participants were given the opportunity to submit revised agents for the final based on the results of the qualifying round. Table 4 summarises the means, standard deviations, and 95% confidence interval bounds for the results

Table 4 Tournament results in the final round

Agent	Score	SD	low CI	high CI
HardHeaded	0.748697121	0.00956806	0.74512475	0.752269492
Gahboninho	0.739685706	0.005244734	0.737727511	0.741643902
IAMhaggler2011	0.686419417	0.004663924	0.684678075	0.688160759
AgentK2	0.680806502	0.004701551	0.679051111	0.682561893
TheNegotiator	0.680365635	0.004348087	0.678742215	0.681989055
BRAMAgent	0.679967769	0.005026365	0.678091104	0.681844434
Nice Tit for Tat Agent	0.678179856	0.007599692	0.675342403	0.68101731
ValueModelAgent	0.616818076	0.006890059	0.614245575	0.619390578

Table 5 Detailed scores of every agent in the final round

Agent	adg	NorD	Energy	IS_BT	Grocery	AMS	laptop	Camera
HardHeaded	0.954	0.500	0.549	0.744	0.724	0.867	0.664	0.805
Gahboninho	0.942	0.511	0.676	0.751	0.673	0.914	0.745	0.671
IAMhaggler2011	0.872	0.300	0.558	0.824	0.741	0.787	0.767	0.727
AgentK2	0.922	0.375	0.483	0.797	0.727	0.762	0.673	0.734
TheNegotiator	0.931	0.317	0.533	0.749	0.732	0.797	0.641	0.745
BRAMAgent	0.822	0.500	0.452	0.737	0.724	0.795	0.608	0.741
Nice Tit for Tat Agent	0.784	0.445	0.518	0.754	0.745	0.750	0.630	0.756
ValueModelAgent	0.941	0.193	0.326	0.748	0.758	0.852	0.607	0.777

of each agent. In common with the approach used in the qualifying round, all agents use both of the profiles that are linked to a scenario. The average score achieved by each agent in each scenario is given in Table 5. In the finals, *HardHeaded* proved to be the clear winner, with a score of 0.749.

In more detail, we can consider the performance of the agents in each scenario. Table 5 gives the average score of each agent in each scenario. It shows that *HardHeaded* and *Gahboninho* won by a big margin in most of scenarios. Usually, some agents lost the utility in scenarios with a large size and high opposition, however, *HardHeaded* and *Gahboninho* could get the higher utility in such “tough” scenarios. In addition, *IAMhaggler2011* won the *Company Acquisition* and *Laptop* scenarios with low discount factor, therefore, *IAMhaggler2011* has a high capacity to the discount factors. The differences among *BRAMAgent*, *AgentK2*, *TheNegotiator* are very small.

5 Conclusion

This paper describes the second automated negotiating agents competition. 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 bilateral multi-issue closed negotiation. The competition has achieved just that. Eighteen teams have participated in the competition and we hope that many more will participate in the following competitions.

One of the successes of ANAC lies in the development of state-of-the-art negotiation strategies that co-evolve every year. This incarnation of ANAC already yielded

seven new strategies and we hope that next year will bring even more sophisticated negotiation strategies. ANAC also has an impact on the development of GENIUS. We have released a new, public build of GENIUS⁵ 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 are the first step towards creating autonomous bargaining agents for real negotiation problems. We plan to organize the second ANAC in conjunction with the next AAMAS conference in 2012.

Acknowledgements. The authors would like to thank the team of masters students at Nagoya Institute of Technology, Japan for their valuable help in the organisation of the ANAC2011 competition. Moreover, the authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work. Furthermore, this research is supported by the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs. It is part of the Pocket Negotiator project with grant number VICI-project 08075. In addition, this research is based upon work supported in part under NSF grant 0705587, by the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-08-1-0144 and by ERC grant #267523.

References

1. Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (eds.): *New Trends in Agent-Based Complex Automated Negotiations*. SCI, vol. 383. Springer, Heidelberg (2012)
2. Aumann, R.J., Hart, S. (eds.): *Handbook of Game Theory with Economic Applications*, vol. 1. Elsevier (March 1992)
3. Bazerman, M.H., Neale, M.A.: Negotiator rationality and negotiator cognition: The interactive roles of prescriptive and descriptive research. In: Young, H.P. (ed.) *Negotiation Analysis*, pp. 109–130. The University of Michigan Press (1992)
4. Erev, I., Roth, A.: Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibrium. *American Economic Review* 88(4), 848–881 (1998)
5. Faratin, P., Sierra, C., Jennings, N.R.: Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems* 24(3-4), 159–182 (1998)
6. Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make negotiation trade-offs. *Journal of Artificial Intelligence* 142(2), 205–237 (2003)
7. Fatima, S.S., Wooldridge, M., Jennings, N.R.: Multi-issue negotiation under time constraints. In: *AAMAS 2002: Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 143–150. ACM, New York (2002)
8. Ito, T., Hattori, H., Klein, M.: Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces (2007)

⁵ <http://mmi.tudelft.nl/negotiation/index.php/Genius>

9. Jonker, C.M., Robu, V., Treur, J.: An agent architecture for multi-attribute negotiation using incomplete preference information. *Journal of Autonomous Agents and Multi-Agent Systems* 15(2), 221–252 (2007)
10. Kraus, S.: *Strategic Negotiation in Multiagent Environments*. MIT Press (October 2001)
11. Kraus, S., Wilkenfeld, J., Zlotkin, G.: Multiagent negotiation under time constraints. *Artificial Intelligence* 75(2), 297–345 (1995)
12. Lin, R., Kraus, S., Tykhonov, D., Hindriks, K., Jonker, C.M.: Supporting the Design of General Automated Negotiators. In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T., Yamaki, H. (eds.) *Innovations in Agent-Based Complex Automated Negotiations*. SCI, vol. 319, pp. 69–87. Springer, Heidelberg (2010)
13. McKelvey, R.D., Palfrey, T.R.: An experimental study of the centipede game. *Econometrica* 60(4), 803–836 (1992)
14. Osborne, M.J., Rubinstein, A.: *Bargaining and Markets (Economic Theory, Econometrics, and Mathematical Economics)*. Academic Press (April 1990)
15. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press (1994)
16. Rubinstein, A.: Perfect equilibrium in a bargaining model. *Econometrica* 50(1), 97–109 (1982)