

The First Automated Negotiating Agents Competition (ANAC 2010)

Tim Baarslag, Koen Hindriks, Catholijn Jonker, Sarit Kraus and Raz Lin

Abstract Motivated by the challenges of bilateral negotiations between people and automated agents we organized the first automated negotiating agents competition (ANAC 2010). The purpose of the competition is to facilitate the research in the area bilateral multi-issue closed negotiation. The competition was based on the GENIUS environment, which is a **General Environment for Negotiation with Intelligent multi-purpose Usage Simulation**. The first competition was held in conjunction with the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10) and was comprised of seven teams. This paper presents an overview of the competition, as well as general and contrasting approaches towards negotiation strategies that were adopted by the participants of the competition. Based on analysis in post-tournament experiments, the paper also attempts to provide some insights with regard to effective approaches towards the design of negotiation strategies.

Tim Baarslag¹, Koen Hindriks¹, Catholijn Jonker¹, Sarit Kraus^{2,3} and Raz Lin²

¹ Man Machine Interaction Group

Delft University of Technology

e-mail: {T.Baarslag, K.V.Hindriks, C.M.Jonker}@tudelft.nl

² Computer Science Department

Bar-Ilan University

e-mail: {linraz,sarit}@cs.biu.ac.il

³ Institute for Advanced Computer Studies

University of Maryland

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 economics, social science, game theory and artificial intelligence (e.g., [?, ?, ?]). 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 [?, ?, ?, ?, ?]. However, the results of the different implementations are difficult to compare, as various setups are used for experiments in ad hoc negotiation environments [?, ?]. 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 [?]. 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 [?, ?, ?, ?]. 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** [?], 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.

With GENIUS in place, we organized ANAC with the aim of coordinating the research into automated agent design and proficient negotiation strategies for bilateral multi-issue closed negotiation, similar to what the Trading Agent Competition (TAC) achieved for the trading agent problem [?].

We believe ANAC is an important and useful addition to existing negotiation competitions, which are either aimed at human negotiations or have a different focus, as we explain in Section 2.8.

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 ANAC 2010 followed by Section 4 that layouts the results of competition. In Section 5 we discuss proposed outline for the future ANAC competitions, and finally, Section 6 outlines our conclusions and our plans for future competitions.

2 General Design of ANAC

One of the goals of ANAC is to encourage the design of agents that can negotiate in a variety of circumstances. This means the agents should be able to negotiate against any type of opponent within arbitrary domains. Such an open environment lacks a central mechanism for controlling the agents' behavior, and the agents may encounter human decision-makers who make mistakes and whose behavior is diverse, cannot be captured by a monolithic model, is affected by cognitive, social and cultural factors, etc. [?, ?]. Examples of such environments include online markets, patient care-delivery systems, virtual reality and simulation systems used for training (e.g., the Trading Agent Competition (TAC) [?]). The use of open environments is important as the automated agent needs to be able to interact with different types of opponents, who have different characteristics, e.g. people that origin from different countries and cultures. Automated negotiation agents capable of negotiating proficiently with people thus must deal with the fact that

people are diverse in their behavior and each individual might negotiate in a different manner.

The design of the competition was focused on the development of negotiating strategies, rather than other aspects of the negotiation process (though not less important aspects) such as preference elicitation, argumentation or mediation. The setup of ANAC was designed to make a balance between several concerns, including:

- Strategic challenge: the game should present difficult negotiation domains in a real-world setting with real-time deadlines.
- Multiplicity of issues on different domains, with a priori unknown opponent preferences.
- Realism: realistic domains with varying opponent preferences.
- Clarity of rules, negotiation protocols, and agent implementation details.

We specify the general choices that were made for ANAC with regard to negotiation model and the tournament setup.

2.1 Negotiation Model

In order to define the setup of the negotiation competition, we first introduce the model of negotiation that we use. In this competition, we only consider *bilateral* negotiations, i.e. a negotiation between two parties. The parties negotiate over *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.

We further assume that both parties have certain preferences prescribed by a *preference profile* over Ω . These preferences can be modeled by means of a utility function U that maps a possible outcome $\omega \in \Omega$ to a real-valued number in the range $[0, 1]$. In contrast to the domain, the preference profile of the players is private information.

Finally, the interaction between negotiating parties is regulated by a *negotiation protocol* that defines the rules of how and when proposals can be exchanged. We use the alternating-offers protocol for bilateral negotiation as proposed in [?], in which the negotiating parties exchange offers in turns. The alternating-offers protocol conforms with our criterion to have simplicity of rules. It is widely studied and used in literature, both in game-theoretic and heuristic settings of negotiation (a non-exhaustive list includes [?, ?, ?, ?, ?]).

2.2 *Tournament Platform*

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) [?]. 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 analyze 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.

2.3 *Domains and Preference Profiles*

The specifications of the domain and preferences, such as the constitution and valuation of issues, can be of great influence on the negotiation outcome. We assume that all agents have complete information about the domain, but the preference profile of the players is private information. Thus, if a strategy attempts to tailor its offers to the needs of the opponent, it is required to model the opponent. As the amount of information exchanged during the negotiation is limited in a closed negotiation, the size of the domain has a big impact on the learning capabilities of the agents.

For example, for ANAC 2010, we used a domain named Itex-Cypress [?], in which a buyer and a seller of bicycle components negotiate about issues such as the price of the components and delivery times. There are few possible values per issue, creating a domain of only 180 potential offers. Such a small domain simplifies the task of getting a good picture of the opponent's preferences by studying its proposal behavior.

Due to the sensitivity to the domain specifics, negotiation strategies have to be assessed on negotiation domains of various sizes and of various complexity [?]. Therefore, we selected several domains for ANAC, with different characteristics.

Negotiation strategies can also depend on whether preferences of the negotiating parties are opposed or not. The notion of weak and strong opposition

can be formally defined [?]. Strong opposition is typical of competitive domains, when a gain for one party can be achieved only at a loss for the other party. Conversely, weak opposition means that both parties achieve either losses or gains simultaneously.

Negotiation strategies may depend on the opposition of the preferences. For example, in the case of Itex–Cypress the opposition is strong as it concerns a manufacturer negotiating with a consumer. In such a case the parties have naturally opposing requirements. Hence, the selection of preference profiles should also take into account that the preference profiles have a good variety of opposition.

As stated in the negotiation model, we assume that the negotiating parties have a certain preference profile which can be modeled by a utility function U that maps a possible outcome to a real-valued number in $[0, 1]$.

There are various ways to represent such a utility function (cf. [?]). For ANAC, we have chosen domains without (preferential) dependencies between issues, i.e.: the contribution of every issue to the utility is linear and does not depend on the values of other issues. An advantage of independence between issues is that algorithms that search for a proposal with a particular utility can be implemented in a computationally efficient way. It also makes it easier for negotiation strategies to efficiently model the opponent’s preferences, as it reduces the amount of information that is to be learned by a preference learning technique.

When the tradeoffs between issues are (preferentially) independent, then an additive scoring system is appropriate [?]. Therefore, we assume that utility functions are additive [?, ?, ?]. That is, in a domain Ω with an outcome $\omega = (\omega_1, \dots, \omega_m)$, we assume the utility function has the following form:

$$U(\omega) = \sum_{i=1}^m w_i \cdot u_i(\omega_i),$$

where the w_i are normalized weights (i.e. $\sum w_i = 1$) and $u_i(\omega_i)$ is an evaluation function with range $[0, 1]$ for every individual issue x_i .

2.4 Protocol and Deadline

To add to the realism of the protocol, we can supplement it with a deadline and discount factors. We impose a real-time deadline on the negotiation process for both theoretical and practical reasons. The pragmatic reason is that without a deadline, the negotiation might go on forever, especially without any discount factors. Secondly, with unlimited time an agent may simply try a huge amount of proposals to learn the opponent’s preferences. Another reason for introducing a real-time deadline in the alternating offers protocol is the various challenges it poses to the competitors, as described in Section 2.7.

We believe that a negotiation model with a real-time deadline comes closer to realistic negotiation environment.

2.5 Scoring

We now move on to a formal description of the utility and final scoring functions. Let \mathcal{D} be our set of domains. For every domain $D \in \mathcal{D}$ two preference profiles exist, $\mathcal{P}^D = \{P_1^D, P_2^D\}$. Let \mathcal{A} be the set of competing agents, with $|\mathcal{A}| = n$. Every agent competes against all other agents on all domains (see Section 2.3), alternating between the two preference profiles defined on that domain.

Suppose agent A negotiates with B on domain $D \in \mathcal{D}$, where A has the first preference profile P_1^D and B uses P_2^D . If they reach a certain outcome ω , in which A receives the associated utility $U(\omega)$, then we denote this utility with

$$U_{A \rightarrow B}^D.$$

Our evaluation metric is defined as follows. Every agent A plays against all agents $B \in \mathcal{A}$, with the exception that A will not play itself. The score for A is averaged over all trials, playing with both preference profiles P_1 and P_2 (e.g., on the Itex-Cypress domain, A will play both as Itex and as Cypress against all others). That is, for each profile $P \in \mathcal{P}^D$ an average utility $\bar{u}_D(A, P)$ is calculated for each agent:

$$\bar{u}_D(A, P_1) = \frac{1}{n-1} \sum_{B \in \mathcal{A} \setminus \{A\}} U_{A \rightarrow B}^D,$$

and

$$\bar{u}_D(A, P_2) = \frac{1}{n-1} \sum_{B \in \mathcal{A} \setminus \{A\}} U_{B \rightarrow A}^D.$$

The average utility is then re-normalized using the maximum and minimum utility achieved by all other agents for that profile:

$$\tilde{u}_D(A, P) = \frac{\bar{u}_D(A, P) - \min_{B \in \mathcal{A} \setminus \{A\}} \bar{u}_D(B, P)}{\max_{B \in \mathcal{A} \setminus \{A\}} \bar{u}_D(B, P) - \min_{B \in \mathcal{A} \setminus \{A\}} \bar{u}_D(B, P)}.$$

This gives a score per profile, which is averaged over the two profiles in the domain to give an overall normalized domain-score. The domain-score is then averaged over all trials and the final score $s(A)$ of an agent is determined as an average of the domain-scores:

$$s(A) = \frac{1}{|\mathcal{D}|} \sum_{D \in \mathcal{D}} \left(\frac{1}{2} \tilde{u}_D(A, P_1) + \frac{1}{2} \tilde{u}_D(A, P_2) \right).$$

2.6 ANAC Rules

To enter the competition the teams had to register and upload their agents by a given deadline. The time frame for the submission allowed teams to upload their agents for compliance checks prior to the official deadline.

The tournament itself consists of several rounds. If more than 16 agents are submitted, a qualifying round is held to select the best 16 agents. If an insufficient number of agents is submitted we preserved the right to either run one final round, or to add agents from the repository available at GENIUS at the moment in which we run the qualifying rounds. During the qualifying rounds, tournaments will be played with four agents each, the winners of those tournaments go to the next qualifying round until there are 16 agents in the competition. The top 16 agents proceed to the quarter finals. During the quarter finals rounds, four tournaments are played with four agents each, the top two agents then proceed to the semi-finals. During the semi-finals, two tournaments are played with four agents each, the best two agents proceed to the finals. The finals consist of one tournament with the best four agents of the semi-finals. The best agent wins the competition. In case two agents share the best score, the average results of the two agents when playing against each other will be used to determine the winner.

The domains and preference profiles used during the competition are not known in advance and were designed by the organizing team. The size of the domains can be up to 10 issues. The structure of a negotiation domain is fixed before the tournament and cannot be changed during negotiation.

An agent's success is measured using the evaluation metric (see Section 2.5) in all negotiations of the tournament for which it is scheduled.

Agents can be disqualified for violating the spirit of fair play. In particular, the following behaviors were strictly prohibited:

- Design an agent in such a way that it benefits some specific other agent.
- Communication with the agent during the competition.
- Alter the agent during the competition.
- Denial-of-service attacks.
- Agents employ API operations for the purpose of occupying or loading the game servers.

2.7 Challenges

Our aim in designing ANAC was to provide a strategic challenge on multiple accounts. According to our goals set forward in Section 1, we selected real-life domains containing multiple issues and preference profiles that are unknown to the opponent. Our setup guarantees the participating teams have to deal with the following challenges:

Incomplete information

Suppose an agent wants to model the utility function of the opponent. Because of incompleteness of information about the opponent, it has to learn the opponent's preferences during the negotiation by studying the proposal behavior. Our protocol only allows proposals to be exchanged, so the communication between the agents is very limited. This prevents agents to share information about their preferences, other than by their proposal behavior. Consequently, if agents want to use an opponent model to make effective proposals, they have to make use of a sophisticated learning technique.

Domain complexity

Analyzing the domain for beneficial outcomes is essential when implementing an efficient proposal strategy. Even when an agent has gathered information about its opponent, it still has to be able to find win-win situations, for example by computing the Pareto frontier.

Real-time deadline

Dealing with all of the above while taking into account a real-time deadline. Agents should be more willing to concede near the deadline, as a break-off yields zero utility for both agents. A real-time deadline also makes it necessary to employ a strategy to decide when to accept an offer. Deciding when to accept involves some prediction whether or not a significantly better opportunity might occur in the future.

Some parts of ANAC are less demanding: the utility functions are (linearly) additive functions, so there are no dependencies between issues. This means the utility of an offer can be effectively computed and conversely, a proposal of a certain utility can be easily generated. Moreover, additive utility functions make it fairly easy to enumerate proposals from best to worst. Secondly, agents are completely aware of their own preferences and the corresponding utility values.

2.8 Related Competitions

Some parts of the Trading Agent Competition (TAC) and the Agent Reputation Trust (ART) Competition also relate to automated agents negotiation. We provide a short description of both competitions and outline the differences with ANAC. For 2010, TAC is divided into three games [?, ?, ?]:

TAC SCM

TAC Supply Chain Management was designed to simulate a dynamic supply chain environment. Agents have to compete to secure customer orders and components required for production. In order to do so, the agents have to plan and coordinate their activities across the supply chain. Participants face the complexities of supply chains, which admits a variety of bidding and negotiation strategies.

TAC/AA

In the TAC Ad Auction, game entrants design and implement bidding strategies for advertisers in a simulated sponsoring environment. The agents have to bid against each other to get an ad placement that is related to certain keyword combinations in a web search tool. The search engine simulates clicks and sale conversions, yielding revenues for the advertiser. The advertiser strategies have to decide which keywords to bid on, and what prices to offer. Therefore, the strategies have to optimize their data analysis and bidding tactics to maximize their profit.

CAT

The CAT Competition or TAC Market Design is a reverse of the normal TAC game: as an entrant you define the rules for matching buyers and sellers, while the trading agents are created by the organizers of the competition. Entrants have to compete against each other to build a robust market mechanism that attracts buyers and sellers. The market needs to make profit by setting appropriate commission fees and at the same time attract enough traders, while also adapting to changing environmental conditions.

Some elements of TAC have similar challenges as posed by ANAC, especially the games of TAC SCM and TAC/AA. The games can get very complex and the domains of the games are specifically chosen to model a certain scenario of a trading agent problem. On the other hand, the entrants of ANAC have to consider very general negotiation domains when they design their agents. This makes it very easy to implement a simple agent that can participate in ANAC.

We believe that the ANAC is already very challenging despite its seemingly simple setup. The current design of ANAC poses a lot of challenges, such as a real timeline and the incomplete information of the opponent's preferences

(see also Section 2.7). Moreover, a transparent setup makes it easier to allow insight into the implications of design choices.

The Agent Reputation Trust (ART) Competition [?, ?] is also a negotiating agent competition with a testbed that allows the comparison of different strategies. The ART competition simulates a business environment for software agents that use the reputation concept to buy advices about paintings. Each agent in the game is a service provider responsible for selling its opinions when requested. The agent can exchange information with other agents to improve the quality of their appraisals. The challenge is to perceive when an agent can be trusted and to establish a trustworthy reputation.

Compared to ANAC, the focus of ART is more on trust: the goal is to perceive which agents can be trusted in a negotiation process and what reputation should be attributed to each agent.

3 Design of ANAC 2010

ANAC 2010 was held at the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10) in Toronto, Canada, with presentations of the participating teams and a closing discussion (see Section 5). AAMAS is a well-suited platform to host the competition, as it is the premier scientific conference for research on autonomous agents and multiagent systems, which includes researchers on automated negotiation. It brings together an international community of researchers that are well-suited to tackle the automated agents negotiation challenges posed by ANAC.

3.1 Teams

ANAC 2010 had seven participating teams from five different universities, as listed in Table 1.

IAMhaggler	University of Southampton
IAMcrazyHaggler	University of Southampton
Agent K	Nagoya Institute of Technology
Nozomi	Nagoya Institute of Technology
FSEGA	Babes Bolyai University
Agent Smith	TU Delft
Yushu	University of Massachusetts Amherst

Table 1 The participating teams of ANAC 2010.

3.2 *Domains*

We approached the design of ANAC to comply with the goals that were described in Section 1. Because ANAC is aimed towards multi-issue negotiations under uncertainty in open environments, we selected the following domains and profiles after the participating agents had been submitted. We aimed for a good spread of the relevant parameters, such as the number of issues, the number of possible proposals and the opposition of the domain (see Table 2).

Itex–Cypress

Our first scenario, taken from [?], is a buyer–seller business negotiation for one commodity. It involves representatives of two companies: Itex Manufacturing, a producer of bicycle components and Cypress Cycles, a builder of bicycles. There are four issues that both sides have to discuss: the price of the components, delivery times, payment arrangements and terms for the return of possibly defective parts. An example outcome would be:

(\$3.98, 45 days, payment upon delivery, 5% spoilage allowed).

The opposition is strong in this domain, as the manufacturer and consumer have naturally opposing needs and requirements. Altogether, there are 180 potential offers that contain all combinations of values for the four issues.

Zimbabwe–England

The second domain taken from [?, ?] involves a case where England and Zimbabwe are negotiating to reach an agreement in response to the world’s first public health treaty: the World Health Organization’s Framework Convention on Tobacco Control. The leaders of both countries must reach an agreement on five issues:

1. Funding amount

The total amount to be deposited into a fund to aid countries that are economically dependent on tobacco production. This issue has a negative impact on the budget of England and a positive effect on the economy of Zimbabwe. The possible values are no agreement, \$10, \$50 or \$100 billion. Thus, this issue has a total of four possible values.

2. Other aid programs

The impact on other aid programs. If other aid programs are reduced, then this will create economic difficulties for Zimbabwe. Possible values are:

- a. No reduction;
- b. Reduction equal to half of the fund;
- c. Reduction equal to the whole size of the fund;
- d. No agreement.

Thus, a total of four possible values are allowed for this issue.

3.&4. Trade barriers

Trade issues for both countries. Zimbabwe and England can use trade barriers such as tariffs (taxes on imports) or they can abstain from restrictive trade barriers to increase imports from the other party.

There is a trade-off in revenue of these policies: tariffs increases short-time revenue, but can lead to higher consumers prices. Decreasing import is good for local industries but it can decrease costumer welfare due to the increase in costumer costs. There are actually two issues here: the trade barriers that either side decides to use. Zimbabwe's possible values are divided between

- a. Reducing tariffs on imports;
- b. Increasing tariffs on imports;
- c. No agreement.

While England can choose between

- a. Reducing imports;
- b. Increasing imports;
- c. No agreement.

Thus, a total of three possible values are allowed for each of the two issues.

5. Creation of a forum

A forum can be created to explore other arrangements for health-issues. Zimbabwe would like to establish such a fund, to be able to apply to other global health agreements in the future, while this would be costly for England. The four possible values are:

- a. Creation of a fund;
- b. Creation of a committee that will discuss the creation of a fund;
- c. Creation of a committee that will develop an agenda for future discussions;
- d. No agreement.

Consequently, the domain has a total of $4^3 \cdot 3^2 = 576$ possible agreements. England and Zimbabwe have contradictory preferences for the first two issues,

but the other issues have options that are jointly preferred by both sides, making it a domain of medium opposition.

Travel

Our final domain has two persons negotiating to go on holiday to a location. From a small travel recommendation system we obtained multiple real-life profiles of travelers. They can each list their preferences on properties of a holiday destination:

1. Atmosphere.
2. Amusement.
3. Culinary.
4. Shopping.
5. Culture.
6. Sport.
7. Environment.

There are seven issues to discuss, all with a fairly large amount of choices. This leads to a big offers space of 188,160 possibilities. A sample negotiation outcome reads:

(Hospitable, Nightlife and entertainment, International cuisine, Small boutiques, Art galleries, Outdoor activities, Parks and gardens).

The opposition is weak in this domain, because traveling friends may have very compatible interests. Still the challenge is to find this optimal outcome in such a big search space.

	Itex–Cypress	Zimbabwe–England	Travel
Number of issues	4	5	7
Size	180	576	188,160
Opposition	Strong	Medium	Weak

Table 2 The domains used in ANAC 2010.

3.3 Deadline

We impose a real-time deadline on the negotiation process for reasons stated in Section 2.4. In ANAC 2010, the agents are bound to three minutes *each* to deliberate. This allowed every agent to utilize the same CPU time, but

it forces agents to keep track of the time that the opponent has left. This feature may undergo a small change in the next ANAC (see Section 5).

While the domains of the first ANAC competition did not include any discount factors, we do plan to add this feature to the next ANAC competition to be held in 2011 (see the discussion in Section 5).

4 Tournament Results

We describe the normalized domain scores of every agent in ANAC 2010 in Table 3. The normalized domain score is obtained by averaging the score against the other agents on multiple trials. All agents use both of the profiles that are linked to a domain (see Section 2.5 for more details on the scoring). The final score is listed in the last column, thus making *Agent K* the winner of ANAC 2010.

4.1 Overall Scoring

Rank	Agent	Score per domain			Avg
		Itex-Cyp	Eng-Zimb	Travel	
1	Agent K	0.901	0.712	0.685	0.766
2	Yushu	0.662	1.0	0.250	0.637
3	Nozomi	0.929	0.351	0.516	0.599
4	IAMhaggler	0.668	0.551	0.500	0.573
5	FSEGA	0.722	0.406	0	0.376
6	IAMcrazyHaggler	0.097	0.397	0.431	0.308
7	Agent Smith	0.069	0.053	0	0.041

Table 3 Normalized scores of every agent per domain.

For every domain, due to the normalization of the scores, the lowest possible score is 0 and the highest is 1. The fact that the maximum and minimum score are not always achieved, can be explained by non-deterministic behavior of the agents: the top-ranking agent on one domain does not always obtain the maximum score on every trial.

Agent K has won by a big margin, but it only managed to dominate on the Travel domain. On both Itex-Cypress and England-Zimbabwe, it earned second place after *Nozomi* and *Yushu*, respectively. Its consistent high scoring made *Agent K* the winner of ANAC. Only *IAMhaggler* managed to mirror this consistent scoring on all three domains.

4.2 Negotiation Strategies

We present a discussion of the strategies used by the participating agents. We compare the strategies by highlighting both common and contrasting approaches taken in the general strategic design. We are concerned with the following aspects of proposal strategies:

Proposal behavior

For every agent, we give a brief overview of the basic decisions that comprise the agents' inner proposal loop. We also describe the criteria for accepting an offer. Either of the two can be decided in a deterministic or non-deterministic manner.

Learning

In order to reach an advantageous negotiation agreement, it is beneficial to have as much information about the preference profile of an opponent as possible. If an agent can take into consideration the opponent's interests and learn during their interactions, then their utility might increase [?]. Because of the closed negotiation setting of ANAC, the negotiating parties exchange only proposals, but they do not share any information about their preferences. To overcome this problem, a negotiating agent may try to obtain a model of the preference profile of its opponent by means of learning.

For the participating agents, we are concerned how their strategies model the opponent

Timing aspects

There are substantial risks associated with delaying the submission of a proposal at the end of the negotiation. These risks arise from unpredictable delays and can cause proposals to be received when the game is already over. Agents can try to estimate the length of their negotiation cycles to cope with these risks. The agents can then concede in the final phase of the negotiation, or place their proposals in some calculated amount of time before the end. We examine whether the agents make any predictions on how many time is left and how they use this information.

Table 4 gives an overview of the strategies of all agents. In the "Time dependent" column, we address whether the proposal strategies keep track of the time that is left and change their proposals accordingly. The next column specifies what kind of learning method the agents use to generate the next offer. When agents decide to accept an offer, all take the offer's utility in account (U), but some of them also consider the remaining time (T).

Finally, most of the agents are non-deterministic. For example, *Agent K* may decide on a certain proposal target. But if it previously received even better offers *B*, then it will counteroffer a random offer taken from *B*. Otherwise, it will also select a random proposal; in this case it will choose any

offer that satisfies its proposal target. Most agents have this same mechanism: when they are indifferent between certain offers, they will choose randomly.

	Time dependent	Learning method	Acceptance Criteria	Deterministic
Agent K	Yes	All proposals	T/U	No
Yushu	Yes	Best proposals	T/U	No
Nozomi	No	Compromises	T/U	No
IAMhaggler	Yes	Bayesian	U	No
FSEGA	Yes	Bayesian	U	Yes
IAMcrazyHaggler	No	None	U	No
Agent Smith	Yes	Weights	T/U	Yes

Table 4 Strategies of the agents participated in ANAC.

We continue to report on the individual strategies of the ANAC agents, starting with the winner.

4.2.1 Agent K

The proposal mechanism of *Agent K* [?] works as follows: based on the previous proposals of the opponent and the time that is left, it sets a so-called *proposal target* (initially set to 1). If it already received an offer that matches at least the utility of the proposal target, it will offer this proposal to improve the chances of acceptance. Otherwise, it searches for random proposals that are at least as good as the proposal target. If no such proposals are found, the proposal target is slightly lowered.

The agent has a sophisticated mechanism to accept an offer. It uses the mean and variance of the utility of all received offers, and then tries to determine the best offer it might receive in the future and sets its proposal target accordingly. It then accepts or rejects the offer, based on the probability that a better offer might be proposed. For more information and technical details on Agent K, see [?].

4.2.2 Yushu

Yushu [?] is a fairly simple agent that makes use of a target utility to make its next offer. As a learning mechanism, it uses the ten best proposals made by the opponent, called *suggested proposals*. It also makes an estimate of how many rounds are still left for the negotiation. Combining this information, *Yushu* obtains the target utility. It also keeps track of the acceptability-rate:

the minimum utility it is willing to accept. To set the acceptability-rate, *Yushu* first finds the best possible utility that can be obtained in the domain, and accepts no less than 96% of it. When the number of estimated future rounds becomes short, this percentage is lowered to 92%.

The agent can only accept a proposal when the offered utility is above the target utility or when the utility reaches the acceptability-rate. Provided that either of the two is the case it accepts, when there are less than eight rounds left. When there is more time, it will accept only if it cannot find a suggested proposal with a better utility. If a better suggested proposal is available, it will offer that instead.

4.2.3 Nozomi

The proposal strategy of *Nozomi* [?] starts with an offer of maximum utility. It defines the gap between two parties as the differences in utility of their last offers. Depending on the gap and time that is left, it then chooses to make a certain proposal type, such as making a compromise, or staying put. *Nozomi* keeps track of the compromises made, but the agent does not model the utility function of the opponent.

The agent splits the negotiation into four intervals around 50%, 80% and 90% of the negotiation time. Based on previous offers, the gap between the two parties, and the time that is left in the negotiation, it will choose whether to accept an offer or reject it.

4.2.4 IAM(crazy)Haggler

IAMhaggler and *IAMcrazyHaggler* (cf. [?]) are both implementations of a framework called *SouthamptonAgent*, thus creating a lot of similarity between the two agents. The *SouthamptonAgent* provides standard methods for handling offers, proposing offers and keeping track of time. The framework is the only one that also keeps track of the time that the opponent uses.

IAMcrazyHaggler is a very simple take-it-or-leave-it strategy: it will make random proposals with a utility that is above a constant threshold, set to 0.9 (without discount factors it is set to 0.95). The proposal is done without regard to time or opponent moves.

IAMHaggler, on the other hand, is a fully fledged negotiation strategy, which incorporates a model of the opponent using Bayesian learning. It starts with a proposal of maximum utility and successively sets a target utility based on multiple factors, such as: the utility offered by the opponent, the time left for both agents, and the perceived opponent's profile, such as hard-

headedness. Upon receiving an offer, it analyzes the previous proposals of the opponent and adapts the hypotheses on the opponent's utility function. With this opponent model, it tries to find trade-offs that satisfy the target utility.

Let u be the utility of the last opponent's offer. Both agents accept an offer depending on u , namely when either of the following three conditions is met:

1. When u is at least 98% of the utility of its own previous offer.
2. When u is at least 98% of a *maximum aspiration* constant. The default value is 0.9, but if there are discount factors it is set to 0.85 for *IAMcrazy-Haggler* to make it reach an agreement sooner.
3. When u is at least 98% of the utility of its own upcoming offer.

Note that the three conditions only depend on the utility of the offer and not on the available time.

4.2.5 FSEGA

Similar to *Nozomi*, the *FSEGA* strategy [?] splits the negotiation into three intervals of time and applies different sub-strategies to each interval:

1. The first interval consists of the starting 85% of the negotiation time and is mainly used to acquire the opponent's profile from the counter-offers.
2. In the next 10%, the proposal strategy still does not concede, but relaxes some conditions for selecting the next proposal to improve the chances that the opponent accepts. The agent makes only small concessions and still tries to learn the opponent's profile.
3. In the final 5%, *FSEGA* considers the time restrictions and employs a concession-based strategy to select the next offer up to its reservation value.

In the first phase of the negotiation, the accept mechanism will admit any opponent offer that is 3% better than the utility of *FSEGA*'s last proposal. It will also always accept the best possible proposal. Otherwise, it selects a new proposal, but if the previous opponent's offer is better than the upcoming proposal it will accept it instead. After interval 1, it will also accept when it cannot find a better proposal for the opponent.

4.2.6 Agent Smith

Agent Smith [?] constructs an opponent model that represents the importance and preference for all values of each issue. The agent starts by making a first

proposal of maximum utility and subsequently concedes slowly towards the opponent.

The agent accepts an offer given the following circumstances. The agents' threshold for acceptance slowly decreases over time. In the last 10 seconds of the negotiation session, *Agent Smith* will propose the best proposal that the opponent already proposed (even when the offer is very bad for itself). Since it previously proposed it, it is likely for a rational opponent to accept this proposal. However, an error was made in the implementation, resulting in the fact that the agent already shows this behavior after two minutes instead of three. This explains the poor performance of the agent in the competition.

4.3 Timing Aspects

All agents of ANAC 2010, except for *IAMcrazyHaggler*, make concessions when the deadline approaches. Because a break-off yields zero utility for both agents, an agent that waits until the end of the negotiation takes a substantial risk. The other agent may not know that the deadline is approaching and may not concede fast enough. In addition, either the acceptance of a proposal or the (acceptable) counter-offer may be received when the game is already over. In the same manner, a real-time deadline also makes it necessary to employ a mechanism for deciding when to accept an offer.

We study the inclination of the agents of ANAC 2010 to exhibit either risk averse or risk seeking behavior regarding the timing of their proposals. In order to get a good picture of the risk management of the agents, we consider the number of break-offs that occur for every agent. Table 5 lists for each agent the percentage of negotiations that result in a break-off. All break-offs occur due to the deadline being reached or an occasional agent crash on a big domain.

Agent	Break-off percentage per domain			Avg
	Itex-Cyp	Eng-Zimb	Travel	
Agent K	22%	6%	63%	30%
Yushu	36%	0%	90%	42%
Nozomi	25%	17%	75%	39%
IAMhaggler	11%	0%	63%	25%
FSEGA	22%	0%	100%	41%
IAMcrazyHaggler	72%	23%	83%	59%
Agent Smith	0%	0%	98%	33%

Table 5 Percentage of all failed negotiations of every agent per domain.

The number of break-offs in the Travel domain stands out compared to the other domains. Recall that this is the biggest domain of ANAC 2010, with

188,160 possible proposals. Most of the agents had a lot of problems dealing with this domain. In a large domain it takes too much time to enumerate all proposals or to work with an elaborate opponent model. For example the *FSEGA* agent was unable to finish a single negotiation. Only *Agent K*, *Nozomi* and *IAM(crazy)Haggler* were able to effectively negotiate with each other on this domain.

With respect to the number of break-offs, *IAMHaggler* performs very well on all domains, while *IAMcrazyHaggler* ranks as the worst of all agents. This is to be expected, as its proposal generating mechanism does not take into account the time or the opponent (see Section 4.2.4 for an overview of its strategy). There is an interesting trade-off here: when *IAMcrazyHaggler* manages to reach an agreement, it always scores a utility of at least 0.9, but most of the time it scores 0 because the opponent will not budge.

The exact opposite of *IAMcrazyHaggler* is the strategy of *Agent Smith*. Because of an implementation error, *Agent Smith* accepts any proposal after two minutes, instead of three minutes. This explains why it did not have any break-offs on IteX–Cypress and England–Zimbabwe. The reason for the break-offs on the Travel domain is due to crashing of its opponent model. The importance of the timing aspects is underlined by the performance of *Agent Smith*: a small timing error resulted in very poor scoring on all three domains.

5 Design of Future ANAC

After ANAC 2010 was held at AAMAS-10, the participating teams had a closing discussion. This discussion yielded valuable suggestions for improving the design of future ANAC competitions. The consensus among participants was that the basic structure of the game should be retained. In the discussion between the participating teams and interested parties, we decided to leave further complicating factors out and not introduce too many innovations for the next year. This includes issue interdependencies, a richer negotiation protocol, different evaluation criteria, repeating scenarios (i.e.: multiple negotiation sessions), self-play and changes to the real-time deadline setup.

For the next ANAC in 2011 we decided that the teams participated in the first ANAC agree on the modifications to the rules and thus it was jointly agreed that the following modifications should be made into effect:

- Domains with discount factors should be included in the tournament.
- Changes should be made to the deadline setup and the selection criteria of the domains, that is, how to select a wide variety of domains without bias.

We detail the specific changes below.

5.0.1 Domains

ANAC 2011 will have domains that have discount factors. Without discount factors the current negotiation setup offers no incentive to accept an offer, except for right before the deadline. Waiting until the end of the round is an optimal strategy, except in the rare case that the opponent makes a mistake that might be retracted in the following round. Because of the lack of discount factors, 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 future ANAC setup could also be made more challenging by adding domains that contain continuous issues, such as real-valued price issues.

5.0.2 Issue Predictability

When learning the opponent's preference profile, a learning technique usually makes assumptions about the structure of the domain and preference profile (e.g., [?, ?, ?]). Negotiation strategies can try to exploit the internal structure of the issues in order to improve their proficiency. For example, a learning technique benefits from the information that a certain issue is *predictable*. Informally, an issue is called predictable when the global properties of its evaluation function is known. To illustrate, let us consider the discrete issue "Amount of funding" from the Zimbabwe–England domain (cf. Section 3.2). Its values are: *no agreement*, *\$10 billion*, *\$50 billion*, or *\$100 billion*. Even when we do not know which party we are dealing with, we can be confident that the utility of a particular value is either increasing or decreasing in the amount of funding. A price issue like this is typically predictable, where more is either better or worse for a negotiating party. Other issues, e.g. color, can be less predictable and therefore learning the preferences of color is more difficult. In order to improve the efficiency of the learning algorithms, we intend to eventually introduce (un)predictability labels.

5.0.3 Deadline

The real-time deadline of ANAC is considered a nice challenge to the competitors. 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 a future ANAC setup, we may choose a simpler protocol where both agents have a shared time window of three minutes.

6 Summary and Conclusion

This paper describes the first 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. Seven teams have participated in the first 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 GENIUS¹ 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 2011.

Acknowledgements

We would like to thank all participating teams of the University of Southampton, Nagoya Institute of Technology, Babes Bolyai and TU Delft and their representatives at AAMAS. We thank the ACAN and AAMAS organizers for enabling us to host ANAC. We are indebted to Dmytro Tykhonov his work on GENIUS and for helping organize this event. Furthermore, we would like to thank the sponsors of Bar-Ilan University and TU Delft.

¹ <http://mmi.tudelft.nl/negotiation/index.php/Genius>