

Towards a Quantitative Concession-Based Classification Method of Negotiation Strategies

Tim Baarslag, Koen Hindriks, and Catholijn Jonker

Man Machine Interaction Group
Delft University of Technology

{T.Baarslag,K.V.Hindriks,C.M.Jonker}@tudelft.nl

Abstract. In order to successfully reach an agreement in a negotiation, both parties rely on each other to make concessions. The willingness to concede also depends in large part on the opponent. A concession by the opponent may be reciprocated, but the negotiation process may also be frustrated if the opponent does not concede at all.

This process of concession making is a central theme in many of the classic and current automated negotiation strategies. In this paper, we present a quantitative classification method of negotiation strategies that measures the willingness of an agent to concede against different types of opponents. The method is then applied to classify some well-known negotiating strategies, including the agents of ANAC 2010. It is shown that the technique makes it easy to identify the main characteristics of negotiation agents, and can be used to group negotiation strategies into categories with common negotiation characteristics. We also observe, among other things, that different kinds of opponents call for a different approach in making concessions.

Keywords: Automated bilateral negotiation, Classification, Concession, Cooperation, Competition, Negotiation strategy.

1 Introduction

In bilateral negotiation, an opening offer is usually met with a counteroffer, and this then defines the initial bargaining range [12] of the negotiation. Sometimes the other party will immediately accept the offer, or will state that the set of demands is unacceptable, breaking off the negotiation. But usually, after the first round of offers, the question is: what concession is to be made next? One can choose to signal a position of firmness and stick to the original offer. Or one can take a more cooperative stance, and choose to make a concession. If one side is not prepared to make concessions, the other side must capitulate, or more commonly, the negotiation will end up in a break off.

The making of concessions is therefore central to a successful negotiation. Without them, negotiations would not exist [12]. Negotiation can even be defined in terms of making concessions: Pruitt [16] defines it as a process by which a joint decision is made by two or more parties that first verbalise contradictory

demands and then move towards agreement by a process of concession making or search for new alternatives.

Many of the classic negotiation strategies are characterized by this process of concession making throughout the negotiation. For example, the time dependent tactics such as Boulware and Conceder [5] are characterised by the fact that they steadily concede throughout the negotiation process. Other strategies, like behaviour dependent tactics (such as Tit for Tat) [1,4] base their concessions on the concessions of the other negotiating party.

The choice for what concessions to make depends in large part on the opponent. A concession by the opponent may be reciprocated by another concession, leading to a whole progression of concessions. On the other hand, the negotiation process can easily be frustrated if the opponent adopts a take-it-or-leave-it approach. Against this background, this paper studies negotiation strategies according to the way they concede towards different types of opponents.

This work advances the state-of-the-art in automated negotiation in the following ways. We present a new classification method for negotiation strategies, based on their pattern of concession making against different kinds of opponents. We introduce a definition of Concession Rate (CR) which measures the cooperativeness of an agent. We present a technique to quantitatively measure the CR against two extreme types of strategies: a take-it-or-leave-it strategy, and a conceding tactic. We then apply this technique to classify some well-known negotiating strategies, including the agents of ANAC 2010. This gives, for the first time, insight into the negotiation strategy space of the ANAC 2010 agents. It also aids our understanding of what concession making strategies are effective in settings such as ANAC.

In the discussion of our experimental results, we conclude that our technique has the desirable property of grouping negotiation strategies into categories with common negotiation characteristics. Among other things, we observe that different kinds of opponents call for a different approach in making concessions. For instance, a successful negotiating agent should behave competitively, especially against very cooperative strategies.

The remainder of this paper is organized as follows. Section 2 provides an overview of concession making in negotiation, including our adopted model of negotiation, and the definition of concession rate. In Section 3 we outline a method to compute the concession rate, followed by Section 4 that presents our experimental results. In Section 5 we discuss our findings, and finally, Section 7 presents our conclusions and our plans for future work.

2 Concession Making in Negotiation

In earlier work on conflict management through negotiation, the negotiation stance was characterized by two orientations: cooperative and competitive [3]. The theory relates to two basic types of goal interdependence that negotiators might have. It is either positive, where the negotiators' goals are linked in such

a way that their goal attainments are positively correlated (‘sink or swim together’), or the interdependence is negative, namely when the goal attainments are negatively correlated (‘when one swims, one sinks’).

However, a negotiator’s stance is usually not limited to one of the two orientations, because negotiation is a dynamic process and the position of the negotiators can change in response to the other party’s information or behavior [12]. In this paper, we take the stance that negotiators can exhibit a mixture of the two orientations, mainly depending on the type of opponent (see Figure 1). For example, a negotiator may cooperate with a cooperative opponent, but the same negotiator may be very competitive when facing competition. That is, in this case it *matches* the behavior of the opponent.

Conversely, a negotiator can be cooperative towards a competitive opponent and at the same time exploit cooperative opponents by playing competitive against them. In that case, it *inverts* the opponent’s behavior.




Fig. 1. The diagram of conceding behavior against both cooperative and competitive opponents

This way, we distinguish four types of negotiation orientations depending on the behavior against the opponent (see Table 1): *Inverter*, *Conceder*, *Competitor*, and *Matcher*. Every negotiation orientation corresponds to a different stance towards either of the two types of opponents. The main contribution of this paper is to define a formal, mathematical procedure for classifying agents into one of the four categories.

Table 1. Four types of negotiation orientations

Orientation	vs. Conceder	vs. Hardliner
Inverter	Exploiting	Yielding
Conceder	Cooperating	Yielding
Competitor	Exploiting	Competing
Matcher	Cooperating	Competing

2.1 Negotiation Model

We consider *bilateral* negotiations, i.e. a negotiation between two parties or agents A and B . The agents negotiate over *issues* that are part of a negotiation *domain*, 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 *outcome space*. The outcome space 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 normalized utility function U , which maps a possible outcome $\omega \in \Omega$ to a real-valued number in the range $[0, 1]$. In contrast to the outcome space, the preference profile of the agents 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 [17] for bilateral negotiation, in which the negotiating parties exchange offers in turns.

As in [18], we assume a common global time, represented here by $\mathcal{T} = [0, D]$. The alternating-offers protocol is supplemented with a deadline D at the end of the time line, so for any $t \in \mathcal{T}$, we stipulate that the deadline has been reached when $t = D$, at which moment both agents receive utility 0. This is the same setup as [6], with the exception that issues are not necessarily real-valued and both agents have the same deadline. We represent by $x_{A \rightarrow B}^t$ the negotiation outcome proposed by agent A to agent B at time t . A *negotiation thread* (cf. [5,18]) between two agents A and B at time $t \in \mathcal{T}$ is defined as a finite sequence

$$H_{A \leftrightarrow B}^t := \left(x_{p_1 \rightarrow p_2}^{t_1}, x_{p_2 \rightarrow p_3}^{t_2}, x_{p_3 \rightarrow p_4}^{t_3}, \dots, x_{p_n \rightarrow p_{n+1}}^{t_n} \right),$$

where

1. $t_k \leq t_l$ for $k \leq l$, the offers are ordered over time \mathcal{T} ,
2. $p_k = p_{k+2} \in \{A, B\}$ for all k , the offers are alternating between the agents,
3. All t_i represent instances of time \mathcal{T} , with $t_n \leq t$,
4. $x_{p_k \rightarrow p_{k+1}}^{t_k} \in \Omega$ for $k \in \{1, \dots, n\}$, the agents exchange complete offers.

Additionally, the last element of $H_{A \leftrightarrow B}^t$ may be equal to one of the particles $\{Accept, End\}$. We will say a negotiation thread is *active* if this is not the case.

When agent A receives an offer $x_{B \rightarrow A}^t$ from agent B sent at time t , it has to decide at a later time $t' > t$ whether to accept the offer, or to send a counter-offer $x_{A \rightarrow B}^{t'}$. Given a negotiation thread $H_{A \leftrightarrow B}^t$ between agents A and B , we can express the action performed by A with an *decision function* [6,18]. The resulting action is used to extend the current negotiation thread between the two agents. If the agent does not accept the current offer, and the deadline has not been reached, it will prepare a counter-offer by using a bidding strategy or *tactic* to generate new values for the negotiable issues.

Tactics can take many forms, e.g. time-dependent, resource dependent, imitative, and so on [18]. In our setup we will consider the tactics as given and try to categorize them according to their willingness to concede.

2.2 Concession Rate

In this section we introduce the notion of concession rate which quantifies the amount an agent has conceded towards the opponent during a negotiation. It is generally not enough to simply consider the utility of the agreement as a measure for the concession rate. For instance, a negotiator may not get an agreement before the deadline. In that case, both parties receive zero utility, but this gives no information about the concessions that were made. Therefore, we define the concession rate in terms of the minimum utility a negotiator has demanded during the negotiation.

Suppose a player A has a utility function U_A , mapping any outcome in Ω into the range $[0, 1]$. As we have assumed that the utility function is normalized in our setting, there will exist an *optimal outcome* $\omega_A^{\text{opt}} \in \Omega$ for which $U_A(\omega_A^{\text{opt}}) = 1$. In typical negotiation domains, the corresponding utility $U_B(\omega_A^{\text{opt}})$ of this outcome is far from optimal for player B , because the best outcome for A is typically not the best outcome for B . Player B should be able to always obtain at least this outcome in a negotiation, as A will always be inclined to accept it. We shall refer to this utility as the *full yield utility* (FYU_B) of player B (see Fig. 2). Intuitively, it is equal to his bottom line utility.

Note that an optimal outcome ω_A^{opt} is not necessarily unique, but typical domains (including those considered in ANAC and hence, in this paper) all have unique optimal outcomes for both players, so that the full yield utility is well-defined.

For any $t \in \mathcal{T}$, let

$$H_{A \rightarrow B}^t = \{x_{A \rightarrow B}^s \in H_{A \leftrightarrow B}^t \mid s \leq t\}$$

denote all bids offered by A to B until time t in an active negotiation thread. We can now formulate the minimum utility that agent A demanded during the negotiation thread $H_{A \rightarrow B}^t$. That is to say, we consider the largest concession the player has made so far:

$$\text{MIN}_A^t = \min\{U_A(x) \mid x \in H_{A \rightarrow B}^t\}$$

Informally, MIN_A^t denotes the lowest that A is willing to bid up until time t . The inverse of this is called the *yield* of player A . The lower player B is willing to go,

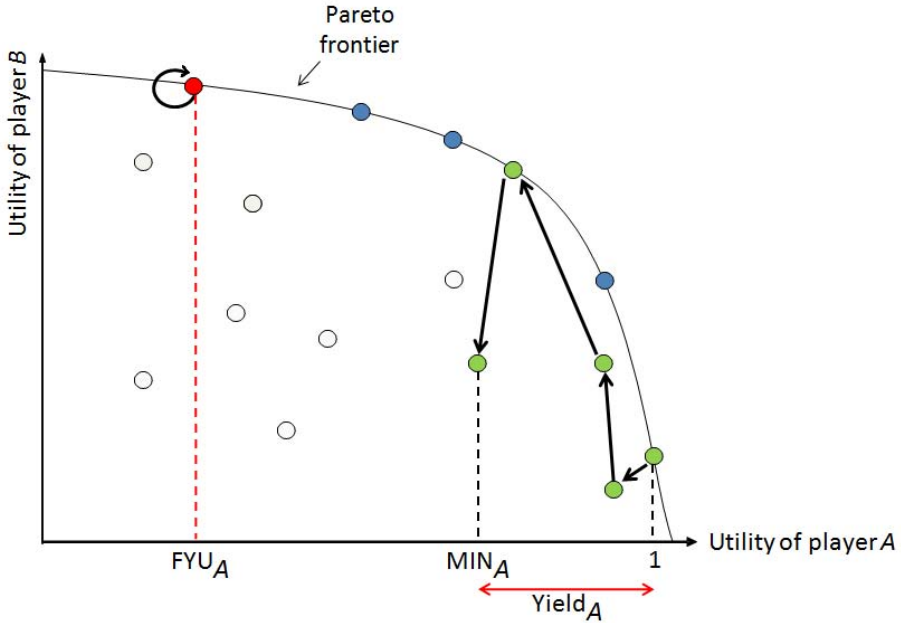


Fig. 2. The yield of player A is determined by MIN_A^t

the larger the yield. A yield of zero means the player has made no concession whatsoever (and therefore his demanded utility remains equal to one); A yield of $1 - \text{FYU}_A$ means the player has yielded fully (see Fig. 2). That is, it is defined as:

$$\text{Yield}_A^t = 1 - \max(\text{MIN}_A^t, \text{FYU}_A).$$

The Concession Rate $\text{CR}_A^t \in [0, 1]$ of player A up until time t is then simply the normalized yield:

$$\text{CR}_A^t = \frac{\text{Yield}_A^t}{1 - \text{FYU}_A}.$$

By normalizing, it is guaranteed that if $\text{CR}_A^t = 0$, then A has not conceded at all, while for $\text{CR}_A^t = 1$, player A has conceded fully (i.e., up to its full yield utility). Normalizing has the added benefit of reducing domain bias: in a typical competitive domain such as Itex-Cypress (defined in the experimental section below) players may obtain utilities anywhere between 0.2 and 1, while in very cooperative domains utilities may vary between 0.9 and 1. Normalization ensures that the concession rate can be compared over such different domains.

This paper only deals with the concession rate CR_A^D of a player A during the entire negotiation thread. We shall denote this simply by CR_A . We also omit the subscript A when it is clear from the context.

3 Method

In order to classify agents according to their concession rate, we considered a negotiation setup with the following characteristics. We selected a set of agents (introduced later) and let them negotiate against both a very cooperative and a very competitive opponent. The opponent tactics that we use to measure concession rates are simple, non-adaptive negotiation tactics. This ensures that the results depend as much as possible on the agent's own negotiating tactic. To be more precise, we aim for three opponent characteristics:

1. Simplicity

If the opponent negotiation tactic is simple and easy to understand, then the results depend on the agent's own negotiating tactic, which makes them easier to interpret.

2. Regularity

We want to give the agent enough time to show its bidding behavior; therefore, the opponent should not end the negotiation prematurely by either reaching an agreement too fast or breaking off the negotiation.

3. Deterministic behavior

In order to reduce variance in experimental results, we prefer deterministic agents to agents that demonstrate random bidding behavior.

For the *competitive opponent*, we chose Hardliner (also known as take-it-or-leave-it or Hardball [12]). This strategy simply makes a bid of maximum utility for itself and never concedes. This is the most simple competitive strategy that can be implemented and it fits the other two criteria as well: it is deterministic, and it gives the agent the full negotiation time to make concessions.

For the *cooperative opponent*, we selected Conceder Linear, i.e. the Time Dependent Tactic adapted from [6,5] with parameter $e = 1$. Depending on the current time $t \in [0, 1]$, this strategy makes a bid with utility closest to

$$P_{min} + (P_{max} - P_{min}) \cdot (1 - F(t)), \quad (1)$$

with

$$F(t) = k + (1 - k) \cdot t^{1/e}.$$

In this experiment, we selected the standard values $k = 0$, and P_{max}, P_{min} are respectively set to the maximum and minimum utility that can be obtained in the domain. With these values, and setting $e = 1$, we obtain a very simple conceding tactic. It reduces equation (1) to

$$P_{min} + (P_{max} - P_{min}) \cdot (1 - t),$$

so that it linearly reduces its demanded utility (from maximum P_{max} to minimum P_{min}) as time passes.

There exist even simpler conceding tactics such as Random Walker (which generates random bids), or an agent that accepts immediately. However, both opponent strategies are not *regular* in the sense that they do not give the agent

enough time to show its bidding behavior. Random Walker has the added disadvantage of not being deterministic. Therefore, we believe Random Walker can serve as a useful base line strategy, but not as a useful opponent to measure an agent’s willingness to concede. Consequently, we selected Conceder Linear as the cooperative opponent, as it fulfills the three requirements listed above.

We measured the concession rate of an agent A playing against the two agents in the following way. Suppose agent A negotiates with either Conceder Linear, or Hardliner. The two parties may attain a certain outcome, or reach the deadline. In both cases, at the end of the negotiation, A has reached a certain concession rate as defined in Section 2.2. The concession rate is then averaged over all trials on various domains (see Section 4.1), alternating between the two preference profiles defined on that domain. E.g., on the negotiation scenario between England and Zimbabwe, A will play both as England and as Zimbabwe.

4 Experiments

For our experimental setup we employed GENIUS (General Environment for Negotiation with Intelligent multi-purpose Usage Simulation) [13]. This environment, which is also used in ANAC, facilitates the design and evaluation of automated negotiators’ strategies. It can be used to simulate tournaments between negotiating agents in various negotiation scenarios, such as the setup described in this section. It supports the alternating offer protocol with a real-time deadline as outlined in our negotiation model. The default negotiation time in GENIUS and in the setup of ANAC is 3 minutes per negotiation session; therefore, we use the same value in our experiments.

4.1 Detailed Experimental Setup

Agents

In our experimental setup we included all the negotiation tactics that were submitted to The Automated Negotiating Agents Competition (ANAC 2010) [2]. ANAC is a negotiation competition aiming to facilitate and coordinate the research into proficient negotiation strategies for bilateral multi-issue negotiation, similar to what the Trading Agent Competition (TAC) has achieved for the trading agent problem [20]. The seven agents that participated in ANAC 2010 have been implemented by various international research groups of negotiation experts. We used these strategies in our experiments as they are representative of the current state-of-the-art in automated negotiation. In order of final ranking the strategies are: *Agent K*, *Yushu*, *Nozomi*, *IAMhaggler*, *FSEGA*, *IAMcrazy-Haggler*, and *Agent Smith*.

Table 2 gives a short overview of the variety of agent strategies used in our experiments. The “Time dependent strategy” column shows whether the 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 (more details are provided in [2]).

Table 2. Short overview of the strategies employed by the ANAC 2010 agents

	Time dependent strategy	Learning method	Deterministic
Agent K	Yes	All proposals	No
Yushu	Yes	Best proposals	No
Nozomi	No	Match compromises	No
IAMhaggler	Yes	Bayesian learning	No
FSEGA	Yes	Bayesian learning	Yes
IAMcrazyHaggler	No	None	No
Agent Smith	Yes	Learning weights	Yes

In addition to the ANAC agents, we included some well-known agents to explore some extreme cases. First, we included the Hardball strategy described in Section 3, which consistently makes the maximum bid for itself.

We also studied three members of the Time Dependent Tactics family [5] as defined above, namely: Boulware ($e = 0.2$), Conceder Linear ($e = 1$), and Conceder ($e = 2$). We included a variant of the Relative Tit-for-Tat agent from the same paper. This strategy, called Simple Nice Tit for Tat, tries to reproduce the behaviour that its opponent performed in the previous step.

Finally, we included the Random Walker strategy, also known as Zero Intelligence strategy [7], which randomly jumps through the negotiation space. It does not employ any information about its own preferences to make an offer.

Domains

The specifics of a negotiation domain can be of great influence on the negotiation outcome [8]; therefore, negotiation characteristics such as concession rate have to be assessed on negotiation domains of different size and competitiveness (or *opposition* [9]). With this in mind, we aimed for two domains (with two preference profiles each) with a good spread of negotiation characteristics. We picked two out of the three domains that were used in ANAC 2010 [2]. We omitted the third domain (Travel) as some of the ANAC agents did not scale well and had too many difficulties with it to make it a reliable testing domain.

Our first scenario is taken from [10], which describes a buyer–seller business negotiation. 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, including the price of the components, delivery times, etc. The opposition between the parties is strong in this domain, as the manufacturer and consumer have naturally opposing requirements. Altogether, the Itex–Cypress domain contains 180 potential offers.

The second domain taken from [13,14] involves a case where England and Zimbabwe negotiate an agreement on tobacco control. The leaders of both countries must reach an agreement on five issues. The England–Zimbabwe is of medium opposition, because the parties have contradictory preferences for some issues, but other issues have options that are jointly preferred by both sides. The domain has a total of 576 possible agreements.

Table 3. The four preference profiles used in experiments

	Itex-Cyp	Eng-Zim
Size	180	576
Opposition	Strong	Medium

4.2 Experimental Results

We present the results of the experiments in Table 4 and its graphical representation is depicted in Fig. 3.

Table 4. An overview of the concession rate of every agent in the experiments

Agent	CR vs. Conceder	CR vs. Hardliner
<i>Agent K</i>	0.12	0.18
<i>Agent Smith</i>	0.46	1.00
Boulware	0.14	1.00
Conceder Linear	0.43	1.00
Conceder	0.63	1.00
<i>FSEGA</i>	0.33	0.76
Hardliner	0.00	0.00
<i>IAMcrazyHaggler</i>	0.05	0.05
<i>IAMhaggler</i>	0.02	0.27
<i>Nozomi</i>	0.20	0.22
Random Walker	0.97	1.00
Simple Nice Tit for Tat Agent	0.42	0.01
<i>Yushu</i>	0.11	0.95

Extreme Cases

The Hardliner strategy and the Random Walker strategy are at the opposite sides of the spectrum. Hardliner will not concede to any type of strategy, so by definition it has $CR = 0$ against both Hardliner and Conceder. Consequently, Hardliner defines the most competitive strategy possible.

On the other hand, Random Walker will make arbitrary concessions given enough time. This makes Random Walker one of the most cooperative strategies possible. Against Hardliner, there is plenty of time for Random Walker to randomly produce a bid with which it fully concedes, so it has $CR = 1$ against Hardliner. Against Conceder, it may not have time to fully concede, but it generally will produce offers of very low utility in this case as well, resulting in a CR of 0.97.

We considered three members of the Time Dependent Tactics family: Boulware ($e = 0.2$), Conceder Linear ($e = 1$), and Conceder ($e = 2$) who are all in the top of the chart because they have a CR equal to 1 versus the Hardliner. In addition to the time dependent tactics, at the top of the chart we see two more strategies: *Agent Smith* and Random Walker. This means all these strategies give in fully to Hardliner and are thus fully exploited by a strategy that does




Fig. 3. A graphical overview of the concession rates of the agents

not give in at all. All of these five strategies have a very simple bidding strategy and are apparently not optimized to deal with very uncooperative opponents.

5 Discussion

This section makes observations regarding the clustering of different strategies in Fig. 3, and then classifies them into the four negotiation orientations we have discussed previously.

5.1 Clustering

Agent Smith and *Conceder Linear* are very close in the chart and this is no coincidence: *Agent Smith* uses essentially the same strategy as the linear *Conceder*, by first making a proposal of maximum utility and subsequently conceding linearly towards the opponent.

The same holds for *Yushu* and *Boulware*: the strategies are very similar, as is indicated by their close vicinity in the chart. Like *Boulware*, *Yushu* adopts a very competitive time dependent strategy, making larger concessions when the negotiation deadline approaches. Both adopt a conservative concession making strategy and are not willing to make large concession at the beginning, but prefer to wait for their opponent to make concessions.

These two examples show that this chart can be useful to cluster strategy types, as similar strategies have similar concession characteristics.

Clustering also occurs on lines in the chart. For example, the three Time Dependent Tactics all share the same behavior against Hardliner: all three ultimately give in to it. This is to be expected, as any agent from the time dependent family will offer the reservation value when the deadline is being reached [6], resulting in full concession to the opponent. In general, all Time Dependent Tactics will lie on the line $CR = 1$ against Hardliner.

Against a more cooperative strategy like Conceder, the results are also intuitively clear: concessions get bigger when the parameter $e \in (0, \infty)$ gets bigger, so Boulware concedes the least, while Conceder concedes the most. More generally, when $e \rightarrow 0$, then $CR \rightarrow 0$. Conversely, when $e \rightarrow \infty$, then $CR \rightarrow 1$.

Finally, there is a big cluster of strategies in the left part of the chart, which is populated by the top four strategies in ANAC: *Agent K*, *IAMhaggler*, *Nozomi* and *Yushu*. The better performing strategies of ANAC have different approaches towards the Hardliner, but they seem to have one trait in common: they all concede very little to the Conceder. In other words: they exploit the Conceder by waiting for even bigger concessions. The fact that these strategies did very well in ANAC seems to indicate that in order to be successful in an automated negotiation competition, an agent should behave competitively, especially against very cooperative strategies.

5.2 Four Negotiation Orientations

We classify the different agent strategies of Fig. 3 into the four negotiation orientations of Fig. 1. This procedure is necessarily arbitrary; nevertheless, we propose the following grouping.

The top left agents in the diagram can be considered to be Inverters: *Yushu*, *Boulware*, and *FSEGA*. The remaining agents in the top right are then Conceders, namely: *Conceder Linear*, *Agent Smith*, *Conceder*, and *Random Walker*.

The Simple Nice Tit for Tat strategy is the only strategy that can be considered a Matcher, i.e.: it does not concede to a Hardliner, but it does concede to the Conceder. Clearly, this is to be expected from a Tit for Tat strategy, as it is based on cooperation through reciprocity: it matches whatever the other player did on the preceding move. The fact that this type of strategy does not occur naturally in ANAC can be explained by our previous comments on clustering: following a Tit for Tat strategy is not as successful in negotiation, because it does not exploit the conceding strategies.

All of the remaining strategies are Competitors, i.e. they do not concede much, whether it is against a cooperative or a competitive agent. The majority of strategies that performed well during ANAC are located in this region. Again, we observe that the successful strategies are very competitive.

6 Related Work

This paper is inspired by ideas presented in [11] (of which parts originally appeared in unpublished work by Kersten in 2005). In [11], four dual negotiation

orientations are distinguished, depending on the negotiator's own orientation and that of the negotiating partner. Both orientations can be either competitive or cooperative, leading to four different labels: Competitor, Yielder, Exploiter, and Cooperator. We re-use these labels to name the stance of a negotiator against different kinds of opponents (see Fig. 1). However in our work, the negotiators are assumed to have different responses to different observed behavior by the other party. Therefore, instead of the negotiator having one particular stance during the negotiation, the position of the negotiators can change in response to the competitiveness of the opponent. For example, a negotiator may be *both* an Exploiter (against a Cooperator), and a Yielder (against a Competitor). The negotiator would then be called an Inverter, as he takes on the reverse role of his opponent.

In [15], a classification scheme is given for electronic commerce negotiation, including characteristics of the negotiating agents. It is argued that agents can act in a self-interested way, or altruistically, or strike a balance in between. This choice is then seen as a component of the bidding strategy of the agent, which ultimately decides how and when to place offers, or when to withdraw, etc. Although the paper makes this distinction in bidding characteristics, it does not provide a definition or a way to quantify them.

Thomas [19] redefines five conflict-handling modes that can be applied to negotiation: competing, collaborating, compromising, avoiding and accommodating. Similar to our work, the classification method uses two underlying dimensions. However, the underlying dimension are different, namely: assertiveness (attempting to satisfy one's own concerns), and cooperativeness (attempting to satisfy other's concerns). This classification method is phrased in qualitative, intentional terms of the conflict-handler. Similarly, [22] distinguishes negotiation strategies into two strategy types: distributive and integrative. This description also focuses on the approach used by the negotiators. Our paper has a different focus from both papers, centering around quantitative negotiation characteristics *in response* to agents having either high and low concession rates. Furthermore, we do not classify negotiation strategies in a binary way (either cooperative or non-cooperative), but we employ a continuous spectrum in our approach.

Currently, there are two papers that analyze the results of ANAC 2010. Baarslag et al. [2] give a short overview of all negotiation tactics and their rankings in the tournament, but they do not provide an in-depth analysis of the bidding behavior of the ANAC participants. Williams et al. [21] consider self-play scores of the agents and also perform an empirical game theoretic analysis of the tournament results. This work focuses on stability of a strategy in a tournament with different mixes of opponent strategies, but unlike our work, it does not discuss or aim to classify the characteristics of the agent's negotiation strategies.

7 Conclusion and Future Work

Making concessions during a negotiation is vital for getting an agreement. Successful negotiations are only possible when both parties employ an effective

conceding strategy. Designing a good strategy of when and how much to concede is challenging, and that is why there are many current negotiation implementations that concede in very different ways.

In this paper, we aimed to classify a selection of current automated negotiation strategies according to their pattern of concession making in order to gain insight into their negotiation characteristics. We first formally defined the notion of concession rate which gives a normalized measure of the largest concession that was made during the negotiation. This formalizes the concept of an agent's willingness to concede against different opponents.

We then presented an empirical method to effectively compute the concession rate of agents, and then applied our approach to a selection of well-known agents (including all participants of ANAC 2010) in an experimental setting. For the first time, this gives insight into the strategy space of negotiation tactics employed in ANAC 2010. We subsequently used our method to classify the agents into four categories types of concession behavior.

In addition to classifying agent strategies, we have drawn various conclusions based on charting the experimental results. We have seen that there is indeed a wide spread in concession rates of current agents. We established that the chart can be useful to cluster strategy types, as similar strategies have similar concession characteristics. Secondly, it makes it easy to understand the agent's main negotiation characteristics at a glance.

Some extreme agents are located in the extreme regions of the chart, while the stronger agents form a cluster in the competitive corner. The results indicate that in order to be successful in an automated negotiation competition, an agent should not concede much, especially not to very cooperative strategies.

While making a number of contributions, this paper also opens up some lines of future work. We plan to conduct a deeper investigation of the impact that concession rates have on tournament results. The focus of this paper is on the tournament setting of ANAC 2010, but it would be interesting to extend the ideas presented in this paper to the results of ANAC 2011, especially because the 2011 competition contains negotiation domains that have discount factors. As discount factors devalue utility with the passing of time, they require the negotiating agents to give even more consideration to effective concession-making.

Secondly, the focus of this paper has been on bidding behavior, and not on acceptance strategy. In general, this is an important part of a negotiator's strategy that also highly influences the outcome of a negotiation. We believe the same interactions between cooperation and competition play a role when agents decide when and whether to accept. This could provide an interesting addition to our work, which we plan to examine in future research.

Acknowledgements. 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.

References

1. Axelrod, R.: *The Evolution of Cooperation*. Basic Books (1984)
2. Baarslag, T., Hindriks, K., Jonker, C.M., Kraus, S., Lin, R.: The first automated negotiating agents competition (ANAC 2010). In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T., Yamaki, H. (eds.) *Innovations in Agent-Based Complex Automated Negotiations*. SCI, vol. 319. Springer, Heidelberg (2010) (to appear)
3. Deutsch, M., Coleman, P.T., Marcus, E.C.: *The Handbook of Conflict Resolution: Theory and Practice*, 1st edn. Jossey-Bass (April 2000)
4. Faratin, P., Sierra, C., Jennings, N., Buckle, P.: Designing flexible automated negotiators: Concessions, trade-offs and issue changes. Tech. rep. (1999)
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. Fatima, S.S., Wooldridge, M.J., Jennings, N.R.: Optimal Negotiation Strategies for Agents with Incomplete Information. In: Meyer, J.-J.C., Tambe, M. (eds.) *ATAL 2001*. LNCS (LNAI), vol. 2333, pp. 377–392. Springer, Heidelberg (2002)
7. Gode, D.K., Sunder, S.: Allocative efficiency in markets with zero intelligence (zi) traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101(1), 119–137 (1993)
8. Hindriks, K.V., Tykhonov, D.: Towards a Quality Assessment Method for Learning Preference Profiles in Negotiation. In: Ketter, W., La Poutré, H., Sadeh, N., Shehory, O., Walsh, W. (eds.) *AMEC 2008*. LNBIP, vol. 44, pp. 46–59. Springer, Heidelberg (2010)
9. Kersten, G., Noronha, S.: Rational agents, contract curves, and inefficient compromises report. Working papers, International Institute for Applied Systems Analysis (1997), <http://econpapers.repec.org/RePEc:wop:iasawp:ir97050>
10. Kersten, G.E., Zhang, G.: Mining inspire data for the determinants of successful internet negotiations. *InterNeg Research Papers INR 04/01 Central European Journal of Operational Research* (2003)
11. Lai, H., Doong, H.S., Kao, C.C., Kersten, G.: Negotiators' communication, perception of their counterparts, and performance in dyadic e-negotiations. *Group Decision and Negotiation* 15, 429–447 (2006)
12. Lewicki, R.J., Saunders, D.M., Minton, J.W.: *Essentials of Negotiation*. McGraw-Hill, Boston (2003)
13. Lin, R., Kraus, S., Tykhonov, D., Hindriks, K., Jonker, C.M.: Supporting the design of general automated negotiators. In: *Proceedings of the Second International Workshop on Agent-based Complex Automated Negotiations, ACAN 2009* (2009)
14. Lin, R., Kraus, S., Wilkenfeld, J., Barry, J.: Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence* 172(6-7), 823–851 (2008)
15. Lomuscio, A., Wooldridge, M., Jennings, N.: A Classification Scheme for Negotiation in Electronic Commerce. In: Dignum, F., Sierra, C. (eds.) *AgentLink 2000*. LNCS (LNAI), vol. 1991, pp. 19–33. Springer, Heidelberg (2001)
16. Pruitt, D.G.: *Negotiation Behavior*. Academic Press (1981)
17. Rubinstein, A.: Perfect equilibrium in a bargaining model. *Econometrica* 50(1), 97–109 (1982), <http://www.jstor.org/stable/1912531>
18. Sierra, C., Faratin, P., Jennings, N.: A service-oriented negotiation model between autonomous agents. In: Boman, M., Van de Velde, W. (eds.) *MAAMAW 1997*. LNCS (LNAI), vol. 1237, pp. 17–35. Springer, Heidelberg (1997)

19. Thomas, K.W.: Conflict and conflict management: Reflections and update. *Journal of Organizational Behavior* 13(3), 265–274 (1992)
20. Wellman, M.P., Wurman, P.R., O'Malley, K., Bangera, R., de Lin, S., Reeves, D., Walsh, W.E.: Designing the market game for a trading agent competition. *IEEE Internet Computing* 5(2), 43–51 (2001)
21. Williams, C.R., Robu, V., Gerding, E.H., Jennings, N.R.: Using gaussian processes to optimise concession in complex negotiations against unknown opponents. In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. AAAI Press (January 2011)
22. Zachariassen, F.: Negotiation strategies in supply chain management. *International Journal of Physical Distribution and Logistics Management* 38, 764–781 (2008)