Predicting Agents Tactics in Automated Negotiation

Chongming Hou
Knowledge Media Institute
The Open University
Milton Keynes, MK7 6AA
c.m.hou@open.ac.uk

Abstract

This paper presents a learning mechanism that applies nonlinear regression analysis to predict a negotiation agent's behaviour based only the opponent's previous offers. The behaviour of negotiation agents in this study is determined by their tactics in the form of decision functions. Heuristics based on estimates of an agent's tactics are drawn from a series of experiments. The findings of this empirical study show that this approach can be used to obtain better deals than existing decision function tactics. The learning mechanism can be used online, without any prior knowledge about the other agents and is therefore, very useful in open systems where agents have little or no information about each other.

1. Introduction

Negotiation is a process of joint decision making between two or more parties in an effort to resolve their conflicting demands. Negotiation has been treated formally by researchers in economics and game theory, and informally (i.e. based on observations) by researchers in industrial relations, international relations and counselling. This paper focuses on the study of two-party negotiation, which is the subject of a great deal more empirical research than the multiparty case [6]. Moreover, multiparty negotiation can be described as multiple, mutually influencing, two-party negotiations over multiple issues [3]. In two-party negotiation, the two agents play opposing roles, such as buyer and seller. Our research on negotiation lies between the fully co-operative and fully competitive negotiations.

In electronic commerce, the task of negotiation can be delegated to a software agent in order to save human users time on activities which are either routine or demanding. To get better individual or social outcomes, the software agents require appropriate tactics. A tactic is the decision policy for choosing actions in different situations. Because negotiation is an interactive process, the outcome is not only determined by an agent's own tactic but it is also influenced by the other agent's choices. This characteristic makes it difficult to find an optimal tactic. The research presented here focuses on the online prediction of the other agent's tactic in order to reach better deals in negotiation.

The following section reviews some of the related research on negotiation tactics. Section 3 explains the motivation for predicting an opponent's negotiation tactic. The use of nonlinear regression for estimating the family, form and parameters of an opponent's tactic is introduced in Section 4. Section 5 presents a set of heuristics for identifying an opponent's negotiation deadline and reservation value using the results of the nonlinear regression analysis. To test the performance of the prediction mechanism against other prevailing tactics, a set of two-party negotiations were carried out and the results are described in Section 6. Section 7 discusses the use of the proposed prediction mechanism. Section 8 highlights the main conclusions to be drawn from this work and Section 9 discusses some potential areas for future research.

2. Related work

The related work carried out in game theory is presented in the following subsection. This work typically assumes that each agent has complete knowledge of the other agent's actions and that each agent has unbounded computational power to explore the space of acceptable deals. However, these theoretical assumptions are not necessarily true in practice. Due to the failing assumptions, decision functions were proposed that would enable an agent to generate offers according to the time remaining, the resources available, or, the behaviour of an opponent. The three families of decision functions and their common forms are described. The section concludes with an overview of some related work on exploring optimal strategies and effective tactics.

2.1. Game theory

Game theory treats the negotiation as a kind of game and negotiating agents as the players in a game. Game theory provides formal concepts to analyse the strategic interaction among agents in negotiation [1]. However, game theory has two fundamental assumptions: common knowledge and perfect computational rationality. In the first assumption, all information about the possible strategies, the outcome with each configuration of strategies, an agent's attitude towards time-related risks, etc., are common knowledge known to each agent. Perfect computational rationality assumes a negotiation agent has unbounded computational power, that it seems no computation is required to find the space of mutually acceptable solutions, because agents in game theory can instantly find the optimal strategies at the beginning of the game with this unbounded computational power. These assumptions do not necessarily hold for real-world negotiation and therefore make it difficult to apply game theory in practice.

2.2. Decision functions as negotiation tactics

Since an agent has limited computational power and incomplete knowledge about other agents and, it is necessary for an agent to produce offers based on their own criteria, such as, time limits or resource availability. Using this approach, [2] proposed a negotiation model and three families of negotiation tactics, namely: time-dependent, resource-dependent and behaviour-dependent tactics.

2.2.1. Negotiation model. In this negotiation model, two parties negotiate on an issue, such as price, delivery time, quality, etc. This paper focuses on single issue that has continuous value. The two parties adopt two conflicting roles, such as the buyer and seller of goods or services. The negotiation is a process of two parties making alternate offers. Let $x^t_{a\to b}$ is the offer proposed by seller *s* to buyer *b* for a negotiation issue at time *t*. $x^t_{s\to b}$ is delimited by [min^s, max^s], the range of all possible offers by seller. The negotiation is to determine a value *x* ($x \in [min^s, max^s] \cap [min^b, max^b]$) which is mutually acceptable to buyer and seller. max^b is actually the reservation value of buyer that any value larger than max^b won't be accepted by buyer. min^s is the reservation value of seller and any value smaller than min^s won't be acceptable to seller. Each agent *a* has a scoring (or utility) function $V^a : D^a \to [0,1]$ that assigns a score to value *x* in D^a. The scoring functions are either monotonically increasing for seller or decreasing for buyer.

A negotiation thread between agents *a* and *b*, at time t_n, noted $X^{t_n}_{a\leftrightarrow b}$, is a sequence of alternating offers of the form ($x^{t_1}_{a\to b}, x^{t_2}_{b\to a}, x^{t_3}_{a\to b}, \cdots$) where $t_{i+1} > t_i$. The negotiation is terminated by *accept or withdraw.* Agent *a*'s response at t_n to agent *b*'s offer $x^{t_{n-1}}_{b\to a}$ sent at time t_{n-1} is defined as:

$$\text{response}^a(t^n, x^{t_{n-1}}_{b\to a}) = \begin{cases} withdraw(a,b) & if\, t^n > t^a_{\max} \\ accept(a,b,x^{t_{n-1}}_{b\to a}) & if\, V^a(x^{t_{n-1}}_{b\to a}) \geq V^a(x^{t_n}_{a\to b}) \\ offer(a,b,x^{t_n}_{a\to b}) & otherwise \end{cases}$$

$x^{t_n}_{a\to b}$ is the counter-offer *a* offers to *b* when offer $x^{t_{n-1}}_{b\to a}$ is rejected by *a*. $t^a_{\max}$ is *a*'s deadline by which *a* must have completed the negotiation.

(Counter-) offers are generated by functions called tactics. A *tactic* generates a value for a single negotiation issue based on a single criterion such as the time remaining, the resource remaining, or the opponent's behaviour.

2.2.2. Time-dependent family of tactics. In this family of tactics time is the predominate factor. These tactics model the fact that the agent is likely to concede more rapidly as the deadline approaches. All of the tactics in this family prescribe an agent *a* to concede to their reservation value by their time deadline $t^a_{\max}$. What differentiates them is the shape of their concession curve (see Figure 1). The offer of agent *a* to agent *b* at time $t \leq t^a_{\max}$, is modelled by a function α^a which is dependent on time:

$$x^t_{a\to b} = \begin{cases} \min^a + \alpha^a(t)(\max^a - \min^a) & if\, V^a\ decreasing \\ \min^a + (1-\alpha^a(t))(\max^a - \min^a) & if\, V^a\ increasing \end{cases}$$

where $0 \leq \alpha^a(t) \leq 1$, and can be in one of the following two forms:

(1) Polynomial: $$\alpha^a(t) = k^a + (1-k^a)\left(\frac{\min(t, t^a_{\max})}{t^a_{\max}}\right)^{\frac{1}{\beta}}$$

(2) Exponential: $$\alpha^a(t) = e^{(1-\frac{\min(t,\, t^a_{\max})}{t^a_{\max}})^\beta \ln k^a}$$

where k^a ($\in [0,1]$) determines the initial offer at t=0 in $[min^a, max^a]$ since $\alpha^a(t=0) = k^a$.

Both forms are parameterised by a value $\beta \in \Re^+$ that determines the rate an agent approaches to their reservation value. The expressions above represent an infinite number of possible tactics, one for each value of β (see Figure 1).

There are three types of tactics in this family: *Boulware* (*β* « 1) where an agent does not start conceding until the deadline is nearly up. *Conceder* (*β* » 1) where an agent will start giving ground fairly quickly. And, *Linear* (*β* = 1) where an agent concedes the same amount in each round of the negotiation.

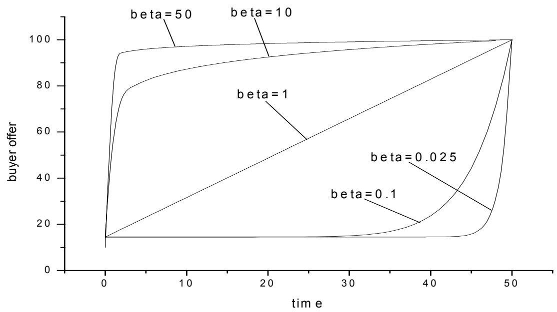

Figure 1. Five examples of concession curves for the polynomial time-dependent family of tactics.

2.2.3. Resource-dependent family of tactics. These tactics generate offers depending on how a particular resource is being consumed. They become progressively more conciliatory as the quantity of resource diminishes.

$$\alpha^a(t) = \kappa^a + (1-\kappa^a)e^{-resource^a(t)}$$

where $resource^a(t) = \frac{\mu^a}{|X^t_{a\leftrightarrow b}|}$, and μ^a is the time agent a considers reasonable to spend on negotiation, and $|X^t_{a\leftrightarrow b}|$ is the number of messages exchanged in the negotiation, i.e. the communication cost.

2.2.4. Behaviour-dependent family of tactics. These tactics base their actions on their opponent's behaviour. They differ in which aspects of the opponent's behaviour they imitate and to what degree.

Relative Tit-For-Tat (RelTFT): Agent *a* reproduce, in percentage terms, the behaviour that their opponent *b* exhibited in the previous $\delta_j \geq 1$ rounds.

$$x^{t_{n+1}}_{a\to b} = \min(\max(\frac{x^{t_{n-2\delta}}_{b\to a}}{x^{t_{n-2\delta+2}}_{b\to a}} x^{t_{n-1}}_{a\to b}, min^a), max^a)$$

Random Absolute Tit-For-Tat (RndTFT): Same as Relative Tit-For-Tat, except that the behaviour is imitated in absolute terms.

$$x^{t_{n+1}}_{a\to b} = \min(\max(x^{t_{n-1}}_{a\to b} + (x^{t_{n-2\delta}}_{b\to a} - x^{t_{n-2\delta+2}}_{b\to a}) + (-1)^s R(M), min^a), max^a)$$

where R(*M*) is a function that generates a random value in the interval [0, *M*], and $s_j = 0$ if V^a_j is decreasing and $s_j = 1$ if V^a_j is increasing.

Average Tit-For-Tat (AvgTFT)*:* Uses the average of the percentage change in a window of $\lambda_j \geq 1$ of the opponent's history:

$$x^{t_{n+1}}_{a\to b} = \min(\max(\frac{x^{t_{n-2\gamma}}_{b\to a}}{x^{t_n}_{b\to a}} x^{t_{n-1}}_{a\to b}, min^a), max^a)$$

2.3. Optimal strategies and effective tactics

Since decision functions were proposed in [2] as negotiation tactics, others have extended the research on the use of these three families of tactics.

The work in [4] concerns the effect of time on the negotiation outcome between agents adopting the time-dependent family of tactics. In [4] an agent holds a set of possible values for the opponent's reservation value and deadline along with a binary probability distribution over these values. Based on the expected utility under the probability distributions, [4] give conditions for the convergence of optimal strategies. The effectiveness of this approach relies heavily on the quality of the probability distribution used.

The research in [2] examined the relative performances of the three families of tactics using a genetic algorithm. They produced a set of average pair-wise comparative performances between tactics. These performance values provide a basis for choosing an appropriate tactic for an agent, when negotiating with another agent whose tactic is known. Since the tactic is private information, an agent needs to learn this information during the negotiation in order to apply this approach.

3. Proposed Approach

It is argued here, that in order to obtain a better outcome in negotiation, an agent needs to find out some information about their opponent. However, in competitive bilateral negotiation, due to the diversity and complexity of negotiation in practice, a negotiator has little information about their opponent. Under this uncertainty, although an agent can choose a tactic with best performance in average, such as resource tactic (with larger μ) [5] or Linear tactic [3], both tactics can not guarantee an agent a best deal. Since there is no optimal tactic for all negotiation environments, an agent has to choose the most effective tactic when facing agents with different tactics. Typically, an agent does not declare their tactic. In fact, the only available information in most cases is the previous offers from the other negotiating agent.

The research here explores the following questions: If an opponent adopts one of the above widely researched tactics, can the tactic's family and form be identified, and furthermore, can the decision parameters in the tactic be estimated from their previous offers? If an agent can predict these kinds of information, the information can then be used to determine the agent's responses in order to achieve the optimal outcome for the predicting agent.

The proposed approach uses nonlinear regression to learn an opponent's tactic. Using the results obtained an agent can:

Discriminate which of the three families of tactics their opponent has adopted (i.e. time-dependent, resource-dependent or behaviour-dependent tactics).

Predicting the opponent's offers.

Identify the opponent's tactic type in the early stages of negotiation (e.g. Boulware, Conceder or Linear in the time-dependent family).

Using the identified tactic type, estimate the opponent's reservation value and deadline.

The information about the opponent's tactic type can be used to guide the predicting agent's choice of tactic, based on the existing research results such as the pair-wise performance table presented in [5] or the hypotheses described in [3].

Furthermore, by identifying the opponent's deadline and reservation value, the predicting agent can either avoid breakdowns or terminate unprofitable negotiations.

For example, in the case where the buyer's predicted offer at the seller's deadline is less than the seller's reservation value breakdown cannot be avoided, therefore, the seller should withdraw from the negotiation. In cases where negotiation carries a communication cost and the buyer's predicted offer at the seller's deadline, minus the communication cost involved, is less than the buyer's current offer, terminating the unprofitable negotiation will avoid any further costs.

By adopting an effective tactic, and minimising the number of negotiation breakdowns and terminating unprofitable negotiations, it is argued that an agent can significantly improve their performance.

4. Nonlinear regression to estimate an agent's tactics

If there is sufficient number of previous offers from the other negotiating agent who adopts a tactic in one of the forms described in Section 2.2, this section will discuss how to use nonlinear regression to predict the opponent's tactic and estimates its parameters.

Among the three families of tactics in Section 2.2, the time-dependent and resource-dependent tactics are decision functions which can be represented as

$$y_i = f(\beta_0, \beta_1, ..., \beta_n; t_i) + e_i \qquad (1)$$

where y_i is the offer generated by a decision function f at time t_i; $\beta_0, \beta_1, ..., \beta n$ are the estimated parameters in the decision function. e_i is the *residual* (or *error*) between offer y_i from the opponent and the offer calculated by the predicting agent with the estimated parameters $\beta_0, \beta_1, ..., \beta n$. Equation (1) is *nonlinear* if f is nonlinear with respect at least one parameter β_r. Both time and resource dependent tactics are nonlinear. Equation (1) represents a system of nonlinear models.

The problem here is how to fit the curve f to the opponent's previous offers. In other words, how to adjust the parameters $\beta_0, \beta_1, ..., \beta n$ to minimise the residuals $e_0, e_1, ..., e_n$ collectively? The matrix form of equation (1) is:

$$\mathbf{Y} = \mathbf{F}(\beta_0, \beta_1, ..., \beta_n) + \mathbf{e} = \mathbf{F}(\boldsymbol{\beta}) + \mathbf{e} \qquad (2)$$

where $\boldsymbol{\beta}$ is a vector of parameters. From equation (2), $\mathbf{e} = \mathbf{Y} - \mathbf{F}(\boldsymbol{\beta})$, the residual $\mathbf{e}$ is distance between offers $\mathbf{Y}$ from the opponent and the calculated offers $\mathbf{F}(\boldsymbol{\beta})$ by the predicting agent. To minimise the distance $|\mathbf{e}|_2$ is to minimise its square $|\mathbf{e}|_2^2 = \mathbf{e'}\,\mathbf{e} = \sum_{i=1}^{m} e_i^2$, which is often referred to as *sum of squared residuals* (*SSR*). To solve the minimum is to solve the nonlinear "normal" equations: $\mathbf{X'F}(\boldsymbol{\beta}) = \mathbf{X'Y}$ where $X = \frac{\partial F}{\partial \beta}$ is the Jacobian matrix. The solving for $\boldsymbol{\beta}$ is an iterative process: 1) Choose a starting value $\boldsymbol{\beta}_0$ for $\boldsymbol{\beta}$; 2) Compute a $\boldsymbol{\Delta}$ such that $SSR(\boldsymbol{\beta}_0 + \boldsymbol{\Delta}) < SSR(\boldsymbol{\beta}_0)$ to improve SSR; 3) If the convergence measures are satisfied, the iteration process stops, $\boldsymbol{\beta} \leftarrow \boldsymbol{\beta}_0 + \Delta$, and $\boldsymbol{\beta}$ is the estimated parameter vector; otherwise, go to step 2). Marquardt method is chosen to compute $\boldsymbol{\Delta}$ in step 2) because it does not require $\boldsymbol{\beta}_0$ the initial values of parameters to be very close to the opponent's actual value of $\boldsymbol{\beta}$. In Marquardt method, update Δ is calculated by $\Delta = (X'X + \lambda diag(X'X))^{-} X'e$

The convergence measures chosen include R with criterion level 1E-5 such that the iteration will stop when $R \leq$ 1E-5, SSR with criterion level 1E-13 and no change in estimated parameters in consecutive iterations [7].

Mapping of the range of a negotiation issue [*min, max*]: The range of negotiation issues varies as the negotiation environment changes. Since the time and resource tactics are linear with respect to parameters *min* and *max* that when [*min, max*] is linearly scaled to [*min'*, *max'*], the estimate of *min* and *max* will be scaled in the same degree. Hence, the experiments in this research can be designed with *min* and *max* in a specific range, say, [0,1] or [0, 100].

Although Marquardt Method does not require the initial guess $\boldsymbol{\beta}_0$ very close to actual $\boldsymbol{\beta}$, this method does not work either if $\boldsymbol{\beta}_0$ is far from $\boldsymbol{\beta}$. For this reason, "grid search" facility is provided to relieve the demanding requirement on the predicting agent which allows the predicting agent to provide multiple initial guess $\boldsymbol{\beta}_0$'s.

5. Heuristics about predicting the other agent's tactics

Since the estimates of $\boldsymbol{\beta}$ are not influenced by the role assignment of buyer or seller except for the direction of approaching their reservation values, the predicting agent is assumed to be seller, and the opponent is buyer hereafter unless specifically noted otherwise.

By fitting well to the opponent's previous offers, the estimates of $\boldsymbol{\beta}$ can predict the opponent's future offers.

The estimates of **β** are usually different over the same history of the opponent's offers since these estimates are obtained with different β_0's. However, for the predicting agent, any of these estimates can be the opponent's actual **β** since the curve determined by each estimate fits to the opponent's previous offers so well (with SSR<1E-13). But those estimated *max* (denoted *~max* where "~" means estimated) and estimated deadline $t_{\max}$ (denoted $\sim t_{max}$) close enough to *max* and $t_{\max}$ are very important in avoiding negotiation breakdown when the opponent's offer is at *max* or t_{max} comes. However, *max* and $t_{\max}$ are different in the various estimates of **β**. A decision problem arises as to how to choose the estimates of **β**? The estimate of **β** with the smallest $\sim max^b$ or $\sim t_{\max}$ is chosen if the opponent is buyer, or the largest $\sim min^s$ is chosen if the opponent is seller. This choice can avoid risky $\sim max^b$ and $\sim t^b_{\max}$, but on the other hand, it makes the deal less ideal for the predicting agent.

To find out heuristics about the opponent's real *max*, t_{max}, and *type*, three groups of experiments have been designed. The heuristic knowledge obtained is about three typical tactics: Boulware, Conceder and resource-dependent tactics. Some of the settings are common to all the three groups of experiments: *min* ={10}, *max* ={20,50, 80,100}, *k*= {0.05,0.2,0.5,0.7,0.9}, t_{max} ={10,20,50,100}.

5.1. Heuristics about Boulware tactics

The Boulware tactics are determined by β ={0.02, 0.05,0.067,0.1,0.2,0.5}. There are 480 fully combinations in these settings, and each combination is a **β** which determines an individual tactic. As each offer is generated by a tactic at round from 12 (at least 6 offers) to t_{max}, the estimate of **β** with the smallest estimated *max* is recorded in tables in a relational database. When a pattern is found it is checked against the estimates made from round 12 to t_{max} in all the other tactics. The heuristics below indicates that the buyer's offer $offer(t_n)$ at current round t_n is very close to *max* and t_n is at t_{max} when 1) $\sim t_{max} \geq 10$ and $|offer(t_n) - \sim max(t_n)| < 0.2$ and $|t_n - \sim t_{max}| < 0.2$ and $offer(t_n)/\sim max(t_n) > 0.98$; or 2) $\sim t_{max} \geq 10$ and $0 \leq t_n - \sim t_{max} < 0.1$ and $|offer(t_n) - \sim max| < 10\% * offer(t_n)$; or 3) $\sim t_{max} < 10$ and $0 \leq t_n - \sim t_{max} < 0.1$; or 4) $\sim t_{max} < 10$ and $0 \leq t_n - \sim t_{max} < 0.9$ and $|offer(t_n) - \sim max(t_n)| < 10\% * offer(t_n)$; or 5) $\sim t_{max} < 10$ and $offer(t_n) * 90\% > \sim max(t_{n-1})$

There are only 4.5% of the 480 tactics for which the estimated *max* is not very close to *max* in terms of $|\sim max - max| > 0.5*\log(max - min)$.

5.2. Heuristics about Conceder tactics

There are 400 fully combinations or tactics in this group of experiments with β={5,10,15,20,40}.The experiments are conducted in the same procedures as in Section 5.1. Any of the following conditions can indicate the buyer is approaching to their reservation value *max.*

1) $|offer(t_n) - \sim max| < 0.5*\log(\sim max)$; or 2) $\sim\alpha > 0.90$ to 0.98; or 3) $|t_n - \sim t_{max}| < 1$, or 4) $|offer(t_n) - offer(t_{n-1})| < \sim max * 1\%$.

5.3. Heuristics about resource-dependent tactics

In this group of experiments, μ = {1,2,3,4,5,6,7,8,9,10} [3]. 800 tactics are tested with conclusion that if 1) $max - min \leq 30$ and SSR <3; or 2) $max - min > 30$ and SSR<0.1, then, $\sim max \approx max$ with $|\sim max - max| < \log((max - min)*0.7)$. For instance, if [*min, max*] = [10,100], *~max* will be in the area 100±1.8.

The type of resource tactics, indicated by μ, can be identified by the following empirical relation between ~ α and μ at the early stage of a negotiation in Table 1.

Table 1. Relation between μ and ~α

~ α	0.4~1.0	0.3~0.4	0.2~0.3	0~0.2
μ	1~5	6~7	8~9	10
Tactic type	Impatient or Steady	Patient	Very Patient	Very Patient

6. Application of heuristics to negotiation

With the nonlinear regress capability and obtained heuristics, a predicting agent’s extended reasoning mechanism is designed and its performance is examined in this section.

6.1. The reasoning mechanism of the predicting agent

The predicting agent and other agents follow similar protocol as described in Section 2.2.1. Before the point when a sufficient number of the opponent's offers are available to apply the learning approach, the predicting agent chooses a time-dependent tactic slightly tougher (β=0.5~0.8) than Linear tactic. After this point in a negotiation and no deal has been made so far, the predicting agent starts to estimate the opponent’s tactics.

First, predict which tactic family the opponent’s tactic belongs to. Since the opponent, in this research, is assumed to adopt a tactic in the forms described in Section 2.2 in this research, the polynomial, exponential and resource forms of function are simultaneously to be predicted by the nonlinear regression method. Experiments show that the forms other than the opponent’s actual tactic form will be estimated with a larger SSR (>5 when $[min, max] \subset [0,100]$). This finding will identify the opponent’s polynomial, exponential and resource forms. If it is not in these three forms, by assumption, the opponent’s tactic is from Tit-For-Tat family.

Then, the prediction is focused on the tactic parameters if the opponent's tactic form is identified to be time or resource decision function. With each new offer available from the opponent, newly obtained estimate of the opponent's parameters are used to predict the opponent's offer or the heuristics is applied to find the opponent's actual *max* and t_{max}, and an appropriate response ensues.

6.2. Performance of the predicting mechanism

To test the performance of this learning approach and compare with tactics as suggested by other researches, experimental negotiations have been designed with similar settings as in [3]:

The settings of tactic types: The three families of tactics are classified into 15 tactic types: CONceder: β ={20,30,40}; LINear: β={1}; BoulWare (BW): β = {0.025, 0.1, 0.2}; IMpatient: μ ={1}, STeady: μ ={2,3,5}, PAtient: μ ={6,8,10}, AvgTFTIM: δ=1 (as in resource tactic, Impatient is denoted by IM); AvgTFTST: δ={2,3,5}; AvgTFTPA: δ={6,8,10}. RndTFT and RelTFT are classified in the same way: RelTFTIM, RelTFTST, RndTFTPA, RndTFTIM, RndTFTST. RndTFT's additional parameter M ={1,3}. The default β for any TFT is 2. The buyer and the seller each has the 15 types of tactics and there are 225 (=15x15) type-pairings between the buyer and the seller.

The buyer's min^b={10}, max^b={20,50,100}, k^b={0.1, 0.9}, $t^b_{\max}$ ={30,45,60}. Φ is the overlapping degree between [min^b, max^b] and [min^s, max^s]. As in [3], max^b-min^b = max^s-min^s, and $k^s = k^b$. $min^s = min^b+\Phi(max^b-min^b)$, $max^s=min^s+(max^b-min^b)$. Since the buyer's max^b and seller's min^s are private information, the fully overlapping is not a common case, therefore in addition to Φ = 0(fully overlapping) in [3], two other overlapping degrees are included: Φ = 0.67 (1/3 partial overlapping), Φ = 1.33 (non-overlapping), which make the performance of tactics more comprehensive.

A *performance function* is defined as the non-subjective, cost-adjusted function in [5] to measure the success of an agent *a* with negotiation thread $X_a^{t_n}$:

$$f(X_a^{t_n}) = \begin{cases} V^a(x_a^{t_{deal}}) - V^a(x_a^N) - \tau(X_a^{t_n}) & \text{if last of } X_a^{t_n} = accept \\ -V^a(x_a^N) - \tau(X_a^{t_n}) & otherwise \end{cases}$$

where the larger the value of the function *f*, the more successful the agent's tactic. The Nash equilibrium point X_a^N is included for fairness, and communication cost is modelled as $\tau(X_a^{t_n})$ = tanh($comm\_k \times |X_a^{t_n}|$) where *comm_k* is a cost constant [3]. *comm_k* is chosen as {0.002(low), 0.01(middle), 0.05(high), 0.1(very high)}. For each of the 225 type-parings, there will be many negotiations to be held in order to observe the performance of each tactic type. With the above experimental settings, there will be 11664 negotiations between a BW buyer and a ST seller. To eliminate the influence of negotiation initiators on the performance of a tactic, both buyer and seller will initiate a negotiation once with the same setting. Another 6 groups of negotiations are held between the predicting mechanism and the 6 tactic types: CON, LIN, BW, IM, ST and PA. The experiment results are shown in Figure 2.

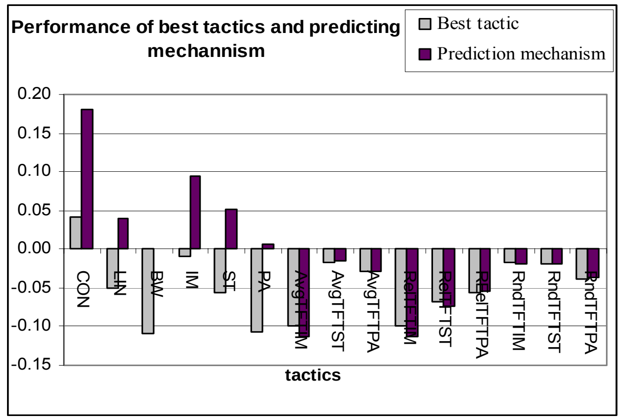

Figure 2. The relative success of predicting mechanism against other best tactics with non-high cost constant

If the cost constant *comm_k* is not high ($comm\_k \in$ [0.002, 0.01]), the cost is 0.06 with *comm_k* =0.002 and 0.3 with *comm_k* =0.01, at round 30. In Figure 2, against each of the 15 tactic types which are printed across the X-axis, the performance of the averagely best tactic type among the 15 tactic types and the performance of the predicting mechanism are printed as two adjacent bars. Figures 2 shows that the predicting mechanism do much better than the best of all the 15 tactic types against CON, LIN, BW, IM, ST and PA tactics.

Since Tit-For-Tat tactic bases on the opponent's offers, it is difficult to estimate the parameter δ, λ and M in a Tit-For-Tat tactic in which case, an averagely best tactic among the 15 tactic types against the whole Tit-For-Tat family is chosen for the predicting agent. Therefore, as Figure 2 shows, the predicting mechanism's performance is almost the same as that of the best from the 15 tactic types in the negotiation with a Tit-For-Tat tactic agent.

7. Discussion

The success of the predicting mechanism comes from its capability to predict the opponent's offer, and the heuristic knowledge in finding the opponent's reservation value and deadline. When the opponent's offer at the predicting agent's deadline is no better than the predicting agent' reservation value, withdraw from the negotiation will save cost since the breakdown is doomed. If the utility gained between the opponent's offer at predicting agent's deadline and the opponent's current offer is less than the cost involved as shown below:

$V^s(\sim x_b^{t_{\max}^s}) - V(x_b^{t_n}) \le \tanh(t_{\max}^s * comm\_k) - \tanh(t_n * comm\_k)$ (The predicting agent is assumed to be seller here for explanation purpose), accepting the current offer of the opponent is economically rational. The predicting agent can also wait for, or be aware of, the opponent's deadline when it comes, then accepting the offer at the opponent's deadline will reach a deal at the opponent's reservation, and yet avoid breakdown of the negotiation.

If the cost constant is high ($comm\_k \in [0.05, 0.1]$), for example, the cost is 0.905 with $comm\_k$=0.05, and 0.995 with $comm\_k$ =0.1, at round 30, whilst 30 is the smallest deadline in the performance test settings in Section 6.2. The cost 0.905 and 0.995 is very high when the maximum score (or utility) an agent can get from an offer is 1.0 Therefore any delay in reaching a deal will be punished with a high cost. Since the learning mechanism will be started until enough number of offers available and learning itself will take time, the gain from the learning process will be much smaller than the cost incurred. The information about the communication cost constant can be obtained from domain knowledge. In the case of high cost constant, the predicting agent can choose those nice tactics such Conceder, Linear, Steady or Tit-For-Tat with δ=1 to make deals in earlier rounds.

The original definition of performance function in [5] is not adopted in the performance tests presented above. The range of the cost definition $\tau(X_a^{t_n}) = comm\_k \times |X_a^{t_n}|$ in [5] is not within [0,1] and the cost with $comm\_k$=0.1 in 15 rounds, for example, will be 1.5, which is very high relative to the maximum utility 1.0. Our experiments with settings similar to [5] do not show their finding. Our findings show that in those conditions, Conceder rather than resource tactics is the overall best tactic since the cost is very high relative to utility.

8. Conclusion

Decision functions have been proposed as negotiation tactics to partly overcome the limitations of game theory. These tactics produce offers based on time remaining, resource remaining or the opponent's behaviour. When applying the findings of the existing research to negotiation, they require information about the other agents, either the probability distribution over the opponent's reservation value and deadline, or the opponent's tactic type. Knowing these kinds of information about the opponent will increase an agent's performance in negotiation. However, these kinds of information are private, particularly in competitive negotiation, where revealing own preference will render themselves in disadvantage position. To handle the uncertainty about other agents, this paper presents an online learning approach to eliciting information about other agents with only the opponent's previous offers.

Although the nonlinear regression can directly predict offers and identify tactic type of the opponent, the close estimates of the opponent's reservation value and deadline are not obvious. A large number of experiments were designed to obtain the heuristics about estimates of the other agent's reservation value and deadline. The performance tests show that the nonlinear regression prediction approach, working with the obtained heuristics, can indeed make the predicting agent get better deals than those tactics recommended by existing research. By balancing the future gain and cost, it can also avoid wasting time on unrewarding negotiation. This prediction approach can even get deals at a Boulware opponent's reservation value whilst avoiding the breakdown of the negotiation.

9. Future work

The prediction approach is being considered to apply to an opponent's fixed weighted combination of tactics and changing weighted combination of tactics. When adopting a combination of tactics, an agent's behaviour is perhaps not as characteristic as in the single tactic case and therefore may be more difficult to predict. However, by varying the weights on a set of orthogonal typical tactics, an agent's combination of tactics can be imitated.

Although the prediction approach is discussed around bilateral negotiation in this paper, the possibility of incorporating it into other models of negotiations will be explored in the future.

References

[1] Binmore, K. 1992. *Fun and Games: A Text on Game Theory*. Lexington, Massachusetts: D.C. Heath and Company.

[2] Faratin, P., Sierra, C. and Jennings, N.R. (1998). "Negotiation Decision Functions for Autonomous Agents." *Int. Journal of Robotics and Autonomous Systems* 24 (3-4): 159-182.

[3] Faratin, P. (2000). *Automated Service Negotiation Between Autonomous Compositional Agents*. PhD thesis, Queen Mary & Westfield College, University of London, London, UK, 2000.

[4] Fatima, S. S., Wooldridge, M., et al. (2001). Optimal Negotiation, Strategies for Agents with Incomplete Information. *Intelligent Agents VIII*. J.-J. Meyer and M. Tambe, Springer-Verlag. Lecture Notes in AI.

[5] Matos, N., Sierra, C. and Jennings, N.R. (1998). Determining successful negotiation strategies: an evolutionary approach. In *Proceedings of 3rd Int. Conf. on Multi-Agent Systems (ICMAS-98)*. Paris, France.

[6] Pruit,D.G. (1981). *Negotiation Behaviour*, Academic Press, Inc. London, 1981.

[7] *SAS/STAT User's Guide, Version 8* (3 Volume Set), SAS Publishing, August 2000, ISBN: 1580254942