

Modelling Agents Behaviour in Automated Negotiation

Chongming Hou

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

Abstract. *This paper presents a learning mechanism that applies nonlinear regression analysis to model a negotiation agent's behaviour based only on 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. By applying the nonlinear regression and the obtained heuristic knowledge, an agent can improve their overall performance by predicting the opponent's deadline and reservation value, terminating pointless negotiation, and avoiding negotiation breakdown. The findings of this study show that this approach can be used to obtain better deals than previously proposed 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. This 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 modelling 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. Due to the unrealistic assumptions in game theory, decision functions were proposed that would enable an agent to generate offers according to the time available, the resources remaining, 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, etc., are common knowledge known to each agent. Perfect computational rationality assumes a negotiation agent has unbounded computational power. With this power, agents can actually find the optimal strategies at the beginning of the game. 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, it is necessary for an agent to produce offers based on their own criteria, such as, time limits or resource availability. Using this approach, Faratin et al. proposed a negotiation model and three families of negotiation tactics, namely: time-dependent, resource-dependent and behaviour-dependent tactics [2].

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 the negotiation of a single issue that has a 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 $x_{s \rightarrow b}^t$ is the offer proposed by the seller s to the buyer b for a negotiation issue at time t . $x_{s \rightarrow b}^t$ is delimited by $[min^s, max^s]$, the range of all possible offers by s . The negotiation is to determine a value x ($x \in [min^s, max^s] \cap [min^b, max^b]$) which is mutually acceptable to s and b . max^b is actually the reservation value of b , that is, any value larger than max^b won't be accepted by b . min^s is the reservation value of s , that is, any value smaller than min^s won't be acceptable to s . Each agent a has a scoring (or utility) function $V^a : D^a \rightarrow [0,1]$ that assigns a score to value x in D^a . The scoring functions are either monotonically increasing for the seller or decreasing for the buyer.

A negotiation is an alternating-offer process terminated by *accept* or *withdraw*. Agent a 's response at t_n to agent b 's offer $x_{b \rightarrow a}^{t_{n-1}}$ sent at time t_{n-1} is defined as:

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

where $x_{a \rightarrow b}^{t_n}$ is the counter offer a offers to b when offer $x_{b \rightarrow a}^{t_{n-1}}$ is not accepted by a . t_{\max}^a is a 's deadline by which a must have completed the negotiation.

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 available, the resource remaining, or the opponent's behaviour.

2.2.2 Time-dependent family of tactics. In this family of tactics time is the predominant factor. All of the tactics in this family prescribe that an agent a concedes to their reservation value by their t_{\max}^a . What differentiates them is the shape of their concession curves (see Fig.1). The offer of agent a to agent b at time $t \leq t_{\max}^a$ is modelled by a function α^a which is dependent on time:

$$x_{a \rightarrow b}^t = \begin{cases} \min^a + \alpha^a(t)(\max^a - \min^a) & \text{if } V^a \text{ decreasing} \\ \min^a + (1 - \alpha^a(t))(\max^a - \min^a) & \text{if } V^a \text{ increasing} \end{cases}$$

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

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

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

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

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

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

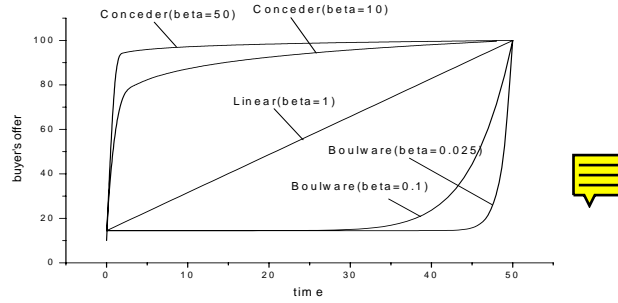


Fig. 1. Examples of Concession Curves for the Polynomial Time-dependent Family of Tactics

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 $resource^a(t) = \frac{\mu^a}{|X_{a \leftrightarrow b}^t|}$, and μ^a is the time agent a considers reasonable to spend on negotiation, and $|X_{a \leftrightarrow b}^t|$ 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.

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

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

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

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

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

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

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

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 Fatima et al. concerns the effect of time on the negotiation outcome between agents adopting the time-dependent family of tactics [4]. In their approach 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, Fatima et al. 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 of Matos et al. examines 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 [5]. 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 on average, such as resource tactic [5] or Linear time-dependent tactic [3], both tactics can not guarantee an agent the 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? Furthermore, if an agent can predict these kinds of information, can the information 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 tactic).
- predict 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).
- estimate the opponent's reservation value and deadline using the identified tactic type.

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 breakdown 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 accepting the buyer's current offer.

By adopting an effective tactic, 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

This section will discuss how to use nonlinear regression to model the opponent's tactic and to estimate parameters in the tactic. Among the three families of tactics in Section 2.2, the time-dependent and resource-dependent tactics are decision functions which can be represented respectively as

$$x_{a \rightarrow b}^{t_i} = f^{\text{time-dependent}}(\min, \max, k, t_{\max}, \beta; t_i) + e_i$$

$$x_{a \rightarrow b}^{t_i} = f^{\text{resource-dependent}}(\min, \max, k, t_{\max}, \mu; t_i) + e_i$$

In general, these two forms can be represented by

$$y_i = f(\theta_0, \theta_1, \dots, \theta_n; t_i) + e_i \quad (1)$$

where y_i is the previous offer ($x_{a \rightarrow b}^{t_i}$) generated by a decision function f at time t_i ($i=0, 1, \dots, m$); $\theta_0, \theta_1, \dots, \theta_n$ are the parameters ($\min, \max, k, t_{\max}, \beta/\mu$) 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 $\sim\theta_0, \sim\theta_1, \dots, \sim\theta_n$, the estimate of parameters $\theta_0, \theta_1, \dots, \theta_n$, respectively. Hereafter, " \sim " means "estimate of". Equation (1) is *nonlinear* if f is nonlinear with respect to at least one parameter θ_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, y_i ($i=0,1,\dots,m$). In other words, how to adjust the parameters $\theta_0, \theta_1, \dots, \theta_n$ to minimise the residuals e_0, e_1, \dots, e_n collectively? The matrix form of equation (1) is:

$$\mathbf{Y} = \mathbf{F}(\theta_0, \theta_1, \dots, \theta_n) + \mathbf{e} = \mathbf{F}(\boldsymbol{\theta}) + \mathbf{e} \quad (2)$$

where $\boldsymbol{\theta}$ is a vector of parameters, i.e., $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_n)$, $f_i(\boldsymbol{\theta}) = f(\boldsymbol{\theta}, t_i)$. From equation (2), $\mathbf{e} = \mathbf{Y} - \mathbf{F}(\boldsymbol{\theta})$, the residual \mathbf{e} is the distance between offers \mathbf{Y} from the opponent and the calculated offers $\mathbf{F}(\boldsymbol{\theta})$ 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 the *sum of squared residuals (SSR)*. To solve the minimum of *SSR* is to solve the nonlinear “normal” equations:

$$\mathbf{X}'\mathbf{F}(\boldsymbol{\theta}) = \mathbf{X}'\mathbf{Y} \quad (3)$$

where $\mathbf{X} = \frac{\partial \mathbf{F}}{\partial \boldsymbol{\theta}}$ is the Jacobian matrix. The solving for $\boldsymbol{\theta}$ is an iterative process:

1. choose a starting value $\boldsymbol{\theta}_0$ for $\boldsymbol{\theta}$;
2. compute a Δ such that $SSR(\boldsymbol{\theta}_0 + \Delta) < SSR(\boldsymbol{\theta}_0)$ to improve *SSR*;
3. if the convergence measures are satisfied, the iteration process stops, $\sim \boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_0 + \Delta$, and $\sim \boldsymbol{\theta}$ is returned as the estimate of parameter vector $\boldsymbol{\theta}$; otherwise, let $\boldsymbol{\theta}_0 \leftarrow \boldsymbol{\theta}_0 + \Delta$ and go to step 2.

There are generally four methods to compute Δ : Steepest descent, Gauss-Newton, Newton, and Marquardt. Marquardt method is chosen in this research to compute Δ in step 2 because it does not require the initial value, $\boldsymbol{\theta}_0$ of parameters to be very close to the opponent's actual value of $\boldsymbol{\theta}$. In Marquardt method, the update Δ is calculated by

$$\Delta = (\mathbf{X}'\mathbf{X} + \lambda \text{diag}(\mathbf{X}'\mathbf{X}))^{-1} \mathbf{X}'\mathbf{e}$$

The convergence measures chosen include the relative offset R measure with criterion level $1\text{E-}5$ (the iteration process will stop when $R \leq 1\text{E-}5$), *SSR* measure with criterion level $1\text{E-}13$ and the measure of no change in estimated parameters in consecutive iterations [7].

Although Marquardt Method does not require the initial guess $\boldsymbol{\theta}_0$ to be very close to actual $\boldsymbol{\theta}$, this method does not work either if $\boldsymbol{\theta}_0$ is far from $\boldsymbol{\theta}$. For this reason, “grid search” facility is designed that allows the predicting agent to provide multiple initial guesses for $\boldsymbol{\theta}_0$.

Mapping of a negotiation issue range [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 so that when [min, max] is linearly scaled to [min', max'], the estimate of min and max will be scaled in the same degree to the estimate of min' and max' . Hence, the experiments in this research can be designed with min and max in a specific range, say, $[0, 1]$ or $[0, 100]$.

5 Heuristics about Predicting the Other Agent's Tactic

The estimate of the opponent's parameter θ can be used to predict the opponent's offer. Furthermore, the estimate of t_{max} in $\sim\theta$ is expected to be used to avoid breakdown of the negotiation, and the estimate of max^b (min^s) in $\sim\theta$ to be used to make a deal at opponent's reservation value. However, the estimates of θ are usually different from each other since they are obtained from the various initial values for θ_0 . Consequently, different $\sim t_{max}$ and different $\sim max$ are obtained. Which $\sim t_{max}$ is the opponent's actual deadline t_{max} , and which $\sim max^b$ ($\sim min^s$) is the opponent's actual reservation value max^b (min^s)?

From the predicting agent's point of view, any $\sim\theta$ of these estimates can be the opponent's actual parameter θ , since the curve determined by each $\sim\theta$ fits to the opponent's previous offers very well (with $SSR < 1E-13$). A decision problem arises as to how to choose the estimates of θ ? One alternative is to choose the $\sim\theta$ with the smallest $\sim max^b$ or the smallest $\sim t_{max}$ if the opponent is a buyer, or to choose the largest $\sim min^s$ if the opponent is a seller. The chosen $\sim max^b$ ($\sim min^s$) and $\sim t_{max}$ are not necessarily the opponent's actual max^b (min^s) and t_{max} . This choice can avoid risky $\sim max^b$ ($\sim min^s$) and $\sim t_{max}^b$, but on the other hand, it makes the deal less ideal for the predicting agent.

In order to find some kinds of indication or heuristics about the opponent's actual max^b (min^s), t_{max} , and *type*, three groups of experiments were carried out. Heuristic knowledge is acquired about the typical tactics: Boulware time-dependent, Conceder time-dependent and resource-dependent tactics. Some settings are common to all the three groups of experiments: $min = \{10\}$, $max = \{20, 50, 80, 100\}$, $k = \{0.05, 0.2, 0.5, 0.7, 0.9\}$, $t_{max} = \{10, 20, 50, 100\}$. Since the estimates of θ are not influenced by the role assignment of the buyer or the seller except for the direction of approaching their reservation values, assume, hereafter, that the predicting agent is the seller, and the opponent is the buyer unless specifically noted otherwise.

5.1 Heuristics about Boulware Time-dependent Tactics

The Boulware tactics are determined by $\beta = \{0.02, 0.05, 0.067, 0.1, 0.2, 0.5\}$. There are 480 combinations in these settings, and each combination is a θ which determines an individual tactic. As each offer is generated by a tactic, the $\sim\theta$ with the smallest estimated max is recorded. When a pattern is found, it is checked against all the $\sim\theta$ obtained in the 480 tactics. The heuristics below indicates that the buyer's offer $offer(t_n)$ at t_n is very close to their reservation value max and t_n is at their deadline 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 $offer(t_n) * 90\% > \sim max(t_{n-1})$

With 95.5% of the 480 tactics the estimated $\sim max$ is very close to max in terms of $|\sim max - max| \leq 0.5 * \ln(max - min)$. For example, if $min = 10$, $max = 50$, the estimate of $\sim max$: $\sim max \approx max \pm 1.84$, i.e., $\sim max \approx 50 \pm 1.84$.

5.2 Heuristics about Conceder Time-dependent Tactics

There are 400 combinations or tactics in this group of experiments with $\beta = \{5, 10, 15, 20, 40\}$. The experiments are conducted in the same procedures as in Section 5.1. Any of the following conditions can indicate the buyer's offer $offer(t_n)$ is very close to their reservation value max .

1. $|offer(t_n) - \sim max| < 0.5 * \ln(\sim max)$; 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, $\mu = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$. 800 tactics are tested with the finding that if

1. $\sim max - \sim min \leq 30$ and $SSR < 3$; or
2. $\sim max - \sim min > 30$ and $SSR < 0.1$

then, $\sim max \approx max$ with the precision of $|\sim max - max| < 0.5 * \ln(max - min)$.

For instance, if $min = 10$, $max = 100$, $\sim max$ will be in the area 100 ± 2.25 .

The empirical relation between $\sim \alpha$ (see Section 2.2.3) and μ is obtained and summarised in Table 1. By this relation, the opponent's resource tactic type can be determined as Impatient ($\mu=1$), Steady ($\mu \in [2,5]$), or Patient ($\mu \in [6,10]$) [3].

Table 1. Relation between the resource tactic type (μ) and $\sim \alpha$

| | | | | |
|---------------|---------------------|---------|--------------|--------------|
| | Impatient or Steady | Patient | Very Patient | Very Patient |
| μ | 1~5 | 6~7 | 8~9 | 10 |
| $\sim \alpha$ | 0.4~1.0 | 0.3~0.4 | 0.2~0.3 | 0~0.2 |

6 Application of Heuristics to Negotiation

With the nonlinear regression ability and heuristics in the previous sections, a predicting agent's extended reasoning mechanism and its performance are presented in the subsections below.

6.1 The Reasoning Mechanism of the Predicting Agent

The predicting agent and the opponent follow the protocol 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 ($\beta=0.5\sim 0.8$) than the Linear tactic. The predicting mechanism is started to estimate the opponent's tactics after this point in a negotiation provided that no deal has been made.

First, the agent predicts which tactic family the opponent's tactic belongs to. Since the opponent, in this research, is assumed to adopt a tactic of the form in Section 2.2, the polynomial, exponential and resource forms of function are simultaneously predicted by the nonlinear regression method. Experiments show that forms other than the opponent's actual tactic will be estimated with a larger *SSR* (>5 when $[min, max] \subset [0,100]$). Therefore, it will enable the predicting agent to identify the opponent's polynomial, exponential or resource tactic form. If it is not of these three forms, by assumption above, the opponent's tactic is from the behaviour-dependent family.

Once the opponent's tactic form is determined, the prediction then focuses on the parameters of the opponent's time-dependent or resource-dependent tactic. With each offer received from the opponent, the estimate of the opponent's parameters is generated and used to predict the opponent's offer. The heuristics is applied to the newly generated estimate in order to find the opponent's actual *max* and t_{max} . Based on the prediction and the information about *max* and t_{max} , an appropriate response is sent to the opponent.

This extended reasoning mechanism with prediction approach is referred to as prediction mechanism (Pred).

6.2 Experimental Negotiations

To test the performance of this learning approach, experimental negotiations have been designed with similar settings as in [3]:

The settings of the normal, non-prediction tactic types: The three families of tactics are classified into nine tactic types as below:

- Time-dependent:
Conceder (CON): $\beta = \{20, 30, 40\}$, Linear (LIN): $\beta = \{1\}$, Boulware (BW): $\beta = \{0.025, 0.1, 0.2\}$
- Resource-dependent:
Impatient (IM): $\mu = \{1\}$, Steady (ST): $\mu = \{2, 3, 5\}$, Patient (PA): $\mu = \{6, 8, 10\}$
- Behaviour-dependent:
Average Tit-For-Tat (AvgTFT): $\gamma = \{1, 2, 3, 5, 6, 8, 10\}$,
Relative Tit-For-Tat (RelTFT): $\delta = \{1, 2, 3, 5, 6, 8, 10\}$,
Random Tit-For-Tat (RndTFT): $\delta = \{1, 2, 3, 5, 6, 8, 10\}$, $M = \{1, 3\}$
Default behaviour of TFT: when the number of the other agent's previous offers is less than δ (γ), a time-dependent tactic with $\beta=2$ is adopted for the behaviour-dependent tactics.

The buyer's $min^b = \{10\}$, $max^b = \{20, 50, 100\}$, $k^b = \{0.1, 0.9\}$, $t_{max}^b = \{30, 45, 60\}$. Φ is the overlapping degree between $[min^b, max^b]$ and the seller's range $[min^s, max^s]$. 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 $\Phi = 0$ (fully overlapping) in [3], two other overlapping degrees are included: $\Phi = 0.67$ (1/3 partial overlapping), $\Phi = 1.33$ (non-overlapping), which make the performance of tactics more comprehensive.

To eliminate the influence of the role of the buyer or the seller in the performance of a tactic, both the buyer and the seller will initiate a negotiation once with the same setting. Therefore, it can be assumed that the opponent is a buyer, who can adopt any of the nine non-prediction tactic types, and that the seller can have the nine non-prediction tactic types and the predicting mechanism (Pred).

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

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

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)}.

Can the seller's predicting mechanism outperform the seller's other nine tactics against each of the buyer's nine tactic types?

To answer this question, the negotiations were carried out between the seller's 10 (9 non-prediction tactic types and the Pred) tactic types and the buyer's 9 non-prediction tactic types. Hence, 90 (= 10 x 9) groups of negotiations were carried out between the buyer and the seller, and each of the 90 groups of negotiations includes a large number of negotiations. For example, there are 11664 negotiations between a BW buyer and a ST seller.

6.3 Performance of the Predicting Mechanism

The results of all the negotiations in the experiments demonstrate that with reasonable communication cost, the seller's predicting mechanism is

1. as good as the seller's best tactic among the nine non-prediction tactic types if the buyer adopts a Tit-For-Tat tactic in {AvgTFT, RelTFT, RndTFT};
2. much better than the best tactic among the nine non-prediction tactic types if the buyer adopts a tactic in {CON, LIN, BW, IM, ST, PA}.

Fig. 2 shows the seller's performance with predicting mechanism (Pred) and with the 9 non-prediction tactic types in negotiation with the buyer who adopts a tactic in {AvgTFT, RelTFT, RndTFT}. For each of the buyer's tactic type, only the performance of the seller's best non-prediction tactic types is shown against the performance of the prediction mechanism, since the comparison between the prediction mechanism and the best non-prediction tactic types is sufficient to show how well the prediction mechanism performs. For example, in negotiation with the buyer adopting an AvgTFT tactic, the seller's best tactic in performance on average is ST. Fig.2 shows, across all the buyer's tactics in {AvgTFT, RelTFT, RndTFT}, the seller's prediction mechanism is as good as the seller's best non-prediction tactics.

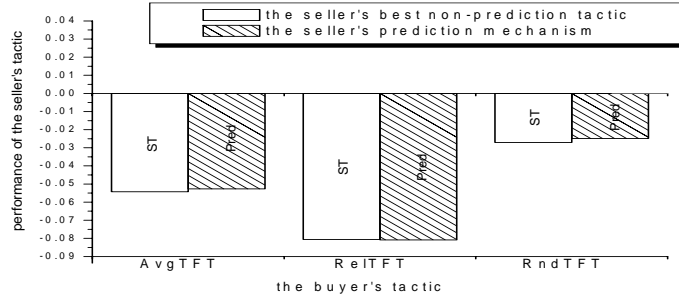


Fig. 2. The Relative Success of the Seller's Predicting Mechanism to the Seller's Non-prediction Tactics in Negotiations with the Buyer's Tit-For-Tat Tactics

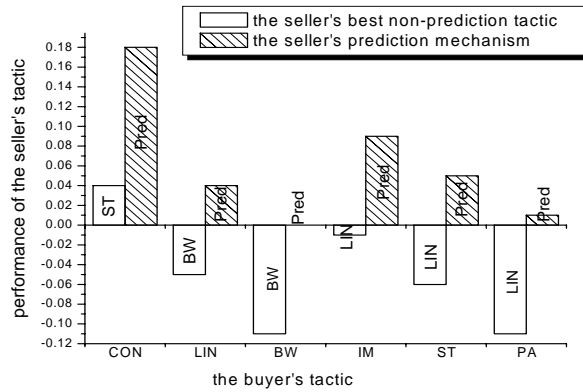


Fig. 3. The Relative Success of the Seller's Predicting Mechanism to the Seller's Non-prediction Tactics in Negotiations with the Buyer's 6 Typical Tactics

Fig. 3 shows that the seller's predicting mechanism does much better than the seller's best tactic among the 9 non-prediction tactic types against each of the buyer's tactic in {CON, LIN, BW, IM, ST, PA}. In Fig. 3, for example, if the buyer adopts a ST tactic, the best non-prediction tactic type available to the seller is LIN with which the seller scores -0.06. However, the seller can get a score of 0.05 by adopting the prediction mechanism (Pred). The increase, 0.11, in the seller's score will be a gain of £1,100 for the seller in the deal, for instance, if the seller's price range is [£10,000, £20,000].



7 Discussion

The success of the predicting mechanism comes from its ability 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's reservation value, withdrawing 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:

$$V^s(\sim x_b^{t_{\max}^s}) - V(x_b^{t_n}) \leq \tanh(t_{\max}^s * comm_k) - \tanh(t_n * comm_k),$$

accepting the current offer of the opponent is economically rational (The predicting agent is assumed to be seller here for explanation purpose). 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.

The performance in Fig.2 and Fig.3 are drawn with the cost constant $comm_k$ in $[0.002, 0.01]$ in which the cost is reasonable. The cost is 0.06 with $comm_k=0.002$ at round 30, and 0.3 with $comm_k=0.01$ at round 30, whilst 30 is the smallest deadline in the performance test setting $t_{\max} = [30, 60]$ in Section 6.2.

If the cost constant is high ($comm_k \in [0.05, 0.1]$), the cost is 0.905 with $comm_k=0.05$ at round 30, and 0.995 with $comm_k=0.1$ at round 30. The costs 0.905 and 0.995 are very high when the maximum score (or utility) an agent can get from a deal is 1.0. Therefore, any delay in reaching a deal will be punished with a high cost. Since the learning mechanism will start until an 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. Nevertheless, 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 as Conceder, Linear, Steady or Tit-For-Tat (with $\delta=1$) to make deals with best performance.

The choice of cost definition will also significantly influence the experiment results. The original definition of the performance function in [5] is not adopted in performance tests presented above because of its cost definition: $\tau(X_a^{t_n}) = comm_k * |X_a^{t_n}|$. The range of this cost definition is not within $[0, 1]$, and the cost with $comm_k=0.1$ at round 15, for example, will be 1.5, which is very high relative to the maximum utility 1.0. The result of an experiment, with settings similar to [5], shows that Conceder, Impatient and Tit-For-Tat ($\delta=1$) are more successful than the resource tactics recommended by [5].

8 Conclusion

Decision functions have been proposed as negotiation tactics to overcome the limitations of game theory. These tactics produce offers based on the time available, the resource remaining or the opponent's behaviour. When applying the findings of the existing research to negotiation, they require information about the other agent, either the probability distribution over the opponent's reservation value and deadline, or the opponent's tactic type. Knowing the 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 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.

The nonlinear regression can directly predict offers and identify tactic type of the opponent. A large number of experiments were designed to obtain the heuristics about the 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 two-party negotiation in this paper, the possibility of incorporating it into other models of negotiations will be explored in the future.

References

1. Binmore, K.: *Fun and Games: A Text on Game Theory*. Lexington, Massachusetts: D.C. Heath and Company (1992)
2. Faratin, P., Sierra, C. and Jennings, N.R.: Negotiation Decision Functions for Autonomous Agents. *Int. J. Robotics and Autonomous Systems* 24 (1998) (3-4) 159-182
3. Faratin, P.: *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. and Jennings, N.R.: Optimal Negotiation, Strategies for Agents with Incomplete Information. In: Meyer, J.-J., Tambe, M. (eds.): *Intelligent Agents VIII: Agent Theories, Architectures, and Languages*. Lecture Notes in AI, Vol. 2333. Springer-Verlag, Berlin Heidelberg New York (2001)
5. Matos, N., Sierra, C. and Jennings, N.R.: Determining successful negotiation strategies: an evolutionary approach. In: *Proceedings of 3rd Int. Conf. on Multi-Agent Systems (ICMAS-98)*. Paris, France (1998)
6. Pruitt, D.G.: *Negotiation Behaviour*, Academic Press, Inc. London, 1981
7. SAS/STAT User's Guide, Version 8 (3 Volume Set), SAS Publishing, August 2000