

Determining Successful Negotiation Strategies: An Evolutionary Approach

Noyda Matos¹ Carles Sierra¹ Nick R. Jennings²

¹IIIA - Artificial Intelligence Research Institute
CSIC - Spanish Council for Scientific Research
Campus UAB, 08193 Bellaterra, Catalonia, Spain.

² Dept. Electronic Engineering
Queen Mary and Westfield College.
University of London, London E1 4NS, UK.

Abstract

To be successful in open, multi-agent environments, autonomous agents must be capable of adapting their negotiation strategies and tactics to their prevailing circumstances. To this end, we present an empirical study showing the relative success of different strategies against different types of opponent in different environments. In particular, we adopt an evolutionary approach in which strategies and tactics correspond to the genetic material in a genetic algorithm. We conduct a series of experiments to determine the most successful strategies and to see how and when these strategies evolve depending on the context and negotiation stance of the agent's opponent.

1. Introduction

Negotiation is a central component of many multi-agent systems. Agents negotiate to coordinate their activities and to come to mutually acceptable agreements about the division of labour and resources. In many cases, the agents involved are required to exhibit a range of different behaviours in a variety of different settings. Thus, an agent negotiating with a competitor needs to behave differently to one negotiating with a member of its own organisation, and an agent with limited time to reach an agreement needs to behave differently to one which has no time pressure. Moreover, to be usable in practical situations the agents need to exhibit such flexible behaviour in settings in which: individuals' preferences are private information, individuals are bound in terms of the amount of communication and computation they can perform, and the complete set of agreements cannot be a priori enumerated. Given such constraints, experience indicates that a heuristic approach to modelling negotiation is often the most appropriate [9].

In line with this view, we have developed a model of negotiation which relies on strategies and tactics to define an agent's behaviour [5, 13]. This model was devised for a real-world business process management system [8] in which the agents need to agree who should perform a particular service and under what terms and conditions. We term

such negotiation *service-oriented*. Since our model has to operate in a wide range of environments and has a large number of parameters, it was decided that the best means of determining its performance characteristics is through empirical evaluation. To this end, we have already reported initial results about which types of tactic are suitable for which types of environment [5]. Our aim in this paper, therefore, is to extend this analysis by considering how the agents' negotiation strategies evolve. The technique we adopt for evolving strategies is Genetic Algorithms (GAs) [6].

The long term goal of this research is to study the dynamics of agent populations – how the relative presence of strategies changes as populations co-evolve, whether populations have stable equilibrium points, and how changes in an agent, or group of agents, alter the equilibria. The work described herein represents a step towards these aims in two ways. Firstly, it performs a systematic evaluation of a range of negotiation strategies. It indicates, for the first time, the relative success of the various strategies and how these strategies evolve over time to better fit their environment. Secondly, it extends the range and sophistication of interactions studied in the field of artificial social systems [2] (see section 5 for more details).

The remainder of the paper is structured in the following manner. Section 2 summarises our service-oriented negotiation model. Section 3 describes the encoding and operation of the GA and how populations were evolved. Section 4 presents our experimental results and discusses their implications. Section 5 places our work in context. Section 6 outlines the avenues of further research.

2. The service-oriented negotiation model

Our multi-lateral negotiation model is based on a set of mutually influencing two parties, many issues negotiations [5, 13]. Such negotiations involve two roles that are, in principle, in conflict: *Sellers* of services and *Buyers* of services.

Here we limit ourselves to consider issues for which negotiation amounts to determining a value between a delimited range (i.e. $x_j \in \mathcal{D}_j = [\min_j, \max_j]$). Each agent has

a scoring function $V_j^a : \mathcal{D}_j^a \rightarrow [0, 1]$ that gives the score agent a assigns to the value x of issue j in the range of its acceptable values \mathcal{D}_j^a . For convenience, scores are kept in the interval $[0, 1]$ and scoring functions are either monotonically increasing or decreasing. When an agent receives an offer, it rates it using a function that combines the scores of the different issues in a linear fashion:

$$V^a(x) = \sum_{1 \leq j \leq n} w_j^a V_j^a(x_j)$$

where w_j^a is the importance of issue j for agent a . We assume the weights are normalised. If the score of the received offer is greater than that of the counter offer the agent would send at this point, then the offer is accepted. If the pre-established deadline (t_{max}^a) by when the negotiation must have completed is reached, the offer is rejected. Otherwise, a counter offer is submitted.

Let $x_{a \rightarrow b}^t$ represent the vector of values proposed by agent a to agent b at time t , and $x_{a \rightarrow b}^t[j]$ be the value for issue j . A *Negotiation Thread* between agents a and b , at time t_n , noted $X_{a \rightarrow b}^{t_n}$, is any finite sequence of length n of the form $(x_{a \rightarrow b}^{t_1}, x_{b \rightarrow a}^{t_2}, x_{a \rightarrow b}^{t_3}, \dots)$ with $t_1, t_2, \dots \leq t_n$, where:

1. $t_{i+1} > t_i$, the sequence is ordered over time,
2. For each issue j , $x_{a \rightarrow b}^{t_i}[j] \in \mathcal{D}_j^a$, $x_{b \rightarrow a}^{t_{i+1}}[j] \in \mathcal{D}_j^b$ with $i = 1, 3, 5, \dots$, and optionally the last element of the sequence being one of the particles $\{accept, reject\}$.

$last(X_{a \rightarrow b}^{t_n})$ is a function returning the last element in a sequence.

Offers and counter-offers are generated by linear combinations of functions called *tactics*. A tactic generates a value for a single negotiation issue based upon a single criterion: time remaining (section 2.1), resources remaining (section 2.2) or the opponent's behaviour (section 2.3). If multiple criteria are important in determining the value of a negotiation issue, then multiple tactics can be applied to the issue. In this case, the tactics are assigned a weight to indicate their relative importance. As the negotiation proceeds, new criteria may become relevant and the relative importance of existing criteria may vary. To reflect this fact, an agent has a *strategy* which varies the weights of the different tactics over time in response to various environmental and negotiation cues (section 2.4).

2.1. Time-dependent tactics

These tactics model the fact that the agent is likely to concede more rapidly as the negotiation deadline approaches. All the tactics in this family put forward their minimum reservation value at t_{max}^a , but what differentiates them is the shape of their concession curve. We model the offer of agent a to agent b for issue j at time $t \leq t_{max}^a$, by a

function α_j^a depending on time.

$$x_{a \rightarrow b}^t[j] = \begin{cases} \min_j^a + \alpha_j^a(t)(\max_j^a - \min_j^a) & \text{If } V_j^a \text{ is decreasing} \\ \min_j^a + (1 - \alpha_j^a(t))(\max_j^a - \min_j^a) & \text{If } V_j^a \text{ is increasing} \end{cases}$$

A wide range of functions can be defined simply by varying the way in which $\alpha_j^a(t)$ is computed. Here we use a family of polynomial functions parameterised by a value $\beta_j \in \mathbb{R}^+$ that determines the convexity of the curve:

$$\alpha_j^a(t) = \left(\frac{\min(t, t_{max}^a)}{t_{max}^a} \right)^{\frac{1}{\beta_j}}$$

This expression represents an infinite number of possible tactics, one for each value of β_j . However to better understand their behaviour, we classify them into two sets which show clearly different patterns of behaviour: **Boulware** (don't start conceding until the deadline is nearly up) with $\beta_j \ll 1$, and **Conceder** (start giving ground fairly quickly) with $\beta_j \gg 1$.

2.2. Resource-dependent 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:

$$\alpha_j^a(t) = e^{-\frac{\mu_j^a}{|x_{a \rightarrow b}^t|}}$$

Here we use a measure of resource which includes both the time μ_j^a agent a considers reasonable to spend negotiating for each issue j and the number of messages involved in the negotiation thread.

2.3. Behaviour-dependent tactics

These tactics base their actions on the behaviour of their negotiation opponent [1]. They differ in which aspect of the opponent's behaviour they imitate and to what degree. By default, if the duration of the negotiation does not permit a tactic to be applied (i.e. $t < 2\delta$), the agent employs a Conceder tactic with $\beta = 2$ (in line with Axelrod's recommendation [1]).

Relative Tit-For-Tat (Relative-TFT)

Reproduce, in percentage terms, the behaviour that its opponent exhibited $\delta_j \geq 1$ steps ago:

$$x_{a \rightarrow b}^{t_n+1}[j] = \min(\max(\frac{x_{b \rightarrow a}^{t_n-2\delta_j}[j]}{x_{b \rightarrow a}^{t_n-2\delta_j+2}[j]} x_{a \rightarrow b}^{t_n-1}[j], \min_j^a), \max_j^a)$$

Random Absolute Tit-For-Tat (Random-TFT)

Same as *Relative TFT*, except that the behaviour is imitated in absolute terms:

$$x_{a \rightarrow b}^{t_n+1}[j] = \min(\max(x_{a \rightarrow b}^{t_n-1}[j] + (x_{b \rightarrow a}^{t_n-2\delta_j}[j] - x_{b \rightarrow a}^{t_n-2\delta_j+2}[j]) + (-1)^{\delta_j} R(M_j), \min_j^a), \max_j^a)$$

where $s_j = 0(1)$ if V_j^a is decreasing (increasing) and $R(M_j)$ is a function that generates a random integer in the interval $[0, M_j]$.

Averaged Tit-For-Tat (Average-TFT).

Uses the average of the percentage change in a window of size $\lambda_j \geq 1$ of its opponent's history:

$$x_{a \rightarrow b}^{t_{n+1}}[j] = \min(\max(\frac{x_{b \rightarrow a}^{t_n-2\lambda_j}[j]}{x_{b \rightarrow a}^{t_n-1}[j]} x_{a \rightarrow b}^{t_n-1}[j], \min_j^a), \max_j^a)$$

2.4. Strategies

When agent a receives an offer from agent b , it becomes the last element in the current negotiation thread between the agents. If a deems the offer unsatisfactory, it generates a counter offer. In generating this response, a may use different weighted combinations of tactics for each negotiation issue. However for a given issue, the agent can employ at most one tactic from each of the aforementioned families. Thus, an agent can use at most six tactics to generate a counter offer for any one issue. Given a negotiation thread between agents a and b at time t_n , $X_{a \leftrightarrow b}^{t_n}$, over domain $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_p$, with $\text{last}(X_{a \leftrightarrow b}^{t_n}) = x_{b \rightarrow a}^{t_n}$, and a finite set of m tactics¹ $T^a = \{\tau_i | \tau_i : MS_a \rightarrow \mathcal{D}\}_{i \in [1, m]}$, where MS_a is the set of all possible mental states for agent a at time t (MS_a^t). A *weighted counter proposal*, $x_{a \rightarrow b}^{t_{n+1}}$, is a linear combination of the tactics given by a matrix of weights:

$$\Gamma_{a \rightarrow b}^{t_{n+1}} = \begin{pmatrix} \gamma_{11} & \gamma_{12} & \dots & \gamma_{1m} \\ \gamma_{21} & \gamma_{22} & \dots & \gamma_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{p1} & \gamma_{p2} & \dots & \gamma_{pm} \end{pmatrix}$$

defined in the following way:

$$x_{a \rightarrow b}^{t_{n+1}}[j] = (\Gamma_{a \rightarrow b}^{t_{n+1}} * T^a(MS_a^{t_{n+1}}))[i, j]$$

With $(T^a(MS_a^{t_{n+1}}))[i, j] = (\tau_i(MS_a^{t_{n+1}}))[j]$, $\gamma_{ji} \in [0, 1]$ and for all issues j , $\sum_{i=1}^m \gamma_{ji} = 1$. The weighted counter proposal extends the current negotiation thread by appending $x_{a \rightarrow b}^{t_{n+1}}$ onto $X_{a \leftrightarrow b}^{t_n}$.

Given a set of tactics, different types of negotiation behaviour can be obtained by weighting the tactics in different ways. That is, by changing the matrix Γ particular to each negotiation thread. Thus a *Negotiation Strategy* for agent a is any function f which, given a 's mental state at time t_n , $MS_a^{t_n}$, and a matrix of weights at time t_n , $\Gamma_{a \rightarrow b}^{t_n}$, generates a new matrix of weights for time t_{n+1} :

$$\Gamma_{a \rightarrow b}^{t_{n+1}} = f(\Gamma_{a \rightarrow b}^{t_n}, MS_a^{t_n})$$

¹This definition uses the natural extension of tactics to the multi-dimensional space of issues' values.

3. Co-evolving the negotiation strategies

GAs generate a sequence of ever improving (fitter) populations as the outcome of a search method modelled by a *selection mechanism*, *crossover* (recombining existing genetic material in new ways) and *mutation* (introducing new genetic material by random modifications) [6]. In our case, the individuals of the population are negotiating agents, and their genetic material is the parameters of the negotiation tactics and their relative importance weightings. Co-evolution refers to the fact that both the buyer and seller populations are simultaneously evolving and that the fitness of an individual in one population is based on direct competition with individuals from the other population [12].

To find strategies which are optimally adapted for particular negotiation situations, we need to indicate: how strategies are encoded as genes (3.1), how the fitness of the strategies is computed (3.2), and how the search algorithm operates in detail (3.3).

3.1. Encoding strategies as genes

Each agent is represented as a string of fixed length. The bits of the string (the gene) represent the parameters of the agent's negotiation strategy²:

- t_{max}^a : *real*. Maximum time the agent has to negotiate.
- Issue-related genes:
 - $\mathcal{D}_j^a = [\min_j^a, \max_j^a]$. Interval of acceptable values for issue j .
 - $Type.V_j^a : \{inc, dec\}$. Scoring function type: monotonically increasing or decreasing.
- Tactic-related genes.
 - *Time Dependent*. β_j : *real*. $\beta_{boulware_j} \ll 1$, $\beta_{conceder_j} \gg 1$.
 - *Resource Dependent*. μ_j^a : *integer*. Time considered reasonable to spend negotiating for issue j .
 - *Imitative Tactics*.
 - * Relative TFT. δ_j : *integer*. Number of steps into the past considered.
 - * Random TFT. δ_j : *integer*. Number of steps into the past considered. M_j : *integer*. Maximum amount by which agent can vary its imitative behaviour.
 - * Averaged TFT. λ_j : *integer*. Size of window over which average is computed.

²To ensure there is no interference between negotiation issues during the evaluation of the GA's results, different tactics are used for each issue. Thus if there are p issues and $q \in [1, 6]$ tactic families we define q concrete tactics (by fixing the parameters in the family expressions) per issue. So the total number of concrete tactics an agent employs in a particular thread is $m = p \times q$. This means there are $(p - 1)q$ zeros in each row of Γ .

5		[10, 100]	inc		[2,6]	dec		0.5	0.1	...	0.1	0.2	0.6	...	0.05		0.02	20
t_{max}^a		\mathcal{D}_1^a	Type-V ₁		\mathcal{D}_2^a	Type-V ₂		$\gamma_{1,1}$	$\gamma_{1,2}$...	$\gamma_{1,6}$	$\gamma_{2,7}$	$\gamma_{2,8}$...	$\gamma_{2,12}$		$\beta_{boulware_1}$	$\beta_{conceder_1}$
5	2	3	3	4	0.2	30	4	3	4	2	2							
$\mu_{resource_1}$	$\delta_{relative_1}$	δ_{random_1}	M_{random_1}	$\lambda_{average_1}$	$\beta_{boulware_2}$	$\beta_{conceder_2}$	$\mu_{resource_2}$	$\delta_{relative_2}$	δ_{random_2}	M_{random_2}	$\lambda_{average_2}$							

Figure 1. Encoding of a strategy which uses a combination of 6 tactic families for 2 issues.

- Strategy-related genes: Values of γ_{ij} for the q tactics used for each issue. The genetic code contains only the values for the non-zero gammas (i.e. those tactics not used to determine a value for a particular issue are not represented.)

By means of an illustration, figure 1 represents an encoding of an agent negotiating over two issues using six families of tactics. The first row indicates that the maximum time available for negotiation is 5, that issue 1 can take a value between 10 and 100 and that issue 2 can vary between 2 and 6. The next genes contain the weights of each tactic —i.e. for issue 1, the relative weight for the Boulware tactic is 0.5 and the remaining five tactics have weight 0.1 (meaning that for this issue the agent behaves in a generally boulwarish fashion). For issue 2, the weight for the Boulware tactic is 0.2, 0.6 for the Conceder, and 0.05 for the other four. The remaining genes contain the values of the parameters of the 12 tactics used —thus $\gamma_{1,1}$ has $\beta_{boulware_1} = 0.02$, $\gamma_{1,2}$ has $\beta_{conceder_1} = 20$, and so on.

3.2. Measuring a strategy's fitness

An agent's fitness value indicates how well it performs in comparison to others in the same population. Following basic evolutionary ideas, fitness also determines the agent's chance of surviving to the next population generation. The greater the fitness, the more likely the agent is to reproduce. To compute an agent's fitness we play a round-robin tournament in which each buyer negotiates with each seller. Then, we score each agent with a value that measures how well it performed. This score becomes its fitness.

Determining the score to attribute to an agent is a non-trivial task. The most obvious start point is the agent's utility function which provides its preference rating over deals. However, adopting this would mean that the fitness rating is subjective and varies between agents. Moreover, the utility function takes an end-result perspective meaning the resources consumed in reaching the deal are not considered. To counter the subjectivity problem, we define a fitness function that compares the utility associated to the deal with that associated to the Nash equilibrium point (the point at which the sellers' and the buyers' scoring functions are equal [3]). The more positive the difference, the more successful the agent's behaviour. The cost involved in attaining

a deal is taken into consideration by associating a charge with each message interchange —the more messages exchanged in coming to a deal, the higher the associated cost. With these two components in place, a non-subjective, cost-adjusted fitness function f can now be defined:

$$f(X_c^{t_n}) = \begin{cases} V^c(x_c^{t_{deal}}) - V^c(x_c^N) - \tau(X_c^{t_n}) & \text{if } \text{last}(X_c^{t_n}) = \text{accept} \\ -V^c(x_c^N) - \tau(X_c^{t_n}) & \text{otherwise} \end{cases}$$

where $x_c^{t_{deal}}$ is the final deal, and x_c^N represents the deal that would be made at the Nash equilibrium point. Communication cost is modelled as: $\tau(X_c^{t_n}) = k * |X_c^{t_n}|$ where k is a constant. Varying k allows the relative importance of communication cost to be changed in the fitness computation.

3.3. Algorithm steps

The details of the GA used are as follows (recall there are two families of agents: sellers S and buyers B):

Randomly create initial seller (P_S^0) and buyer (P_B^0) populations;

While not (Stopping Criterion) **do**

- Make a tournament and calculate the fitness of all the individuals in P_S^t and P_B^t ;
- $MP_S = \text{Tournament_Selection}(P_S^t)$;
- $MP_B = \text{Tournament_Selection}(P_B^t)$;
- $Best_{MP_S} = \text{Best_Individual}(MP_S)$;
- $Best_{MP_B} = \text{Best_Individual}(MP_B)$;
- $R_S = \text{Crossover\&Mutation}(MP_S - Best_{MP_S})$;
- $R_B = \text{Crossover\&Mutation}(MP_B - Best_{MP_B})$;
- $P_S^{t+1} = Best_{MP_S} + R_S$ and $P_B^{t+1} = Best_{MP_B} + R_B$

end while

1. **Generation of the First Population.** This represents the search's starting point and is created by randomly generating genes from the range of specified values.

2. **Selection Process.** All GAs use some form of mechanism to chose which individuals from the current population should go into the mating pool (MP) that forms the basis of the next population generation. To be effective, the selection mechanism should ensure that as diverse a range of fit agents make it into the MP as possible (especially in the early stages). A selection mechanism known to work well in such circumstances is *Tournament Selection* [4]. For this reason, it is the mechanism we employ to select from P_S^t and P_B^t , those individuals that will reproduce. Tournament selection works in the following way: k individuals

are randomly chosen from the population. The individual with the highest fitness among the selected k is placed in the MP. This process is repeated N times, where N is the size of the population. k is called the tournament size (which in our case is 2) and it determines the degree to which the best individuals are favoured [10]. Once the seller and buyer MPs have been created, the individual with the highest fitness in each pool is selected (respectively, $Best_{MP_S}$ and $Best_{MP_B}$). These agents will definitely form part of the new population. The remainder of the individuals in the next population, R_S and R_B , are created by applying crossover and mutation to the rest of the agents in the MP. Thus, the next generation of sellers and buyers (P_S^{t+1} and P_B^{t+1}) are composed of the best of the individuals of the old population plus a number of newly created strategies.

3. Crossover Process. This mechanism exchanges genetic material between individuals. We randomly select two individuals from the population. c crossover points are then randomly chosen and sorted in ascending order. Then the genes between successive crossover points are alternately exchanged between the individuals, with a probability P_c . In our experiments, P_c is 0.5 and the value of c is 6.

4. Mutation Process. Mutation is the other technique for creating individuals in new generations. It works by randomly selecting some of the genes present in the population in order to mutate. If a mutation occurs, a random value is chosen from the domain of the gene. This aims to avoid successive generations leading to local minima by introducing entirely new genetic structures. To provide a means of comparison between different populations, mutation is inhibited for some parameters (genes), specifically t_{max}^a and the issue-related genes. This is because they effectively determine the negotiation context and, therefore, they set the conditions in which a particular hypothesis is to be tested. The remaining genes are given a chance $P_m = 0.002$ of undergoing mutation.

5. Stopping Criterion. The simulations stop when the population is stable (95% of the individuals have the same fitness) or the number of iterations is bigger than a pre-determined maximum (100 in our case).

4. Experiments and Results

The aim of these experiments is to analyse the evolution of negotiation strategies for buying and selling agents. In particular, we want to determine which strategies are appropriate in which situations (given that there is no universal best strategy). Our preliminary evaluation of the service-oriented model indicates that negotiation behaviour is strongly influenced by the time available [5]. Thus, our evolutionary experiments are split into two groups: those in which there is a short time available ($t_{max}^a \in [2, 5]$) and those in which time is plentiful ($t_{max}^a \in [10, 20]$). For both cases, we consider the following types of agent population:

E_1 : All buying (selling) agents employ a tactic from the same family (subsection 4.1). This tactic family is used by the agent for all the issues under negotiation, although the specific family member employed varies for each issue. E.g. all buyers employ a Boulware tactic and all sellers employ a relative TFT tactic.

Type of analysis. Determine, on a per tactic family basis, how the tactic parameters vary against a given type of opponent tactic. E.g., if all buyers employ Conceder tactics and all sellers employ relative TFT tactics, how do the respective values of $\beta_{conceder}$ and δ_j evolve.

E_2 : Same as E_1 except: (i) a given agent can employ a weighted combination of tactics (in E_1 the Γ matrix contains only ones and zeros) and (ii) an agent can use tactics from different families for the different issues in a given negotiation (subsection 4.2).

Type of analysis. Determine for a mixed population which weighted combinations of tactics are successful.

The following parameters, which are fixed for all experiments, complete the environment setting specification: $|P_S| = |P_B| = 100$; k (communication cost) = 0.1; and number of negotiation issues = 3.

4.1. Homogeneous strategies

These experiments provide an indication of the relative merits of the different tactic families when they compete against one another and for a given tactic family they determine how negotiation behaviour evolves over time to better fit the environment and the type of opponent.

4.1.1 Rankings of tactic families (short t_{max}^a)

Figure 2 shows the pairwise rankings for all buyer and seller population combinations when there is a short time available for negotiation. Each matrix cell contains the average fitness for the buyers and sellers in the final population. In cell $F_{ij} = \begin{matrix} a \\ b \end{matrix}$ the two numbers correspond to the average fitness for buyer b using tactic i , and the average fitness for seller a using tactic j . Thus, for instance, we can see that when buyers using Boulware tactics negotiate against sellers using Conceder tactics, the sellers fair worse (-0.53 vs. -0.19).

Considering just the average performance of the tactic families, we observe that Resource are the best and Boulware the worst. Resource tactics do well because they stick firm and extract concessions from opponents when the resource (time) allows them to do so, but they give ground when necessary for the sake of making a deal. Boulware tactics are particularly bad because by sticking firm they make comparatively few deals (although the deals they make give them high utility).

<i>Seller</i> <i>Buyer</i>	Boulware	Conceder	Resource	Relative TFT	Random TFT	Average TFT	Mean
Boulware	-0.37 -0.47	-0.53 -0.19	-0.18 -0.53	-0.51 -0.25	-0.52 -0.24	-0.46 -0.29	-0.33
Conceder	-0.28 -0.3	-0.06 -0.16	0.18 -0.4	0.03 -0.23	0.07 -0.24	-0.12 -0.26	-0.26
Resource	-0.52 -0.19	-0.51 0.14	-0.13 -0.16	-0.42 0.01	-0.41 0.04	-0.54 -0.13	-0.05
Relative TFT	-0.38 -0.32	-0.24 -0.1	0.01 -0.53	-0.17 -0.27	-0.12 -0.12	-0.33 -0.24	-0.26
Random TFT	-0.38 -0.29	-0.23 -0.03	-0.03 -0.46	-0.19 -0.18	-0.17 -0.26	-0.27 -0.17	-0.23
Average TFT	-0.37 -0.42	-0.24 -0.08	-0.02 -0.55	-0.19 -0.22	-0.1 -0.21	-0.26 -0.25	-0.29
Mean	-0.38	-0.3	-0.03	-0.24	-0.21	-0.33	

Figure 2. Matrix F —Mean fitness of individuals in final buyer and seller populations with short t_{max}^a .

Examination of the pairwise tactic family interactions reveals a number of significant points. Firstly, the tactic which wins the highest percentage of the time is Resource, followed by Random TFT, Relative TFT and Boulware. The success of the imitative tactics is explained by the fact that they are exploited only by the Boulware and Resource families, that they are able to exploit Conceders, and that they have a similar behaviour negotiating with agents of their own family. However, the fact that Boulware is so successful is surprising (given its overall ranking). It occurs because opponents occasionally concede to their tough negotiation stance. But since they make comparatively few deals, they pull down the average fitness values for themselves and for their opponents. Secondly, agents using Resource tactics negotiating with agents using Conceder or Imitative tactics can get a better deal than the theoretical Nash equilibrium point. The reason for this is that Resource tactics have the exploitative power of the Boulwares, but also the flexibility to give ground (when time is running out) for the sake of making a deal. Thirdly, agents using Imitative tactics negotiating with Conceder opponents can get a better deal than the Nash equilibrium point. When the thread is so short that the imitative behaviour cannot be applied (δ is bigger than the steps in the thread) the default behaviour is tougher than that of the Conceder (resulting in a certain degree of exploitation). Fourthly, Imitative tactics are exploited by agents using Boulware and Resource tactics because they give ground in the first round of negotiation and, given that t_{max}^a is small, this forms a significant part of the interaction.

4.1.2 Evolution of tactic families (short t_{max}^a)

Here we analyse how strategies evolve *through* populations, rather than just focusing on the end point.

1. When Conceder agents negotiate with Boulware and Resource agents, the value of $\beta_{conceder}$ reduces over time (Figure 3). Thus, Conceders react to being exploited by becoming less conciliatory in nature.

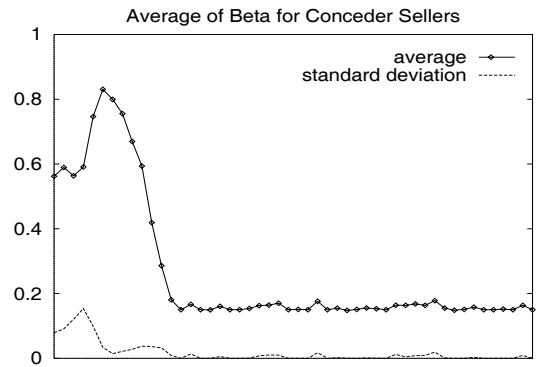


Figure 3. Conceder Sellers against Resource Buyers

2. Agents employing imitative tactics tend to increase the value of δ_j . This happens because if δ_j is large in comparison to t_{max}^a , then in almost all the cases no actual imitation takes place and the default Conceder tactic (with $\beta_j = 2$) is applied. This default is comparatively tough in short negotiations and so the agent is able to exploit more conciliatory agents.

3. Agents using Resource tactics have the best performance against all the other tactic families. They tend to increase the value of μ_j , thus becoming more Boulware like, so they can exploit more conciliatory tactics (Figure 4).

4.1.3 Experiments with large t_{max}^a

We repeated the above set of experiments with a large time ($t_{max}^a \in [10, 20]$) to reach an agreement (Figure 5). In most cases, results are broadly similar to those when t_{max}^a is short. Again Resource tactics are the most successful but for a large t_{max}^a the Conceder tactics improve their rating. With more time available for negotiation, Conceder gives ground more slowly and this forces resource agents to concede more than when t_{max}^a is short (because the length of

<i>Seller</i> <i>Buyer</i>	Boulware	Conceder	Resource	Relative TFT	Random TFT	Average TFT	Mean
Boulware	-1.09 -1.04	-0.86 -0.06	-0.86 -0.99	-0.9 -0.68	-1.05 -0.63	-0.99 -0.54	-0.66
Conceder	-0.12 -0.93	-0.12 -0.04	0.04 -0.44	0.09 -0.37	0.03 -0.37	-0.4 -0.29	-0.41
Resource	-1.12 -0.91	-0.42 0.14	-0.44 -0.46	-0.62 -0.33	-0.56 -0.33	-0.82 -0.28	-0.36
Relative TFT	-0.87 -0.98	-0.36 0.08	-0.43 -0.53	-0.52 -0.4	-0.4 -0.34	-0.43 -0.35	-0.42
Random TFT	-0.81 -0.93	-0.37 0.14	-0.37 -0.58	-0.39 -0.38	-0.49 -0.44	-0.43 -0.36	-0.42
Average TFT	-0.75 -1.01	-0.37 0.01	-0.38 -0.52	-0.44 -0.3	-0.41 -0.39	-0.48 -0.36	-0.43
Mean	-0.79	-0.42	-0.41	-0.46	-0.48	-0.59	

Figure 5. Mean fitness of individuals in the final buyer and seller populations with large t_{max}^a

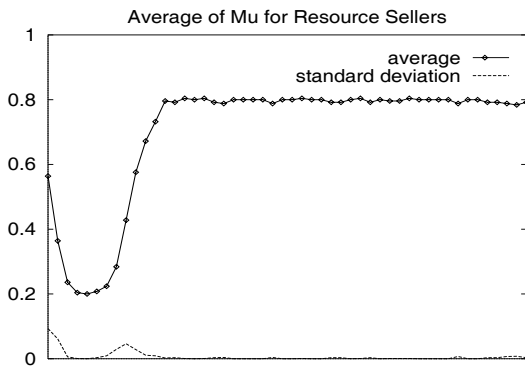


Figure 4. Resource Sellers against Average TFT Buyers

the thread makes the quotient $-\frac{\mu_j^a}{|X_{a \leftrightarrow b}^t|}$ smaller and hence the concession is bigger). The major difference, however, is that imitative tactics begin to exploit their imitative capability. Since negotiations generally involve more interchanges, the imitators have more opportunities to assess their opponents' behaviour and to respond in kind.

4.2. Weighted combinations of tactics

Here we allow agents to employ weighted combinations of tactics to the different issues (recall that previously the Γ matrix contained only ones and zeros). Again we consider environments in which negotiation time is scarce and those in which time is plentiful. Analysis of the Γ matrix evolution reveals the following points of interest. Firstly, equilibrium is always reached with very similar Γ matrixes (whose average is shown below). This indicates the matrix represents a robust solution since it is reached despite wide variations in the starting weights. Secondly, the values inside the matrix (i.e the individual γ_{ijs}) at the equilibrium point are very similar—they have an average of 0.17 and a standard deviation of 0.02. This means no one strategy dominates an agent's negotiation behaviour. Thirdly, we ob-

serve that agents who use a weighted combination of tactics are more robust to varying negotiation situations. For example, the average utility for buyers using a weighted combination in a short t_{max}^a is -0.26 and in a long t_{max}^a it is -0.27. In contrast, for the pure tactic case the average figures for short and long negotiations can vary significantly (eg for conceder buyers it varies between -0.26 to -0.41).

Exp I	Boulware	Conceder	Resource	Relative TFT	Random TFT	Average TFT
<i>issue</i> ₁	0.12	0.21	0.14	0.19	0.19	0.13
<i>issue</i> ₂	0.14	0.21	0.17	0.17	0.15	0.16
<i>issue</i> ₃	0.11	0.28	0.14	0.14	0.12	0.17

5. Related work

Research in negotiation models has been pursued in many different fields of enquiry. Each field has concentrated on different aspects of negotiation, making the assumptions that are pertinent for the goal of their study. For reasons which we have outlined elsewhere [5], much of this extant work is inappropriate for the types of negotiations we encountered in the domain of business process management. Hence we needed to devise a new model to better reflect our needs. In this paper, however, the main focus is **not** on the model per se, but on how we can find the best strategies for negotiation in this model under various conditions.

Axelrod suggests there are three basic ways of determining the best interaction strategy in a given context: tournament, ecological and territorial [1]. The tournament mechanism we apply is basically the same as Axelrod's, but our agents have to make decisions over a much wider range of issues and over continuous variables such as price (he restricts himself to the iterated Prisoner's Dilemma in which agents can only choose to cooperate or defect). It extends Axelrod's work by considering more complex tactics, but, nevertheless, retains the simplicity of the Prisoner's Dilemma environment [7]. Our approach also incorporates Axelrod's ecological model, but we include mutation and

crossover as operations in the evolution of populations. The territorial approach is left as future work.

A number of researchers have attempted to use GAs and evolutionary programming to find optimal interaction strategies. Oliver considers negotiation strategies [11] and his work has several similarities to our own (eg each negotiating agent is a chromosome and the parameters of the negotiation model are genes in the chromosome). However, his model is much simpler. (Offers are accepted if they have a utility over a preset threshold. Also he encodes counter-offers as part of the genetic material, meaning they have limited sophistication.) The extra complexity required by our model means additional work was needed when designing the population evolution mechanism. For example, our encoding necessarily entails that there are gene interdependencies which need to be respected during the crossover and mutation operations (Oliver has no such dependencies).

Rosin [12] considers a competitive, co-evolution in which fitness is based on direct competition between individuals selected from two independently evolving populations of “hosts” and “parasites”. They use: a competitive fitness sharing function, a shared sampling method to select a diverse set of parasites, and a “hall of fame” method to save good individuals from prior generations. Although there are several similarities, our work differs in that we use a fitness function based on the sum of the score across all competitions, we evolve the seller and buyer populations simultaneously, and we perform a tournament between all the sellers and buyers.

6. Discussion and Further Work

This paper presents an empirical evaluation of a range of negotiation strategies and tactics in a number of different types of environment. We take an evolutionary approach — encoding negotiation parameters as genes in a GA. The aim of the evaluation is to assess the operational benefits and drawbacks of a number of negotiation strategies. To this end, we have presented a number of concrete results about the relative merits of particular tactics and strategies. Moreover, we have shown how, in some cases, successive populations of agents following a particular role (buyer, seller) exploit the population of agents using the dual role. These successive generations also give an indication about how strategies can be modified to minimise the opportunity of exploitation. Finally, we have demonstrated the usefulness of agents employing a cocktail of tactics —both for different negotiation issues and, in combination, for a single issue. Our previous theoretical work [13] indicated such a hybrid approach was possible, but had not demonstrated that it was useful.

Our future work will extend the experimentation in two directions. Firstly, we want to introduce more reactivity

into the agents. Here agents have a fixed Γ and participate in a single negotiation round with each of their opponents to determine their fitness. We would like to see how changing Γ within the lifetime of an agent allows it to adapt to repeated negotiations with its opponents. Secondly, we wish to extend our evolutionary approach to cover the territorial case [1]. Here populations are distributed on a 2-D surface and agents can only negotiate with their neighbours. Depending on the outcome of this negotiation, the agent may adopt the strategy of its neighbour if it is more successful. This type of experiment will allow us to observe the changes in the territorial distribution pattern of the strategies.

Acknowledgements.

The research reported in this paper is supported by the ESPRIT LTR 25500-COMRIS *Co-Habited Mixed-Reality Information Spaces* project and by the Spanish CICYT project SMASH, TIC96-1038-C04001.

References

- [1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, Inc., New York, USA, 1984.
- [2] R. Axelrod. *The Complexity of Cooperation*. Princetown University Press, 1997.
- [3] K. Binmore. *Fun and Games*. Heath, 1992.
- [4] T. Blickle and L. Thiele. A comparison of selection schemes used in genetic algorithms. Technical Report 11, TIK, May 1995.
- [5] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. J. of Robotics and Autonomous Systems*, (to appear), 1998.
- [6] D. Goldberg. *Genetic Algorithms in Search, Optimisation and Matching Learning*. Addison-Wesley, 1989.
- [7] A. Ito and H. Yano. The emergence of cooperation in a society of autonomous agents - the prisoner’s dilemma game under the disclosure of contract histories. In *Proc. of the First International Conference on Multi-Agent Systems*, 201-208. San Francisco, USA, 1995.
- [8] N. R. Jennings, P. Faratin, M. J. Johnson, T. J. Norman, P. O’Brien, and M. E. Wiegand. Agent-based business process management. *Int. Journal of Cooperative Information Systems*, 5(2-3):105-130, 1996.
- [9] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, (94):79-97, 1997.
- [10] B. L. Miller and D. E. Goldberg. Genetic algorithms, tournament selection, and the effects of noise. *Complex Systems*, 9(3):193-212, 1995.
- [11] J. R. Oliver. *On Artificial Agents for Negotiation in Electronic Commerce*. PhD thesis, The Wharton School, 1994.
- [12] C. D. Rosin. *Coevolutionary Search Among Adversaries*. PhD thesis, University of California, San Diego, 1997.
- [13] C. Sierra, P. Faratin, and N. R. Jennings. A service-oriented negotiation model between autonomous agents. In M. Boman and W. V. de Velde, editors, *MAAMAW’97*, LNAI 1237, 17-35, 1997.