A REVIEW OF LEARNING METHODS ENHANCED IN STRATEGIES OF NEGOTIATING AGENTS

Marisa Masvoula, Panagiotis Kanellis and Drakoulis Martakos
Department of Informatics and Telecommunications, Nationaland Kapodistrian University of Athens
University Campus, Athens15771, Greece

Keywords: Learning in Negotiations, Negotiation Strategies.

Abstract: Advancement of Artificial Intelligence has contributed in the enhancement of agent strategies with learning techniques. We provide an overview of learning methods that form the core of state-of-the art negotiators. The main objective is to facilitate the comprehension of the domain by framing current systems with respect to learning objectives and phases of application. We also aim to reveal current trends, virtues and weaknesses of applied methods.

1 INTRODUCTION

We view negotiation as an iterative procedure among participants who seek to reach mutually acceptable agreements. Different disciplines are taken to study the negotiation environments and dynamics of the interactions, resulting to a variety of frameworks and perspectives (Raiffa, 1982). On the top of the hierarchy lay behavioral and management sciences that provide descriptive approaches, and focus on "how negotiators behave in reality". Economic theories and formal mathematical models have been used in an attempt to 'quantify' the negotiation problem, find points of equilibrium and suggest optimal behaviors with respect to goals and aspirations of the engaged parties. The most commonly used are game-theoretic tools that have been critiqued because of the unrealistic requirement of unlimited computational power and strong assumption of common knowledge. In reality negotiators have to deal with vague data, limited information, uncertainty and time restrictions. The issue of limited information and the attempt to produce 'suboptimal' strategies has been addressed by heuristic-based approaches. Computer science has contributed to the field of negotiations with the use of information systems that move the negotiation arenas to electronic settings, and with the development of support systems that assist users in the various negotiation stages. Full automation is also supported with the use of agents who represent human users and undertake the process. Specifically, (Kersten and Lai, 2007) identify four kinds of software that have been designed for negotiations; e-negotiation tables (ENT) which are passive systems oriented in facilitating communication of participants, negotiation support systems (NSS), which are software tools that support participants in various negotiation activities, negotiation software agents (NSA), designed with the purpose to automate one or more negotiation activities, and negotiation agent assistants (NAA), agents designed to provide advice and critique, without engaging directly in the negotiation process.

As we move to the field of negotiation analysis, different negotiating behaviors, reflected through the strategies, result to different outcomes, measured with respect to individual or joint gain. A number of research efforts concentrate on conducting extensive experiments to analyze the interactions in different settings. It has been proved that there does not exist a universal best strategy, rather it depends on the negotiation domains, protocols, participants' goals and attitude towards risk, as well as counterparts' strategies. Learning techniques have proved to add value to negotiators since they extend their knowledge and perception of the domain.

Masvoula M., Kanellis P. and Martakos D. (2010).
A REVIEW OF LEARNING METHODS ENHANCED IN STRATEGIES OF NEGOTIATING AGENTS.
In *Proceedings of the 12th International Conference on Enterprise Information Systems - Artificial Intelligence and Decision Support Systems*, pages 212-219
DOI: 10.5220/0002897102120219

2 CONTRIBUTIONS

An overview of negotiation support and e-negotiation systems can be found in Kersten and Lai, 2007, where the authors classify and analyze the use of negotiation software in socio-technical systems. Other surveys concentrate on the behavior of software agents enhanced with learning skills, and analyze their interactions in dynamic and stationary environments (Busoniu, Babuska and De Shutter, 2008). Our research concentrates on learning techniques enhanced in strategies of negotiating agents, therefore an overview more similar to our work, is provided by (Braun et al., 2006), where learning methods of negotiating software agents are presented. In Braun et. al (2006) there is a general presentation of learning techniques adopted by agents in various stages of the negotiation procedure. In this research, we rather concentrate on learning skills enhanced in agents' strategy and list a number of representative implementations, which best describe the domain. We provide a comprehensive frame of the domain by classifying the systems with respect to the learning method and phase of application. In the following sections we define the classification criteria and present a variety of learning methods incorporated in strategies of negotiating agents.

3 NEGOTIATING STRATEGIES ENHANCED WITH LEARNING

Negotiation process model adopted in most frameworks discriminates strategy selection at the planning phase and strategy update during discourse. This has lead to the existence of two schools when it comes to studying negotiation strategies. The first is concerned with the selection of a strategy at a pre-negotiation phase, during formulation of the problem. The second is concerned with strategy update, the change of behavior during discourse, which may be due to changing preferences or environmental parameters. We devise agents to those who intuitively adjust their behavior, and to those who use reasoning skills in the decision-making process. In the former category agents engage in learning methods that differ to the extent of knowledge exploration and exploitation. Specifically, explorative techniques also imply the search for new solutions, while repetitive techniques are based on knowledge reuse. For agents who engage in reasoning processes to decide upon appropriate actions, learning is introduced in the form of predictive decision making, where estimations of factors that influence strategy selection or update serve as input to the agents' decision making. With respect to these factors we discriminate the following three categories: explorative, repetitive and predictive which may be applied either at the planning phase for initial strategy selection or during discourse.

3.1 Explorative Strategies

Explorative strategies are equivalent to search techniques that follow a trial and error learning process until some convergence condition is satisfied. Such techniques are Q-learning and Genetic algorithms. Q-Learning is a reinforcement learning algorithm that maps state-action pairs to values named Q-values. When an agent performs an action, he receives a reward that updates the Q-value of the corresponding state. Exploration of new actions, known as Boltzman explorations, is usually controlled by a temperature parameter. Q-learning may be applied to learn from previous encounters where trials are the previous negotiations, or from the current encounter, where trials are the previous offers.

(Cardoso and Oliveira, 2000) implemented a Q-learning agent who acts in a dynamic environment and tries to estimate which combination of tactics to use in each state. Knowledge is acquired from previous encounters, since the state is defined by environmental parameters that relate to the number of agents and available time of the adaptive agent. Actions are defined as combinations of tactics and are assessed at the end of negotiation, as positive rewards if a deal is achieved, or negative rewards (penalties) if negotiation ends without an agreement. The measure of the reward (Q-value) is determined by the utility or benefit that the procedure incurred to the agent. Experimental results showed that the agents increased their utility with time, though in some cases it took too long to achieve good results.

When Q-learning is applied to the current encounter feedback from the opponent is required after each bid presentation, in order to compute the Q-value. Such an implementation can be found in (Oliveira and Rocha, 2001). The state is defined as the current offer in the form of a sequence of values, and the action specifies how each attribute should change (increase, maintain, or decrease) in order to generate the next offer. If the attribute space is continuous then change is realized by a predefined amount, while if it is ordinal, it moves to the next enumerated value. After sending an offer, the learning agent receives qualitative feedback from the negotiating partner and calculates the reward of its action, which is used to update the Q-value of the corresponding state-action pair. Added to the weakness of many iterations, this approach also suggests the use of opponents' feedback. It is not guaranteed though that the opponent will agree to engage in such protocol or that he will be truthful. Another issue that is left open relates to the ability of Q-Learning technique to deal with large state-action spaces.

The second 'family' of explorative strategies consolidates in Genetic algorithms, optimization techniques inspired by evolution. A population of candidate solutions, encoded into chromosomes is generated and evaluated. The best solutions are assigned the highest fitness and are combined with the use of selection, crossover and mutation techniques, to create new candidate solutions that comprise the next generation. The cycle continues until a stopping condition, usually related to a stable average fitness, is met. This technique is adopted by negotiating agents who seek for robust strategies. Application of GAs at the planning phase is a tool that facilitates analysis of the dynamics of the interaction. It is used to search strategies that are best responses to the counterparts' best strategies, starting from random points. Oliver, (1996) describes a framework where strategies are formed by simple, sequential rules that consist of acceptance thresholds and counterproposals. For each negotiator a random population of strategies is generated. The testing of different negotiation strategies is repeated and the fitness of each one is determined by the utility it incurs to the agent. After a number of strategies have been tested the genetic algorithm is run in order to generate a new population of strategies and this procedure is repeated until an exit condition is satisfied. In (Matos et al., 1998) we find application of genetic algorithms in domains where strategies are defined as a combination of tactics (Faratin et al., 1998). In such approaches the chromosomes comprise of specific strategic information such as deadlines, reservation values, weights of tactics and parameters specifying each tactic. The simulations were repeated until stabilization of populations (95% of the individuals had the same fitness) or until the number of iterations reached a predefined threshold. Gerding, van Bragt and La Poutre, (2000), analyze the negotiation results achieved by GA-based agents, with respect to fairness and symmetry. Such applications of GA are not particularly interesting when viewed in a single negotiation instance. The major drawback is that it requires many iterations, and each iteration is a negotiation instance. On the contrary, in cases where GAs were applied during the current discourse, populations of chromosomes were used to represent the population of feasible offers. Such application can be found in (Lau, 2005) where the fitness of each offer is measured with respect to its distance from the most preferred offer, the distance from the opponents' previous offer and the time pressure. In each round the offers considered fit by the agent may change. This technique aids the agent to gradually learn and adapt to its opponents' preferences. This approach does not assume knowledge of prior negotiations and it could be applied in dynamic environments. An obvious limitation is that the algorithmic complexity increases with the increase of alternatives of each negotiable attribute.

3.2 Repetitive Strategies

In this category we place strategies which follow a routine-based concept; Substance of routines lays on the specific knowledge acquired by the repeated execution of an act combined with the ability to apply this knowledge to specific situations. It has the potential to substitute deliberate planning and decision making since it is used to determine which operations to implement in order to achieve certain intended state. Routinization techniques force agents to develop 'best practices'. The most commonly used is Case-based reasoning (CBR), where previously solved cases are maintained in a case base and when a new problem is encountered, the system retrieves the most similar case and adapts the solution to fit the new problem as closely as possible. CBR is common in negotiations, particularly in the planning phase supporting the process of strategy or supplier selection, or during discourse in argumentative frameworks. A commonly stated risk posed by routinization is the application of ineffective acts. Routines in dynamic environments have proved to be of degrading efficiency, the so called "acting inside the box situation". As stated by Nelson and Winter, (1982) with increasing repetitions, decision making prior to the operation tends to decrease. The use of routines entails rigidity and once a solution is established, it is not further questioned. Another weakness accumulates on the requirement to store the case base and the difficulty to collect the information that best discriminates different situations. In (Sycara, 1988) PERSUADER , a program that acts as a labor mediator, enters in negotiation with each of the parties, the union and the company, proposing and modifying compromises until a final agreement is reached. The PERSUADER's input is a set of conflicting goals and the output is either a plan or an indication of failure. Additionally the system is capable of persuading the parties to change their evaluation of a compromise. CBR is used to keep track of cases that have worked well in similar circumstances. The most suitable case is retrieved from memory and adapted to fit the current situation. If the parties disagree, PERSUADER appropriately repairs the compromise and updates the case base or generates arguments to change the utilities of the disagreeing parties. The system integrates CBR and Preference Analysis, a decision theoretic method, to construct the initial compromise in the planning phase. If previous similar cases are not available, the PERSUADER uses Preference Analysis to find suitable compromises. Another CBR-based approach can be found in (Air Force Research Laboratory, 2003) which describes multi-sensor target tracking, in a cooperative domain, where each agent controls one sensor and consumes resources (cpu, time, memory etc.). The agents are motivated to share their knowledge about the problem, based on their viewpoint, in an effort to arrive to a solution. The model uses case-based reasoning to retrieve the most similar case based on the incurred utility, adapts the case to the current situation and uses the cases' strategy to perform negotiations. An application of CBR to the current discourse can be found in (Wong et al., 2000) who implemented a support system that assists negotiators with agent opponents over used cars. The system matches current negotiation scenario with previous successful negotiation cases, and provides appropriate counter-offers for the user, based on the best-matched negotiation case. A contextual case organization hierarchy is used as an organization structure for categorizing the negotiation cases and similarity filters are used to select the best-matched case from the retrieved set of cases. Strategic moves, concessions and counter-concessions of a past discourse, are adapted to the current situation. If no case is found based on the organization hierarchy, the buyer uses a default strategy. This approach considers a single negotiable attribute, price, and does not consider learning from failure. The virtues of repetitive strategies summarize to saving planning and decision making costs by reusing previously applied solutions. The trade-off, often termed the 'routine trap', relates to the increased risk of applying inefficient acts, if dynamics of the negotiation environment change over time.

3.3 Predictive Strategies

The third group relates to estimating opponents' strategic parameters and preferences, as well as future behaviors, in order to select the most appropriate acts, assessed in terms of individual or joint satisfaction. When such predictions are encountered in the planning phase, the agent may rank his opponents and decide to negotiate only with the most prosperous ones, to save time and resources. In (Brzostowski and Kowalczyk, 2005) the buyer agent uses CBR to predict the outcome of a future negotiation, assuming it is in a particular situation. The situation is characterized by the negotiation strategy and the preferences of the buyer. This approach follows principles of possibility decision theory, and is referred as possibilistic case-based reasoning. The likelihood of successful negotiation is derived from the history of previous interactions in the form of a possibility distribution function. The expected utility of the future negotiation is an aggregate of the distribution function with the current agents' utility and is used to rank the negotiation partners. When it comes to using predictive strategies during the current discourse, the focus lies on the estimation of opponents' strategic parameters and preferences. A significant number of applications use Bayesian learning techniques to update beliefs about the opponents' structure. An early application can be found in (Zeng and Sycara, 1998) who developed Bazaar, a negotiating system which uses a Bayesian network to update the knowledge and belief each agent has about the reservation value of his opponent. Estimation of the opponent's reservation value contributes to approximating his payoff function and provides the agent with the ability to propose more attractive offers to his counterpart. The negotiation domain in Bazaar was rather simplified, as the authors assumed a finite set of offers, and the computational ability of agents to calculate expected payoffs for all possible offers in order to decide the one that maximizes their utility value. Bazaar, as most systems that apply Bayesian methods, has also been critiqued on the requirement of initial knowledge of many probabilities. Probability distributions of hypothesis representing potential reservation prices of the opponents, as well as domain knowledge of previous offers represented as conditional statements, constitute the prior knowledge of the system. These probabilities are estimated based on background knowledge, previously available data and assumptions about the form of the underlying distributions. Nevertheless if the distributions change, the model will no longer produce reliable estimations. To the stated weaknesses we add the fact that illustration was available only for a single attribute (price). Other approaches based on Bayesian learning can be found in (Buffet and Spencer, 2007) where the authors presented a classification method for learning opponents' preference relations during bilateral multi-issue negotiations. Similar candidate preference relations were grouped into classes, and a Bayesian technique was used to determine the likelihood that the opponents' true preference relations lay in a specific class. Negotiation concerned subsets of a set of objects and the goal was to increase knowledge upon the counterparts' preferences, so that an effective strategy could be devised. As the authors suggest, building an initial set of classes is a difficult task, depending on the specifics of the problem and additional information about the other party. Another work using a Bayesian classifier can be found in (Bui et al., 1999) where agents assign probability distributions about their opponents' preference structure, in order to reduce the overall communication cost in a co-operative framework. The system suffers from the difficulty of collecting prior probabilities as all pre-mentioned Bayesian-based approaches.

Estimating opponents' strategic parameters has also been approached by statistical methods, mainly based on non-linear regression. Hou (2004) describes a non-linear regression-based model to predict the opponents' family of tactics and specific parameters. This approach is restrictive in that it relies on the assumption of a known function form that models the concessions of the opponents. The author has assumed two non-linear functions that model time and resource dependant tactics, based on (Faratin et al., 1998). The objective was to fit the function to the opponents' previous offers, by estimating the vector of parameters that minimizes the distance of the actual offer and the estimated one. The optimization problem was dealt with an iterative method combining grid search and the Marquardt algorithm. Non-linear regression was applied in each negotiating round of the predicting agent and the authors adopted a number of heuristics to fix their prediction upon opponents' deadline and reservation value. Although this approach adds value to the negotiating agent, experiments were only conducted with pure strategies, where extreme behaviors are easier to distinguish. An application of non-linear regression with mixed strategies can be found in (Brzostowski and Kowalczyk, 2006 a). The purpose was to predict the opponents' future offers, foresee potential negotiation threads and adopt the strategy that will result to the most beneficial discourse. The authors developed four models to address the issue of mixed strategies that resulted from a combination of time and behavior dependent tactics with various weights assigned. Although this model involved more strategies than the one mentioned earlier, it does not extensively cover the space of possible strategies as discussed by (Faratin et al., 1998). The complexity is expected to increase as the number of models increases, therefore extending this solution would not be an easy task. (Brzostowski and Kowalczyk, 2006 b) take an approach based on the difference method, in order to predict the opponents' future offers. This method has the advantage that the agent does not need to know precisely the opponents' strategic function. The authors assume that the opponent uses a mixture of time and behavior dependent tactics and try to determine to which extent he imitates the predicting agents' behavior and to which extent he responds to a time constraint imposed on the encounter. This was achieved with the use of two criteria combined with time depending and imitation depending predictions, obtained from the previous offers of the opponent, and from a combination of opponents' and predicting agents' offers respectively. Results have proved that the method is not as accurate as the non-linear regression and the accuracy of the weights assessments still needs to be improved. The area of predicting opponents' offers during discourse has attracted much attention, since an agent may refine his strategy and increase individual or overall gain. The current trend concentrates on the use of connectionist approaches. Neural networks are universal function approximators and the proof is based on the well-known Kolmogorov theorem. Oprea (2003) presents a study where a neural network with one hidden layer is used to predict the opponent's next offer. Past opponents' offers are modeled as time-series and the three most recent are used for the networks' input. The agent refined the offer he was about to send in each round based on the prediction of his opponents' next move. Nevertheless it is not explained why the authors selected only the three previous offers of the opponent and why the predicting agents' previous offers are not accounted, (they assumed only time dependency of the responses). A similar approach is followed by (Carbonneau et al., 2006) who developed a predictive model based on neural networks, with the purpose to optimize an agents' current offer. This optimization was achieved by conducting "What-if" analysis over the set of possible alternatives, and selecting the proposal that would result to the most beneficial response. The neural network had thirty nine inputs, resulting from past offers, the current offer, and statistical information. It also had four outputs, one for each predicted attribute of the offer, and ten hidden neurons. As the authors state, the model has been tested for a particular negotiation case in a static domain and the accuracy of its predictions may be less adequate in the general case. Another neural network-based approach can be found in (Lee and Ou-Yang, 2009) who implemented a negotiation support tool of the demander in a supplier selection auction market. The network was used to forecast the suppliers' next bid price, and allow the demander to appropriately choose among a list of alternatives. Nine inputs were used to reflect environment-dependant information as well as bid prices of the three previous offers, twelve hidden neurons obtained by means of trial-and-error experiments, and one output neuron that reflected the predicted bid price. The network was trained with the back-propagation algorithm, which is particularly slow. A different approach, where prediction of opponent's next offer was carried only once during the discourse, in the pre-final round, is found in (Papaioannou, Roussaki, and Anagnostou, 2006). The purpose was to increase the utility of the final agreement. Experiments were conducted over two different types of neural networks, MLPs and RBFs. The latter proved to outperform MLPs in small datasets. Opponent's future moves have also proved valuable in cases where the agents used forecasts to detect unsuccessful negotiations from an early round. Such approaches have been discussed by (Papaioannou et al., 2008) where the decision of the agents to withdraw or not from the current negotiation is supported by determining the providers' offer before the clients' deadline expires. The weakness of current connectionist approaches used in predictive decision making summarizes to the restriction of being tested solely in bounded spaces, where opponents followed static strategies, or negotiations were conducted over fixed, pre-defined alternatives. What happens if opponents also engage in adaptive negotiation strategies and update their behavior during discourse? An open and challenging issue lays in the application of predictive decision making in environments with changing data distributions.

4 CONCLUSIONS

This work provides a review of the learning methods adopted by negotiating agents who either adopt intuitive strategies or engage in predictive decision making. We aimed to provide a categorization with respect to the learning objectives, in order to facilitate comprehension of the domain. We have discriminated explorative, repetitive and predictive strategies applied at a pre-negotiation phase or during discourse. Under this frame we presented various systems that reflect the trends of learning in negotiation strategies, as well as the weaknesses depending on the applied domain. Virtues and weaknesses are summarized to table 1 of the appendix.

REFERENCES

Albesano, D., Gemello, R. Laface, P., Mana, F., Scanzio. S. (2006) Adaptation of Artificial Neural Networks Avoiding Catastrophic Forgetting. In *Proceedings of the International Joint Conference on Neural Networks, IJCNN 2006* (pp. 1554-1561).

Air Force Research Laboratory. (2003). A Case-based Reflective Negotiation Model.

Braun, P., Brzostowski, J., Kersten, G., Kim, J., Kowalczyk, R., Stecker, S., Vahidov, R.(2006). e-Negotiation Systems and Software Agents: Methods, Models, and Applications. In Jatinder N. D. G., Guisseppi A. F. and Manuel M. T. (Eds.), *Intelligent decision-making support systems: Foundations, Applications and Challenges*, (pp. 271-300). London: Springer

Brzostowski, J. and Kowalczyk, R. (2005). On possibilistic case-based reasoning for selecting partners for multi-attribute agent negotiation. In *Proceedings of the Fourth international Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 273-279). New York: ACM

Brzostowski, J. and Kowalczyk, R. (2006a). Adaptive Negotiation with On-Line Prediction of Opponent Behaviour in Agent-Based Negotiations. In *Proceedings of the IEEE/WIC/ACM international Conference on intelligent Agent Technology* (pp. 263-269). Washington, DC: IEEE Computer Society.

Brzostowski, J. and Kowalczyk, R. (2006b). Predicting partner's behaviour in agent negotiation. In Proceedings *of the Fifth international Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 355-361). New York: ACM.

Busoniu, L., Babuska, R., De Shutter, B. (2008). A Comprehensive Survey of Multiagent Reinforcement Learning. *IEEE Transactions on Systems, Man and Cybernetics,* 38(2), 156-172.

Bui, H. H., Venkatesh, S. and Kieronska, D. (1999). Learning other agents' preferences in multi-agent negotiation using the Bayesian classifier. *International Journal of Cooperative Information Systems*, 8(4), 275-294.

Buffett, S. and Spencer, B. (2007). A Bayesian classifier for learning opponents' preferences in multi-object automated negotiation. *Electronic Commerce Research and Applications*, 6(3), 274-284

Carbonneau, R., Kersten, G. E., and Vahidov, R. (2008). Predicting opponent's moves in electronic negotiations using neural networks. *Expert Systems with Applications: An International Journal,* 34 (2), 1266-1273.

Cardoso, H. L. and Oliveira, E. (2000). Using and Evaluating Adaptive Agents for Electronic Commerce Negotiation. In M. C. Monard and J. S. Sichman (Eds.), *Proceedings of the international Joint Conference, 7th Ibero-American Conference on Ai: Advances in Artificial intelligence* (pp. 96-105). London: Springer-Verlag.

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.

Gerding, E. H., van Bragt, D. D. and La Poutre, J. A. (2000). Multi-Issue Negotiation processes by Evolutionary Simulation: Validation and Social Extensions. Technical Report. UMI Order Number: SEN-R0024., CWI (Centre for Mathematics and Computer Science).

Hou, C. (2004). Predicting Agents Tactics in Automated Negotiation. In *Proceedings of the intelligent Agent Technology, IEEE/WIC/ACM international Conference* (pp. 127-133). Washington, DC : IEEE Computer Society.

Kersten, G. E and Lai, H. (2007). Negotiation Support and E-negotiation Systems: An Overview. *Group Decision and Negotiation,* 16(6), 553-586.

Lau, R. Y. (2005). Adaptive negotiation agents for e-business. In *Proceedings of the 7th international Conference on Electronic Commerce* (pp. 271-278). New York: ACM.

Lee, C. C. and Ou-Yang, C. (2009). A neural networks approach for forecasting the supplier's bid prices in supplier selection negotiation process. *Expert Systems with Applications,* 36(2), 2961-2970.

Li, C., Giampapa, J., and Sycara, K. (2004). Bilateral Negotiation Decisions with Uncertain Dynamic Outside Options. In *Proceedings of the First IEEE international Workshop on Electronic Contracting* (pp. 54-61). Washington, DC: IEEE Computer Society.

Matos, N., Sierra, C., and Jennings, N. (1998). Determining Successful Negotiation Strategies: An Evolutionary Approach. In *Proceedings of the 3rd international Conference on Multi Agent Systems* (pp. 182-189). Washington, DC: IEEE Computer Society.

Nelson, R.R and Winter, S. G. (1982). *An Evolutionary theory of economic change.* Cambridge, MA and London: The Belknap Press of Harvard University Press.

Oliveira, E. and Rocha A. P. (2001). Agents Advanced Features for Negotiation in Electronic Commerce and Virtual Organizations Formation Process. In F Dignum and C Sierra, (Eds.), *Agent Mediated Electronic Commerce, the European AgentLink Perspective* (pp.77-97). Heidelberg: Springer Berlin

Oliver, J. R. (1996). A machine-learning approach to automated negotiation and prospects for electronic commerce. *J. Manage. Inf. Syst*., 13 (3), 83-112.

Oprea, M. (2003) The Use of Adaptive Negotiation by a Shopping Agent in Agent-Mediated Electronic Commerce. In: *Proceedings of the 3rd International Central and Eastern European Conference on Multi-Agent Systems, CEEMAS03*. (pp. 594-605). Heidelberg: Springer Berlin.

Papaioannou, I. V., Roussaki, I. G., and Anagnostou, M. E. (2006). Comparing the Performance of MLP and RBF Neural Networks Employed by Negotiating Intelligent Agents. In *Proceedings of the IEEE/WIC/ACM international Conference on intelligent Agent Technology (*602-612). Washington, DC: IEEE Computer Society.

Papaioannou, I., Roussaki, I., and Anagnostou, M. (2008). Detecting Unsuccessful Automated Negotiation Threads When Opponents Employ Hybrid Strategies. In D. Huang, D. C. Wunsch, D. S. Levine, and K. Jo, (Eds.), *Proceedings of the 4th international Conference on intelligent Computing* (pp. 27-39). Heidelberg: Springer Berlin.

Raiffa, H. (1982). *The Art and Science of Negotiation*. Cambridge, MA and London: The Belknap Press

Rau, H., Tsai, M., Chen, C., and Shiang, W. (2006). Learning-based automated negotiation between shipper and forwarder. *Comput. Ind. Eng.*, 51(3), 464-481

Robinson, W. N and Volkov, V. (1998). Supporting the negotiation life cycle. *Communications of the ACM,* 41(5), 95-102.

Van Bragt, D. D. and La Poutré, J. A. (2003). Why Agents for Automated Negotiations Should Be Adaptive. *Netnomics*, 5(2), 101-118.

Wong, Y. W., Zhang, D. M., and Kara-Ali, M. (2000). Negotiating with Experience. Knowledge-Based Electronic Markets, Papers from the AAAI Workshop, Technical Report. AAAI Press

Zeng, D. and Sycara, K. (1998). Bayesian learning in negotiation. *Int. J. Hum.-Comput. Stud*., 48(1), 125-141.

APPENDIX

Table 1: Learning Methods applied in negotiations.

Explorative	**Virtues**	**Weaknesses**
GA (in planning phase)	1. Reach optimal strategy 2. Analyze negotiation interactions	Increased number of iterations due to large strategy space
GA (during discourse)	Adapt to opponent's responses, approach pareto-optimal solutions	Increased complexity as number of alternatives increases
Q-L (in planning phase)	Converges in static environments	Increased complexity in dynamic environments, as state-action pairs increase
Q-L (during discourse)	Adapt to opponent's responses, approach pareto-optimal solutions	Unrealistic assumption of opponents' feedback after each action, or difficulty in estimating the Q-value
Repetitive		
CBR (in planning phase)	Save agents from decision making costs in planning	1.The 'routine trap' 2.Maintain and search large case-base 3.Collect and identify domain-specific information to discriminate situations 4.Accuracy decreases as data distributions change
CBR (during discourse)	1. Decision making shortcuts in state transitions, related to concessions 2. Generation of arguments in argumentative negotiations	
Predictive		
Possibilistic CBR (in planning phase)	Estimate expected utility	1.The 'routine trap' 2.Maintain and search large case-base Collect and identify domain-specific information to discriminate situations 3.Accuracy decreases as data distributions change
Bayesian Learning	Estimate opponents' reservation value 1.Estimate Opponents' preference relations 2.Estimate Opponents' payoff structure	1.A-priori knowledge of many probability distributions 2.Models' accuracy reduces in dynamic environments with changing distributions
Non-Linear Regression	1.Estimate Opponents strategic parameters (reservation value, deadline, concession parameter) 2.Estimate Opponents' future offers 3.Withdraw from unprofitable negotiations	Assumes knowledge of function forms
Difference Method	1.Estimate Opponents' future offers 2.Withdraw from unprofitable negotiations	Weight assessment needs improvement, less accurate compared to non-linear regression and neural networks
Neural Networks	1.Estimate Opponents' future offers 2.Withdraw from unprofitable negotiations	Tested in bounded domains