DESIGN AND EVALUATION OF AGENT-BASED NEGOTIATION HEURISTIC FOR ONLINE NEGOTIATIONS

Kaushal Chari, Manish Agrawal
University of South Florida, 4202 E Fowler Avenue, CIS 1040, Tampa, FL 33620-7800, USA

Keywords: Software agents, electronic markets, automated negotiations.

Abstract: This paper presents a negotiation heuristic for software agents that enable agents to learn about the opponent's behavior and use market information while conducting online negotiations. The heuristic is tested in a pilot experimental study, where the performance of agents is evaluated with respect to human negotiators in a simulated electronic market. Preliminary results indicate that agents may have the potential to do better than humans in multi-issue negotiation settings.

1 INTRODUCTION

Software agents (Chari and Seshadri 2004) have the potential to act as effective surrogates of their human principals during automated negotiations due to their ability to overcome the cognitive and information processing limitations of humans in negotiation tasks. Although many research prototypes have been developed to enable software agents to negotiate (see Section 2), to our knowledge, they all suffer from some weaknesses as pointed out in Section 2. In the quest for developing software agents that are more robust and effective for online negotiations, we propose a learning heuristic for software agents, implement this heuristic in a multi-issue negotiation setting (i.e., an electronic marketplace), and evaluate its effectiveness compared to humans using a pilot experimental study.

The context of agent-based negotiations is an electronic marketplace with a finite open time window for completing transactions. Buyer and seller agents register in the e-marketplace and transact a given quantity of an item within the predefined time window. Both buyer and sellers are cognizant of the market values and use this information during negotiations. An agent can negotiate with multiple opponents sequentially to transact the required quantity, since a single seller for example, may not be able to meet the entire demand of a buyer agent. These conditions are similar to trading in a commodity exchange (Chicago Board of Trade 1998). A transaction is completed successfully when an agreement is reached on all issues under consideration such as price, delivery terms, etc. The preference structure of human principals are elicited and stored in their surrogate agents as utility functions. The agents negotiate on behalf of their principal until either an agreement is reached or the time window expires.

2 LITERATURE REVIEW

Business negotiations have been studied from various perspectives including game theory, economic, socio-psychological, and intelligent approaches. Game theoretic approaches make fairly restrictive assumptions about opponent behaviors, thereby rendering them somewhat impractical in real-life negotiation settings (Kraus and Wilkenfeld 1993). Economic approaches (Zeuthen 1968) treat negotiations as a finite sequence of offers and counter offers that could converge to an agreement if an agreement zone exists, deriving subsequent offers based on expectations about the opponent's behavior. Intelligent approaches, based on artificial intelligence and/or statistical techniques, can facilitate learning of an opponent's behavior, provide efficient search of the negotiations solution space for an agreeable solution, and automate the negotiations process. Examples of intelligent approaches include case-based reasoning, heuristic-searches, automated learning, Bayesian techniques,

Chari K. and Agrawal M. (2004).
DESIGN AND EVALUATION OF AGENT-BASED NEGOTIATION HEURISTIC FOR ONLINE NEGOTIATIONS.
In *Proceedings of the Sixth International Conference on Enterprise Information Systems*, pages 92-97
DOI: 10.5220/0002604500920097
Copyright © SciTePress

and genetic algorithms. Case-based approaches, which match previous recorded instances of negotiations from the case history to the current situation (e.g., PERSUADER system (Sycara 1990)), are not effective when existing cases in the case history database do not match the current negotiation situation. Genetic algorithms pit one negotiation strategy against another, and use the outcome to produce improved strategies from subsequent generations in an evolutionary manner (Oliver 1996). However, they often require a very large number of generations to refine the negotiation strategy. Heuristic techniques search a multi-dimensional space for a point that is agreeable to all negotiating entities. Bayesian approaches provide the ability to learn during negotiations using probability update rules (Zeng and Sycara 1998); however, such probabilities are difficult to define ex ante and may sometimes be inaccurate.

Many agent systems have been developed for automated/semi-automated negotiations. One of the pioneering systems is the Kasbah agent system, which uses a simple negotiation heuristics based on pre-defined price decay or increment functions (Chavez and Maes, 1996). Kasbah agents do not learn and therefore do not adapt to the negotiation environment. Agents developed in the Bazaar project (Zeng and Sycara 1998) use Bayesian update rules to learn and form beliefs about the opponent's behavior. As stated before, this approach is limited by the difficulty in assessing various probabilities used in Bayesian update rules. Faratin et al. (1998) use families of polynomial and exponential functions to model opponent concession behaviors during negotiations (e.g., boulware, conceder and imitative behaviors) and combine them using weights to create a negotiation strategy. This approach requires human intervention to assign weights for alternative negotiation strategies and does not provide agents with any learning capabilities. Chari and Bhattacherjee (2002) present a heuristic for agent negotiations that learns from the opponent's behavior. However, this heuristic suffers from unrealistic requirements such as the need for agents to know the market demand supply ratio, and the lack of robustness for various negotiation settings.

The review of the above literature indicates that: 1) there is no research prototype that is intelligent and robust enough to support automated negotiations in real world negotiation settings, and 2) No research has investigated the performance of existing agents with respect to humans in live e-marketplace negotiations. The current research aims to address these limitations by building a learning-based negotiation heuristic that uses market values in determining bids in real time. We also present results of a pilot experimental study that compares the performance of agents with humans.

3 NEGOTIATION HEURISTIC

A negotiation heuristic determines the scheme for making offers/bids (hence forth referred to as simply offers) during negotiations. Negotiations involving multiple issues (such as price, financing rate, delivery term etc) require the two negotiation partners to agree on all the issues. We present a negotiation heuristic that supports multiple issue negotiations. This heuristic uses the utility function as well as the reservation values of various issues while making offers. The utility functions are generated by eliciting preferences from the human principal of an agent. The heuristic learns from the opponent's behavior, uses market conditions in making offers and handles multiple threads sequentially within a limited time window. We make the following assumptions: (a) bilateral negotiations; (b) the negotiators are always in conflict over each issue; (c) the utility value of an offer for any issue never exceeds the utility value of an earlier offer for that issue.

The central idea behind the heuristic is as follows. An agent implementing the heuristic estimates the number of iterations required to reach the market value by estimating the opponent's concession curve by fitting the best curve on the opponent's observed offer points. Using this information as well as information on market values, the agent estimates the target value that it should strive to reach at the last iteration of the negotiations thread. The agent then determines the concession rate to move from its last offer to the target offer for an issue in the remaining iterations and then accordingly makes an offer subject to some constraints. Notations used in the heuristic are given in Table 1.

Before an agent enters into negotiations with an opponent, its human principal provides: (a) a bound on the maximum number of iterations for that thread, $tmax_{is}$, (b) own reservation values, and c) starting bid values for all issues.

Table1: Notations used in negotiation heuristic

Symbol	Description
y_{lk}	Market value of issue *k*, i.e., value of *k* for last the successful transaction in the marketplace.
$\mathbf{x_{ij}}$	$= [x_{ij1},,\ldots,x_{ijn}]$ is the vector of *n* issue values proposed by agent *i* at iteration *j*.
$Tmax_i$	Total time available to agent *i* for negotiations. Typically the length of the time window of the market place.
$u_i(x_1,..,x_n)$	Utility function of agent *i* as a function of *n* issue values.
Q_i	Quantity required to be transacted by agent *i* during $Tmax_i$
$\mathbf{r_i}$	$=[r_{i1},\ldots, r_{in}]$ reservation value vector of agent *i* containing reservation values for *n* issues.
$\boldsymbol{\tau_{ij}}$	$= [\tau_{ij1},\ldots,\tau_{ijn}]$ is the vector of target values of *n* issues for agent *i* at iteration *j*, that agent *i* strives to achieve as the agreed solution.
$\mathbf{cr_{ij}}$	$= [cr_{ij1,\ldots,}cr_{ijn}]$ is the vector of concession rates of agent *i* at iteration *j* across *n* issue dimensions.
$\mathbf{s_{ij}}$	$= [s_{ij1,\ldots,}\ s_{ijn}]$ is the vector of step sizes of agent *i* at iteration *j* across the *n* issue dimensions.
t_{ijk}	Number of additional iterations estimated by agent *i* at iteration *j* for issue *k* in the current thread.
mn_{jk}	Estimate at iteration *j* of the minimum value for issue *k* that the opponent uses in making offers. This value is used in (1).
mx_{jk}	Estimate at iteration *j* of the maximum value for issue *k* that the opponent uses in making offers. This value is used in (1).
$tmax_{is}$	Bound on the maximum number of iterations of agent *i* in thread *s*. $tmax_{is} \geq j + max_k(t_{ijk})$.
q_s	Quantity transacted during negotiation thread *s* when an agreement is reached with the opponent.
I_i	Set of negotiation issues such that the utility value of agent *i* is non-decreasing with respect to increasing issue values. For example, interest free payment period in case of a buyer.
D_i	Set of negotiation issues such that the utility value of agent *i* is non-increasing with respect to increasing issue values. For example, price in case of a buyer.
δ_I	Constant between 0 and 1 that denotes the fraction of the market value that could be reduced from (denoted by index *i* = 1), or added to (denoted by index *i* = 2), the market value in determining the settlement target.
round(x)	Rounds *x* to the nearest value in the domain of *x*.

The opponent's maximum number of iterations for a given thread is estimated by learning from the opponent's behavior. We use the exponential function proposed in (Faratin et. al. 1998) to model the entire range of concession behaviors of the opponent. According to this function, the opponent's offer at iteration *j* for issue *k* can be computed by (1) as follows.

$$x_{ojk} = \{mn_{jk} + \alpha_{jk}.(mx_{jk} - mn_{jk}) \text{ for } k \in D_o \text{ (i.e., } I_i)$$
$$mn_{jk} + (1 - \alpha_{jk})(mx_{jk} - mn_{jk}) \text{ for } k \in I_o \text{ (i.e., } D_i)$$

where

$$\alpha_{jk} = exp((1 - min(j, tmax_k)/tmax_k)\beta_{jk} \ln k_{jk}),$$
$$\beta_{jk} > 0 \text{ and } 0 < k_{jk} < 1 \qquad (1)$$

In (1) the parameter β_{jk} captures the type of negotiations behavior: boulware ($\beta_{jk} < 1$), conceeder ($\beta_{jk} > 1$) etc; k_{jk} is the estimate at the j^{th} iteration for the starting value of α_{jk}. While parameters β_{jk}, t_{maxk}, mn_{jk} (for $k \in I_o$), mx_{jk} (for $k \in D_o$) and k_{jk} can be estimated by fitting a curve through opponent's observed offer points till iteration *j*: $x_{o1k},\ldots, x_{ojk}$, while minimizing the least squared error, only parameter t_{maxk} is used by the heuristic. Note that the reservation value of opponent can be estimated from parameters mn_{jk} (for $k \in I_o$), mx_{jk} (for $k \in D_o$), however the agent uses the market value instead of the opponent's reservation value estimates in making bids. The estimate of the number of additional iterations at iteration *j* to reach an agreement on issue *k* is given by (2). Note that the total number of iterations in a thread is subject to the bound set by the user of the agent.

$$t_{ijk} = max(\ min(\ tmax_k, tmax_{is}) - j, 1) \qquad (2)$$

A target value is computed for each issue based on the market conditions according to (3). This is the value at which agent *i*, strives to reach an agreement.

$$\tau_{ijk} = min(max\ ((1-\delta_1)y_{lk}, r_{ik}), x_{ij}\text{-}1_k) \text{ for } k \in I_i \ ;$$
$$max(min((1+\delta_2)y_{lk},\ r_{ik}), x_{ij}\text{-}1_k) \text{ for } k \in D_i \qquad (3)$$

For example, when $\delta_2 = 0.05$, and issue *k* is price, then target is set to 105% of the current market value for price, subject to bounds set by own reservation price and previous offer made. The target value is approached at a rate given by the concession rate in (4).

$cr_{ijk} = (\tau_{ijk} - x_{ij}\text{-}1_k)/\, t_{ijk}$ (4)

The concession rate is then used to compute the step size for the move from the previous offer as:

$s_{ijk} = cr_{ijk}\, \Delta_j$ (5)

Note that $\Delta_j = 1$. Agent i's offer can then be computed as follows:

x_{ijk} =
$\{max(min\ (round(x_{ij}\text{-}1_k + s_{ijk}),\ x_{ij}\text{-}1_k),\ \tau_{ijk},\ r_{ik},\ x_{ojk})$ for $k \in Ii$;
$min(max(round(x_{ij}\text{-}1_k + s_{ijk}),\ x_{ij}\text{-}1_k),\ \tau_{ijk},\ r_{ik},\ x_{ojk})$ for $k \in D_i$ } (6)

To reduce the computation time for estimating the opponent's concession curve in order to determine $tmax_k$, a limited enumeration can be performed. The range of values for each parameter of the opponent's estimated concession curve to be searched is bounded by opponent's offer and other parameters.

The heuristic for agent *i* is summarized below.

1. Get the weights w_k for the utility function $u_i(x)$ of the human principal.
2. Get value of Q_i
3. Set $s=0$, $t'=0$ and $q'=0$.
4. If$((q' < Q_i) \wedge (t' < Tmax_i)$ then select an opponent with public information (Q_o, x_{o1}) such that $(Q_o/(Q_i\text{-}q'))(u(x_{o1}))$ is the highest across all non-busy opponents, set $s = s+1$, $j=1$; else stop.
5. Get values for the following: r_i, $tmax_{is}$, x_{i1}
6. If ($x_{i1} = x_{o1}$), agreement is reached, stop thread, transact $min(Q_o, (Q_i\text{-}\ q'))$, set $q' = q' + min(Q_o, (Q_i\text{-}\ q'))$, go to Step 4.

Repeat Steps 7- 16 for iteration *j* of thread *s*

7. Set $j = j+1$. Get offer x_{oj} from the opponent.
8. If $(x_{ij}\text{-}1 = x_{oj})$, agreement is reached, stop thread, transact $min\ (Q_o, (Q_i\text{-}\ q'))$, set $q' = q' + min(Q_o, (Q_i\text{-}\ q'))$, go to Step 4.
9. If $(j \geq 4)$ then
   for each k such that $(x_{ij}\text{-}1_k \neq x_{ojk})$, using $x_{o1},..,x_{oj}$, estimate $tmax_k$, β_k, mn_k, mx_k, by fitting a curve of the form given by (1) and minimizing the sum of squared errors. Use user-supplied bounds while enumerating parameters to search for the best curve and then use $tmax_k$ in (2) to compute t_{ijk}.
10. If $(j < 4)$ then for all k set $t_{ijk} = tmax_{is} - j$.
11. If $(t_{ijk} > 0)$ then compute τ_{ijk} using (3), cr_{ijk} using (4), s_{ijk} using (5), and x_{ijk} using (6).
12. For all k such that $(x_{ij}\text{-}1_k = x_{ojk})$ set $x_{ijk} = x_{ij}\text{-}1_k$.
13. If$(x_{ij} = x_{ij}\text{-}1 \neq x_{oj})$
   then for all $k \in I_i$ such that $(x_{ijk} > y_{lk})$
   set $x_{ijk} = max(y_{lk}, r_{ik}, x_{ojk})$
   for $k \in D_i$ such that $(x_{ijk} < y_{lk})$
   set $x_{ijk} = min(y_{lk}, r_{ik}, x_{ojk})$
14. If$(x_{ij} = x_{oj})$, agreement is reached, stop thread, transact $min\ (Q_o, (Q_i\text{-}\ q'))$, set $q' = q' + min(Q_o, (Q_i\text{-}\ q'))$, go to Step 4.
15. If $(((x_{ij} = x_{ij}\text{-}1) \wedge (x_{oj} = x_{oj}\text{-}1))$ then
   if $(\forall k(\ (x_{ojk} \geq r_{ijk})$ when $k \in I_i$ and $(x_{ojk} \leq r_{ijk})$ when $k \in D_i$ then $x_{ij} = x_{oj}$ and reach an agreement,
   else stop thread, no agreement reached, go to Step 4.
16. If$(((x_{ij} = x_{ij}\text{-}1 = r_i) \vee (j = tmax_{is})) \wedge ((Q_i\text{-}\ q')/ Q_i)/((Tmax_i - t')/ Tmax_i) > 1)$ then
   if $(\forall k(\ (x_{ojk} \geq r_{ijk})$ when $k \in I_i$ and $(x_{ojk} \leq r_{ijk})$ when $k \in D_i$ then $x_{ij} = x_{oj}$ and reach an agreement;
   else stop thread, no agreement reached, go to Step 4.
   Else stop thread, no agreement reached, go to Step 4.

In Step 1, the weights w_k assigned to the negotiation issues (i.e., price and period in the current paper) are obtained from the human principal and then incorporated in a commonly used additive utility function of the form $u(price, period) = w_{price}\ u_1(price) + w_{period}\ u_2(period)$, where u_1 and u_2 are normalized linear functions of *price* and *period* respectively. The quantity to be transacted by the agent in the market place is then specified to the agent in Step 2. In Step 3, the time counter t' is started. The value of t' is constantly updated by a clock. The thread count, s as well as the quantity transacted, q' are also initialized to zero in Step 3. In Step 4, an available opponent is selected for negotiations with the highest value for the metric $(Q_o/(Q_i\text{-}q'))(u(x_{o1}))$ which is the product of the ratio of opponent's quantity and the quantity remaining to be transacted, and the utility value of the starting bid of the opponent. This metric enables the agents to select an opponent in the pool, who has large quantity to transact as well as an attractive starting bid that gives high utility value to the agent. In Step 4, an opponent is only selected if some quantity still remains to be transacted (i.e., $q' < Q_i$) and the time window has not expired ($t' < Tmax_i$). In Step 5, the agent obtains its human principal's reservation values along with the bound on the number of iterations for negotiations for the current negotiation thread as well as the starting bid from its human principal. Step 6 is needed to check if an agreement is reached in the very first iteration.

When iterations are four or higher, then offer points available from the opponent are adequate to run a curve fitting procedure in order to estimate the number of iterations the opponent is targeting. In Step 9, the curve-fitting procedure is run based on the exponential curve in (1) to estimate $tmax_k$. When the iterations are less than four, the number of offer points of the opponents is not adequate to compute good estimates of $tmax_k$. In this case, the total number of iterations is estimated as $tmax_{is}$, the initial bound that is set by the human principal of the agent. The number of iterations remaining is $tmax_{is} – j$, where j is the number of iterations that has elapsed. As long as the number of iterations remaining is one or more, settlement target, concession rate, step size and the new bid is computed in Step 11. For negotiation issues for which agreements have been reached, the last bid value is the current bid value as seen in Step 12.

If the current bid is the same as the last bid and is not equal to the opponent's bid (Step 13), then the current bid value is set to the market value subject to the bounds set by own reservation value and opponent's bid. If now an agreement is reached (Step 14), then quantity can be transacted. If the current bid is the same as the previous bid and the opponent's current bid is also the same as his/her previous bid, then an agreement can be reached if the reservation value constraints are satisfied (Step 15). Finally in Step 16, if the current bid is the same as the previous bid and equals own reservation values, or if the total number of iterations planned $tmax_{is}$, is exhausted, and there is a sense of urgency as given by the metric: $((Q_i - q')/ Q_i)/((Tmax_i - t')/ Tmax_i)$, when its value is greater than 1, then an agreement with the opponent can be reached if the own reservation value constraints are met. Otherwise negotiations are terminated.

4 THE RESEARCH HYPOTHESIS AND EMPIRICAL STUDY

Both human and agent buyers constantly learn from and adapt to each other's behaviors and/or engage in strategic moves in response to opponent behavior. We however conjecture that automated agents implementing our heuristic are likely to have an upper hand in this negotiation process, by virtue of their ability to quickly and accurately estimate uncertain negotiation parameter such as the number of opponent moves. Estimation of such parameter often places substantial cognitive demands on humans. The performance gap between humans and agents will tend to magnify with increasing complexity of the negotiation process, such as the number of negotiation issues. Hence we hypothesize:

H1: Electronic agents will perform significantly better than human negotiators when negotiations involve multiple issues

To test the above hypothesis and the performance of agents, we conducted a pilot experiment to identify the differences between performance and efficiencies achieved by humans and electronic agents while trying to buy fixed quantities of goods. The experiments involved buyers and sellers negotiating over two issues: price and the number of months of interest free payment period. The dependent variable in the experiments was negotiation performance, i.e., the utility gained by the settlement over the utility of own reservation values. The objective of each negotiator was to buy or sell at values that maximize his/her utility.

To make the negotiation environment as realistic as possible, we created experiments similar to a commodities exchange (Chicago Board of Trade 1998). Specifically, to model price discovery, the price of the last trade was displayed to all buyers along with the financing period. All sellers were electronic and used a hybrid Boulware/Conceder algorithm (different from the current heuristic) to make offers. Seller agents used negotiation parameters based on current market values. Human subjects played the role of buyers.

To conduct the experiments, we sought subjects with prior experience or coursework on negotiations. Subjects received token cash incentives based on their negotiation performance based on their utility metric. We elicited weights for utility functions from human subjects. The utility functions used were the commonly used linear additive function of the form presented in Section 3. For the actual negotiation sessions, all subjects were assigned identical reservation values (80 for price and 3 months for period) and were required to buy identical quantities (20 units) of the commodity. Parameters δ_1 and δ_2 were set to 0.05. Experiments were conducted under two market configurations. In the first set of negotiations, market supply was half the total demand. In the second set, supply was twice the actual demand. To make the most efficient use of subjects' time, two rounds of negotiations were conducted for each market configuration. Half the subjects in each round used their surrogate agents implementing the heuristic presented in this paper, and the other half negotiated on their own. In the second round, subjects that used agents in the first round negotiated without agents, while subjects that did not use agents, now used surrogate agents in the second round to buy in the market place. We had eight subjects for the experiments and the supplies were appropriately calibrated for the two market configurations.

The results for successful transactions during the experiments are shown in tables 2 and 3. Specifically, Table 3 contains results from t-tests for differences. As can be seen from Table 3, fewer transactions were made when supplies were limited.

Table 2: Summary Statistics

		Demand/Supply= 2			Demand/Supply = 0.5		
	TYPE	**N**	**Mean**	**Std. Deviation**	**N**	**Mean**	**Std. Deviation**
Close price	Agent	4	74.1700	3.28996	9	73.0667	4.35895
	Human	9	75.7667	5.13323	7	72.7814	4.32834
Close period	Agent	4	4.5425	.50704	9	4.4344	.89605
	Human	9	3.7733	.61640	7	3.4571	.40766
Iterations	Agent	4	7.50	3.697	9	4.44	1.236
	Human	9	9.44	2.128	7	5.14	1.345

Table 3: t-test results

	Demand/Supply =2				Demand/Supply = 0.5			
Variable	**t**	**Df**	**Sig (2-tailed)**	**mean diff**	**T**	**Df**	**Sig (2-tailed)**	**mean diff**
Close price	-0.67	9	0.52	-1.60	0.13	13	0.9	0.29
Close period	2.36	7	0.05	0.77	2.91	11	0.01	0.98
Iterations	-0.98	4	0.38	-1.94	-1.07	12	0.31	-0.7

The results indicate that there were no significant differences between humans and agents in prices. However, agents performed significantly better than humans on negotiating the duration of the financing period. The difference was significant across market conditions, supporting Hypothesis 1 regarding the superiority of agents in multi-issue negotiations. This result also suggests that agents should be preferred when negotiations involve multiple issues.

An obvious limitation of these results is the small sample size of the pilot. In our experiments following this pilot, we will address this limitation by performing more experiments with human subjects. Also, we will vary the number of issues to ascertain the performance of agents with respect to humans as the number of issues change. Some subjects were uncomfortable with the monotonic nature of the negotiations and suggested that they be allowed to lower their offer prices in return for conceding on period. We also plan to improve the human interface based on subject feedback and refine the heuristic further based on the experimental results.

REFERENCES

Chari, K., & Seshadri, S. (2004). Next Generation B2B Commerce Using Software Agents. *Intelligent Enterprises of the 21st Century*, Idea Group Publishing, Hershey, PA 17033.

Chari, K., & Bhattacherjee, A. (2002, December). Multi-Thread Automated Negotiations Using Learning Agents. *Proceedings of the Workshop on Information Technologies and Systems (WITS)*, Barcelona, Spain.

Chavez, A., & Maes, P. (1996, April). Kasbah: An agent marketplace for buying and selling goods. *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London.

Chicago Board of Trade (1998), *Commodity Trading Manual*, Chicago, IL.

Faratin, P., Sierra, C., & Jennings, N. R. (1998). Negotiation Decision Functions for Autonomous Agents. *Robotics and Autonomous Systems*, 24 (3-4): 159-182.

Kraus, S., & Wilkenfeld, J. (1993). A Strategic Negotiation Model with Applications to an International, Crisis. *IEEE Transaction on Systems Man and Cybernetics*, 23(1), 313-323.

Oliver J.R. (1996, Winter). A Machine Learning Approach to Automated Negotiation and Prospects for Electronic Commerce, *Journal of Management Information Systems*, 13, 3 126-164.

Sycara K., (1990). Negotiation Planning: An AI approach. *European Journal of Operational Research*, 46, 216-234.

Zeng D., & Sycara, K. (1998). Bayesian Learning in Negotiation. *International Journal of Human-Computer Studies*, 48, pp. 125-141.

Zeuthen, F. (1968). *Problems of Monopoly and Economic Warfare*, Routledge and Kagan Paul, London.