

Short version appeared in:

Proceedings of the Group Decision and Negotiation 2006
S. Seifert and Ch. Weinhardt (Eds.), Universitätsverlag Karlsruhe, 2006, (pp. 166-167)

Predicting Opponent's Moves in Electronic Negotiations Using Neural Networks [♦]

Réal Carbonneau¹, Gregory E. Kersten² and Rustam Vahidov²

¹ Department of Management Sciences, HEC Montréal
Montréal, Canada

² John Molson School of Business, Concordia University
Montreal, Canada

Abstract

Electronic negotiation experiments provide a rich source of information about relationships between the negotiators, their individual actions, and the negotiation dynamics. This information can be effectively utilized by intelligent agents equipped with adaptive capabilities to learn from past negotiations and assist in selecting appropriate negotiation tactics. This paper presents an approach to modeling the negotiation process in a time-series fashion using artificial neural network. In essence, the network uses information about past offers and the current proposed offer to simulate expected counter-offers. On the basis of the model's prediction, "what-if" analysis of counter-offers can be done with the purpose of optimizing the current offer. The neural network has been trained using the Levenberg-Marquardt algorithm with Bayesian Regularization. The simulation of the predictive model on a testing set has very good and highly significant performance. The findings suggest that machine learning techniques may find useful applications in the context of electronic negotiations. These techniques can be effectively incorporated in an intelligent agent that can sense the environment and assist negotiators by providing predictive information, and possibly automating some negotiation steps.

[♦] This work has been partially supported with grants from the Initiative for New Economy of the Natural Sciences and Engineering Research Council and the Social Sciences and Humanities Research Council, Canada. An earlier version of the paper has been presented at the.

1. Introduction

Negotiations play a crucial role in conducting everyday business activities, such as negotiating contracts with customers, negotiating service level agreements with suppliers or negotiating agreements with unions. Many of these negotiations can have outcomes that impact long-term business relationships, profitability, and reputation of businesses. Electronic negotiations in particular, have gained heightened importance due to the advance of the web and e-commerce (Gregory E. Kersten & Noronha, 1999). While this brings in the challenges associated with conducting negotiations in the global environment, where parties could have little or no knowledge of each other, it also presents some valuable opportunities to employ advanced technologies, such as intelligent agents, in the negotiation process.

Intelligent systems for negotiation support that aim at enhancing the negotiator's abilities to understand the counterparts, their needs and limitations and to predict their moves could be very valuable tools to be used in negotiation tasks (Zeng & Sycara, 1998). The purpose of this paper is to investigate the feasibility of machine learning approaches in modeling the opponent's future offers. The data is obtained from electronic negotiation experiments which provide a rich source of information about the relationships between negotiators, their individual actions, and the negotiation dynamics. This information can be effectively utilized by intelligent agents equipped with adaptive capabilities to learn from past negotiations and assist in selecting appropriate negotiation tactics.

In order to test the applicability of machine learning approaches in modeling negotiation dynamics we use data obtained from bilateral negotiations experiments conducted with the use of the Inspire electronic negotiation system (G.E. Kersten & Lo, 2003; Gregory E. Kersten & Noronha, 1999). In our approach the negotiation process has been modeled in a time-series fashion using an artificial neural network. In essence, the model uses information about past offers and counteroffers, including the most recent offer made by the negotiator, to predict the expected counteroffer. On the basis of the model's prediction, "what-if" analysis of counter-offers can be done with the purpose of optimizing the current offer. The assessment of offers and counter-offers is performed based on the user's utility function.

The purpose of this study is to assess the applicability of machine learning to provide advice to the negotiator. This advice involves simulation of the possible responses to the offer negotiator is contemplating. The subsequent sections present the related work; introduce a neural network-based approach to model the negotiator's counterpart; propose negotiation support which involves offer optimization; discuss the model implementation and results; and provide conclusions and directions for future work.

2. Background

Negotiation is one of key activities of businesses and it has crucial impact on an organization's performance. Researchers in negotiation support and automation seek to facilitate various negotiation-related tasks using the capabilities of technology. Substantial efforts have been expended in attempts to fully automate negotiation processes (Beam & Segev, 1996; A. Chavez, D. Dreilinger, R. Guttman, & Maes P., 1997; Jennings & Faratin, 2001; Maes, Guttman, & Moukas, 1998). While the extensive coverage of automated negotiations is beyond the scope of the

current work, it seems useful to mention several key studies.

Another work employed genetic algorithms to generate rules which relate the negotiators' current offers with the likely subsequent offers (Matwin, Szapiro, & Haigh, 1991). The algorithm evolves multiple classifiers with the higher fitness being assigned to those rules, which more frequently contribute towards "compromise trajectories". The method's limitations include, not taking into account relative importance of various issues and values; outguessing aspiration levels of the counterpart; compromise-orientation rather than orientation on the achievement of values of the supported negotiator; and the substantial increase of the search space size when additional issues and continuous issues are introduced.

The genetic algorithm proposed by Matwin, Szapiro, & Haigh (1991) has never been empirically verified. Chavez and Maes (1996) designed and tested an agent-based marketplace Kasbah in which various agents could be created by the users. These agents engage in bilateral negotiations on behalf of their principals (buyers and sellers). They followed one of the three negotiation strategies defined by price-concession curve over time (Chavez, Dreilinger, Guttman, & Maes, 1997).

More recently, Kwon, Shin, & Kim (2006) used a semantic web-based agent community to introduce the concept of pervasive negotiation support. The system, called "SmartGuide" includes user, supplier, and negotiation agents to provide flexible environment for context-aware automated negotiations.

Dzeng and Lin (2004) propose an agent-based system for procurement negotiations in construction domain. The system employs the coordinator agent that tries to find adequate agreement for the parties while employing genetic algorithms and using the preferences of the parties. While revealing negotiator preferences to a coordinating agent simplifies the task, this may not be appropriate for most business negotiations, where the participants would typically like to keep their preferences private. The experimental results show that the joint payoff of the contractor and supplier improved from 1.5% to 9.8% as compared with conventional human negotiation (Dzeng & Lin, 2005).

An approach that combines recommendation generation and negotiation agents for product selection and purchase is described in (Lee, 2004). The negotiating agents in this work follow a chosen concession strategy while taking into account past offers by the opponent. More specifically, the agents base their concessions according to the "positive" vs. "negative" classification of an opponent's attitude, which is determined using a threshold function.

While the above and numerous other works in automated negotiations may hold a promise in streamlining routine well-structured negotiation tasks, we believe that in most business negotiation contexts humans need to be in control of the process with the agents playing a role of assistants. There has been some research effort in this area to provide solutions for assisting human negotiators. An overview of electronic negotiation and negotiation support systems and negotiation software agents (NSA) is presented in (Kersten & Lo, 2003). The work also discusses the Aspire system which is a combination of the Inspire NSS and the Atin NSA for assisting negotiators. In Aspire the agent uses an inference engine to provide recommendations based on inputs and previously encoded rules.

Another application of agent assistant in commerce negotiations has been implemented in eAgora marketplace (Chen, Kersten, & Vahidov, 2004; Chen, Vahidov, & Kersten, 2005). In

eAgora an agent watches over the shoulder of the negotiator and critiques trial offers. The agent also advises an action upon receiving the counterpart's offer and generates a package of adequate candidate offers for the consideration by the user. An experimental evaluation of eAgora's has revealed that agent-assisted negotiators have been able to achieve better performance and reported higher usefulness and satisfaction levels, especially when completing complex negotiation cases (Vahidov, Chen, & Feng, 2005). These findings encourage further work on using innovative agent solutions in e-negotiations.

In the Aspire and eAgora systems the agents rely on a knowledge base created by the developers. While providing users with help, based on common negotiation related heuristics and rules of thumb is an improvement over unassisted mode, the accuracy of the knowledge and its applicability to concrete cases may be questionable. Machine learning approaches do not require explicit specification of the knowledge; instead they are capable of learning patterns from the existing data.

This paper sets out to investigate the applicability of neural networks to learning the dynamics of negotiations from the past negotiation cases. Being "universal approximators", neural networks are capable of learning and generalization, thus offering an inductive alternative to the expert system-like support. Oprea (2002) proposes a similar approach in which a feed-forward neural nets are used for predicting the counterpart's next offer. The approach is appropriate to a single issue (price) negotiation and only the offers made by an opponent are considered in the model.

There are many sources of information that may be used by a negotiator, which could be categorized as offer and non-offer information. Offer information is self contained in an offer and their sequence, which is always available in any negotiation situation. Non-offer information, such as information from the environment, persuasion tactics and discussions, may be present in negotiation situations, but are not a defining feature. This research focuses on information contained in the offers because this is the information that is fundamental to communication required in any negotiation. In this research we construct a neural network-based model which learns patterns from the past offer exchanges and presents recommendations for future moves.

3. Neural network-based predictive model

Modeling of an opponent in the negotiation process may significantly improve performance of the negotiators. Some of the works mentioned above attempt to incorporate the opponent's moves in the process of offer generation. For example, in (Lee, 2004) past concessions made by the counterpart are used to construct the model of this counterpart. If, on the average, they exceed a pre-defined threshold level, the opponent is modeled as having a "positive" attitude. Zeng and Sycara (1998) use game-theoretic approach with Bayesian belief revision to model a negotiation counterpart. Mudgal and Vassileva (2000) represent the counterpart's decision-making with probabilistic influence diagrams. In (Faratin, Sierra, & Jennings, 2002) an approach to generating offers that are close to the opponent's offer based on a fuzzy similarity measure is employed. In eAgora system an assistant agent generates recommendations based on the chosen strategy, as well as the concessions made by an opponent (Chen et al., 2005). While such approaches provide a certain advantage, they are in some sense "rough" in providing ways to predict next moves by the opponent.

Anticipation of the counterpart's next move can be of critical importance to negotiators as it could help them to better compose their own offer. Therefore, negotiators make efforts to uncover patterns in their counterpart's behavior, in particular the concession making patterns. The feasibility of a predictive model to perform this task can be assessed utilizing advanced machine learning techniques that are effective in pattern detection. In this work we have relied on artificial neural networks (ANN), which have been proven to be universal approximators, provided with sufficient hidden layer neurons and assuming that the activation function is bounded and non-constant (Hornik, 1991).

The predictive model will need to indicate the expected offer from the counterpart. Thus, at the output layer of a neural network each neuron will be associated with one component of the offer, i.e. one negotiation issue. The number of output neurons will be essentially the same as the number of the issues in negotiations. The inputs will be associated with the past offers and counter-offers and the current trial offer. Therefore, the same number of inputs as that of outputs will be allocated to the current considered offer plus the history of past offers and counter-offers. In general, the more of the history of a given negotiation is fed as inputs the more precise one would expect the predictions to be. However, this effectively puts a restriction on when the model can be actually used in the negotiation process. For example, when two latest offers and counteroffers are included, the negotiator may start using the model only in the third round of negotiations.

Our predictive network needs to be trained using past negotiation data to adequately capture the dynamics of negotiations. Although the classical error back-propagation algorithm has been the most popular learning techniques for neural networks, we use a faster training algorithm as well as a framework that will improve the generalization capability of the model. In particular, we use the Levenberg-Marquardt algorithm as applied to neural networks to adjust the weights (Hagan, Demuth, & Beale, 1996; Hagan & Menhaj, 1994). Our choice is due to the fact that this algorithm is one of the fastest training algorithms available with training being 10-100 times faster than simple gradient descent back-propagation of error.

The Levenberg-Marquardt neural network training algorithm is combined into a framework that permits estimation of the network's generalization by the use of a regularization parameter. ANN performance measures typically include the error of the outputs of the network, such as the means squared error (MSE). However, with the large number of hidden units (and thus, weights) neural networks may be too powerful, resulting in overtraining. While improving on MSE, such ANN may end up with inferior generalization capability. Thus, a regularization performance function which includes the sum of the weights and biases can be used instead, combined with a regularization parameter, which determines how much weight is given to the sum of weights and bias in the formula:

$$msereg = \gamma mse + (1 - \gamma) msw.$$

MSE represents the mean of the sum of squares of errors, MSW represents the mean of the sum of squares of the network weights and biases and γ represents the tradeoff ratio. This regularization parameter permits the control of the ratio of impact between reducing the error of the network and the number of weights or power of the network.

The tuning of regularization parameter is automated within the Bayesian framework (MacKay, 1992) and, as combined with the Levenberg-Marquardt training algorithm, results in high

performance training combined with a preservation of generalization by avoiding overfitting of the training data (Foresee & Hagan, 1997). This algorithm helps control overfitting of the target function and it also provides an estimate of how many weights and biases are being effectively used by the network. Larger networks should result in approximately the same performance, since regularization results in a trade off between error and network parameters, which is relatively independent of network size. All Neural Network modeling and training in this work is performed in MATLAB 7.0 and MATLAB's Neural Network Toolbox (Mathworks 2005).

4. Negotiation support with predictive neural model

The predictive neural network-based model introduced in the previous section can be integrated as part of the negotiation support system. One concern with such integration relates to the informational demand required by the model. Ideally, the model should not require extensive information to facilitate the "plug-in" type of integration. Thus, when building the model, we kept a strong focus on using only information that would normally be available to the negotiator.

The most common type of information available to a negotiator any negotiation session include details of an offer and the time it was created. The system monitors the creation of this information and keeps a history, which will be used for predicting the next negotiation steps. Additionally, this permits the negotiation modeling researcher to build models with historical negotiation information from the existing systems; no details on the individual negotiators are necessary. An attractive feature of this model is that no information on the negotiators themselves is required. This allows the intelligent negotiation agent to be built separately from any negotiation system.

The intelligent negotiation assistant discussed here is capable of simulating the opposing negotiator's next offer starting with the third offer. This permits a negotiator to do what-if analysis in order to test various estimated counter offers that will result from a specific offer in the current situation without actually submitting an offer to the counter-part. Apart from using the model as means of learning more about the opponent, this capability would also allow to perform an optimization of all potential offers with the purpose of finding the best offer for the current situation.

The choice of the method of search for the optimal offer depends on the number of issues and options available. For a small number of issues and options a comprehensive search can be performed. For example, if there are two issues and five options, the number of all possible offers is limited to 25. For such small spaces the use of comprehensive search could be justified. However, as the number of possible offers increases, there may be a need for the use of a heuristic optimization method. Such methods could include, for example, hill-climbing, simulated annealing, or genetic algorithms. The agent could choose the appropriate algorithm based on the size of the negotiation space.

5. Model implementation

In order to demonstrate the feasibility and effectiveness of an ANN-based predictive model we have used past data collected by Inspire system (Gregory E. Kersten & Noronha, 1999). The Inspire negotiation system is web-based, which permits two parties located anywhere in the

world, who have internet access, to negotiate on a chosen case. Negotiation issues and issue options are specified in advance for a specific case and each negotiator specifies a rating for each issue and, additionally, specifies a rating for each option of an issue. The sum of all issues must be 100 and the sum of all the options of a specific issue must also be 100. Each selected option rating is multiplied by the issue rating which permits the calculation of the total utility of a package. Package utility ratings are presented to the user for further adjustments since the user may feel that certain package utility ratings should be different than the result of the calculations to correctly reflect his or her preferences. Issue, option and package ratings are specific and confidential to each user, allowing evaluation of all offers submitted or received by a particular user. A negotiator has no information about the preference structure of the counter-part and at no point are these preferences ever revealed to the counterpart. As an additional information source, users may examine a graph of the utility of the history of offers and counter offers.

The negotiation case under study is a simulated scenario where a seller and a buyer want to enter into a business relationship. Specifically, Itex manufacturing is entering the negotiations as a producer of bicycle gears, wishing to sell to Cypress Cycles; a bicycle producer. The case defines the market as competitive, meaning that either party may terminate the negotiations if they do not find the negotiations promising, since they can find other business partners. The four negotiated issues and their options (in brackets) are:

1. Price (3.47\$, 3.71\$, 3.98\$, 4.12\$, 4.37\$);
2. Delivery (20 Days, 30 Days, 45 Days, 60 Days);
3. Payment (60 Days After Delivery, 30 Days After Delivery, Upon Delivery); and
4. Returns (Full Price, 75% Refund with 5% Spoilage, 75% Refund with 10% Spoilage).

Users may send offers and messages to their counterparts through the Inspire system as they work towards a compromise that is acceptable for both parties. At any point users can unilaterally terminate the negotiations.

The Inspire's dataset provides rich information on the negotiation that took place through the use of the system. Of the 6310 offers considered, 2426 (38%) were proposed by females and 3379 (54%) were proposed by males and 505 (8%) offers were from negotiators who did not specify their gender. Negotiators from over 100 different countries are in the sample.

The median negotiation session has 7 offers and the average negotiation session has 6.70 offers. This means that there is enough information to model counter-offers based on information about past offers. At a minimum, we will require the current offer and the next offer, so there is a minimum of two offers in a negotiation session for it to be considered for this type of modeling. To fulfill this minimum requirement, we must also be able to correctly identify the relationship between the offers. We need to know to which offer a counter offer was made, and conversely, to which counter offer an offer belongs. This may seem trivial if offers are forced to be sequential and alternating between the sides, however, in our current data set, one party of the negotiations may make several offers before receiving a counter offer. To reconstruct the correct order relationship from the existing data, we assume that a counter offer is made in relationship to the last offer from the negotiation counterpart.

Adding information about the last buyer's offer and the last seller's offer to the model, which

already contains the current offer and the next offer, provides a model that contains information about past offers. However, it takes a minimum of 4 offers in a negotiation session to learn the patterns and a minimum of 2 past offers plus the current offer, for a total of 3 offers, to estimate the subsequent counter offer. This will result in losing about 27% of the observations because of this time requirement.

The last offers by the buyer and seller are relevant for estimating the subsequent counter offer. The first offer is of interest to us because it sets opening tone of the negotiation as well as the minimum and maximum offers so far, which provides a measure of the absolute range of the current negotiations. To provide information on how much disagreement there is up to date in the current negotiation, we add the standard deviation of the offer values. Finally, the average of the offers will provide information about the mid-point towards which the negotiation is moving. The major inputs and outputs of the model are summarized in Table 1. One extra input is reserved to indicate whether a negotiator is a buyer or a seller.

Table 1 Summary of the model inputs and outputs

Offers	Price	Delivery	Payment	Returns	Timestamp
Inputs					
Last "Buy"	✓	✓	✓	✓	✓
Last "Sell"	✓	✓	✓	✓	✓
Current	✓	✓	✓	✓	✓
First	✓	✓	✓	✓	✓
Minimum	✓	✓	✓	✓	
Maximum	✓	✓	✓	✓	
St. deviation	✓	✓	✓	✓	✓
Average	✓	✓	✓	✓	✓
Outputs					
Predicted	✓	✓	✓	✓	

Thus, the resulting model has 39 inputs and 4 outputs and the final dataset for the above required observations results in 6310 observations. In order to test the generalizability of the model we created two separate sets, 5048 (80%) of the observations have been assigned to the training set and 1262 (20%) of the observations have been kept in the testing set.

Our neural network has 39 inputs, 4 outputs and 10 hidden neurons (Figure 1). The transfer function we used in the hidden layer is the tan-sigmoid function which does non-linear scaling from an infinite range to values between -1 and 1 and the output layer transfer function is a linear function. The input for each neuron is the not only the input signals coming from either the input variables or the results of a previous neural network layer, but it is this signal multiplied by the weight of the connection. There is a total of 444 connection weights, which results in a ratio of observations to weights equal to 11.37 (5048/444), which is an appropriate level to permit powerful modeling and good generalization.

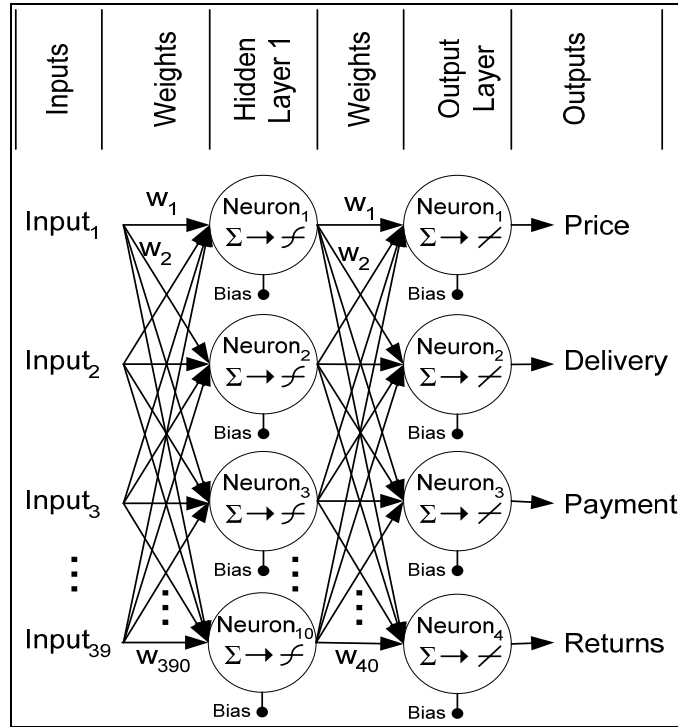


Figure 1. Predictive negotiation neural network model

6. Results

Training of the neural network was executed and the networks converged after 252 iterations through all of the dataset. Convergence is characterized by the stabilization of the sum of square errors (SSE) and the sum of square weights (SSW). For our network of 39 inputs, 10 hidden layer neurons and 4 outputs, we reach a stable SSE of 3642.17, a SSW of 39.5911 and an effective number of network parameters (weight and biases) of 382.143 of the total 444 parameters (Figure 2).

The correlation of the model’s results with the actual output on the testing set are 0.74647 for price, 0.68512 for delivery 0.67421 for payment and 0.69253 for returns which is an average of a 0.6996 correlation of predicted output with actual output. From this we can see that price is the negotiation issue that has the most predictable output based on past negotiation patterns, while delivery has the least predictable pattern.

To further understand the performance of the model we consider the results in terms of the average absolute error and the error obtained after the outputs are rounded to an integer value (rounding is necessary because the negotiation case requires selection of ordinal values). The average absolute error across all outputs of the testing set for one prediction is 1.68 on the 15 possible ordinal levels, which is approximately 11%. When the outputs are rounded to integers, the rounded error is 1.39 for the total of 15 possible ordinal levels which represents approximately a 9% error. The detailed results are presented in Table 2.




Figure 2 - Neural Network training history graph

Table 2. Summary of results

Measure	Price	Payment	Delivery	Returns	Avg/Tot
Testing Set					
R	0.74647	0.68512	0.67421	0.69253	0.6996
R2	0.55722	0.46939	0.45456	0.47960	0.4902
Exact error	0.4753	0.4454	0.3609	0.3972	1.6788
Rounded error	0.4113	0.3796	0.2837	0.3170	1.3914
Training Set					
R	0.76405	0.71638	0.70387	0.7336	0.7295
R2	0.58377	0.51320	0.49543	0.53817	0.5326

To assess the model, several “what-if” scenarios were conducted, a few of them that may be interesting to the seller (in the previously presented negotiation situation) are discussed here. These scenarios provide the seller with information about the expected results of the offers thus helping make a better decision.

The seller may want to see what would the expected counter offer of multiple different offer possibilities be (e.g. resending the same offer as the last time, giving a price concession and giving a concession for issue Returns). Because of the complexity of the option and package ratings in Inspire, we choose a simple package utility rating scheme, where price is 45%,

delivery is 20%, payment is 10% and returns is 25% of the overall utility. Every issue option rating is related to the levels, for example, price has 5 levels, so level 1 would have utility of 0 and level 5 would have utility of 100. From this we can calculate the utilities for any offer which will permit us to evaluate counter offers. As a base for comparison, we know the actual offer made, and we have included this in the comparison of the what-if scenarios in Table 3.

Table 3 Comparison of example “what-if” scenarios

Name	Current Offer				Simulated counteroffer					
	Pr.	Del.	Pay.	Ret.	Pr.	Del.	Pay.	Ret.	Estimated utility	Ordinal utility
Actual	4	3	3	3	2	2	2	1	20.611	18.75
Same as last	5	2	2	3	3	1	1	1	22.138	22.5
Price concession	4	2	2	3	2	1	2	1	19.577	13.75
Returns concession	5	2	2	2	3	1	2	1	22.404	25

Based on these simple “what-if” simulations, we observe that even small variations in the current offer can have important impact on the expected counter offer from the opponent. Using the utility function, we can evaluate different scenarios to better understand the negotiation dynamics and select one that has a high-utility expected counter-offer.

The case used for this research has a total search space of 180 options. Since the search space is small a comprehensive search for an optimal solution is feasible. As mentioned earlier, for larger spaces one of the heuristic optimization methods could be used. Table 4 shows the results of optimization based on the expected counter-offers.

Table 4 Offer optimization results

Current offer					Expected counteroffer					Ordinal utility	Direct utility
Pr.	Del.	Pay.	Ret.	Utility	Pr.	Del.	Pay.	Ret.			
5	4	3	1	65.00	2.5543	1.9644	1.5168	0.91199	30	23.04981	
5	3	2	2	63.75	2.5332	1.6841	1.503	1.0449	30	22.20713	
5	4	2	1	62.50	2.5151	1.9311	1.5415	0.8547	30	22.146	
5	3	3	1	60.00	2.6103	1.7381	1.6238	0.92875	30	22.92056	
5	3	2	1	57.50	2.5929	1.7339	1.6303	0.88159	30	22.42531	
5	2	3	1	55.00	2.6704	1.5069	1.7312	0.94722	30	22.82463	
5	2	2	1	52.50	2.6747	1.5318	1.7095	0.90945	30	22.70719	

The maximum expected ordinal level utility of the counter offer was identified as 30. To achieve this high counter offer, the intelligent agent recommends that the seller changes his offer back to the highest price but at the same time suggests that the seller offer a large returns concession, which is, in most suggested offers, the most generous returns policy level even though it is the issue that has the second highest impact on utility. The other two negotiation issues have more flexibility in offer level ranges before having an impact on the next counter offer.

It is important to note that these recommendations are specific to the current negotiations patterns identified in this specific scenario, they are not general recommendations; the generalized knowledge is embedded in the weights of the neural network. For example, the recommendation to move to a higher price but give concessions on returns is not a general rule, but a result of the patterns identified in the current scenario. Even though the above offers result in the same ordinal counter offer utility level, our agent provides an overall recommendation of the package that maximizes the utility of the current offer, subject to the maximizing of the counter offer utility. Thus the intelligent negotiation assistant agent recommends that the seller offer a price of 4.37\$, 60 days delivery, payment upon delivery and full price returns. Maximizing the current offer subject to maximizing the expected counteroffer means that the first offer of a party on any side of the negotiations has the highest utility since the expected counter offer is not modeled. This means that the starting offer for buyers would be the lowest price, fastest delivery, latest payment and best returns policy; the inverse offer would be submitted by the seller.

The negotiation-assistant agent which uses the proposed approach is capable of providing the seller with information about the dynamics of the current negotiation situation, the best potential offer and the ability to run “what-if” analysis for any considered offer. Thus the seller may directly propose the recommended offer or do further analysis and refine the offer to meet requirements that may have not been thoroughly captured by the utility function or other external factors.

7. Summary and Conclusions

In this paper we have presented a neural networked-based model for predicting the opponent’s offers during negotiation process. The model can be embedded in a system for assisting negotiator in making offers in e-negotiations. The simulation of the predictive model on a testing set has very good and highly significant performance, especially considering the noisy data domain. An examination of “what-if” scenarios and optimization results on a real case shows that the model can exhibit interesting negotiation strategies and, at the very least, provides useful information to a negotiator.

One potential limitation of the current research relates to the generalizability of the model. Specifically, our model was used for a particular negotiation case, and the accuracy of its predictions and recommendations may be less adequate for other types of cases. While this is an important issue to tackle in future research, it is important to note that many e-commerce related negotiation contexts share similar aspects. For example, issues like price, delivery, and returns are common for many real-life commerce negotiations. Thus, to ensure the generalizability of the model, the future research should look into different classes of negotiation cases and recommend a variety of neural-based models as part of the negotiation support toolbox, where the appropriate model could be chosen for particular contexts.

Another important direction for future research includes investigation of the effectiveness of the model in experimental settings. There are two potential experiments that would be interesting in this respect: one featuring fully automated negotiations and the other one involving assisted negotiations. In the first experiment, the human subjects would be assigned to one side of the negotiation case and the opponent would be the negotiation agent incorporating models described in this research. We could then compare the utility of the negotiations outcomes of the two groups, human and agent, to see if there is a difference. In the second experiment, negotiation system users would be randomly assigned to either side of a negotiation case and one of the two negotiators would be assigned to be assisted by the intelligent negotiation agent. As in the previously presented experiment, the two groups could be compared to identify if there is an advance to using the predictive intelligent negotiation assistant agent.

References

- Beam, C., & Segev, A. (1996). *Automated Negotiation: A Survey and the State of Art*. Berkeley, CA: Hass School of Business.
- Chavez, A., Dreilinger, D., Guttman, R., & P., M. (1997). *A Real-life Experience in Creating an Agent Market*. Paper presented at the Application of Intelligent Agents and Multi-Agent Technology PAAM'97, London.
- Chavez, A., Dreilinger, D., Guttman, R., & P., M. (1997). A Real-life Experiment in Creating an Agent Marketplace. In H. S. Nwana & N. Azarmi (Eds.), *Software Agents and Soft Computing* (pp. 160-179): Springer-Verlag.
- Chavez, A., & Maes, P. (1996). *Kasbah: An Agent Marketplace for Buying and Selling Goods*. Paper presented at the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, London.
- Chen, E., Kersten, G. E., & Vahidov, R. (2004, Jan. 14-16, 2004). *An E-marketplace for Agent-supported Commerce Negotiations*. Paper presented at the 25th McMaster World Congress Management of Electronic Business, Hamilton, Ontario, Canada.
- Chen, E., Vahidov, R., & Kersten, G. E. (2005). Agent-Supported Negotiations in the E-marketplace. *Int. J. Electronic Business*, 3(1), 28-49.
- Dzeng, R.-J., & Lin, Y.-C. (2004). Intelligent Agents for Supporting Construction Procurement Negotiation. *Expert Systems with Applications*, 27, 107-119.
- Dzeng, R.-J., & Lin, Y.-C. (2005). Searching for Better Negotiation Agreement Based on Genetic Algorithm. *Computer-Aided Civil and Infrastructure Engineering*, 20, 280-293.
- Faratin, P., Sierra, C., & Jennings, N. R. (2002). Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiations. *Artificial Intelligence*, 142(2), 205-237.
- Foresee, F. D., & Hagan, M. T. (1997, June 9-12, 1997). *Gauss-Newton approximation to Bayesian regularization*. Paper presented at the International Joint Conference on Neural Networks, Houston, Texas.
- Hagan, M. T., Demuth, H. B., & Beale, M. H. (1996). *Neural Network Design*. Boston, MA: PWS Publishing.
- Hagan, M. T., & Menhaj, M. (1994). Training Feedforward Networks with the Marquardt Algorithm. *IEEE Transactions on Neural Networks*, 5(6), 989-993.

- Hornik, K. (1991). Approximation Capabilities of Multilayer Feedforward Networks. *Neural Networks*, 4(2), 251-257.
- Jennings, N. R., & Faratin, P. (2001). Automated Negotiation: Prospects, Methods and Challenges. *Group Decision and Negotiation*, 10, 199-215.
- Kersten, G. E., & Lo, G. (2003). Aspire: An Integrated Negotiation Support System and Software Agents for E-Business Negotiations. *International Journal of Internet and Enterprise Management*, 1(3), 293-315.
- Kersten, G. E., & Noronha, S. J. (1999). WWW-based Negotiation Support: Design, Implementation, and Use. *Decision Support Systems*, 25, 135-154.
- Kwon, O., Shin, J. M., & Kim, S. W. (2006). Context-aware Multi-agent Approach to Pervasive Negotiation Support Systems. *Expert Systems with Applications*, 31, 275-285.
- Lee, W.-P. (2004). Towards Agent-based Decision Making in the Electronic Marketplace: Interactive Recommendation and Automated Negotiation. *Expert Systems with Applications* 27, 665-679.
- MacKay, D. J. C. (1992). Bayesian Interpolation. *Neural Computation*, 4(3), 415-447.
- Maes, P., Guttman, R., & Moukas, A. (1998). Agents That Buy and Sell. *Comm. ACM*, 42(3), 81-91.
- Matwin, S., Szapiro, T., & Haigh, K. (1991). Genetic Algorithms Approach to a Negotiation Support System. *IEEE Transactions on Systems, Man and Cybernetics*, 21(1), 102-114.
- Mudgal, C., & Vassileva, J. (2000, July 2000). *Bilateral Negotiation with Incomplete and Uncertain Information: A Decision-Theoretic Approach Using a Model of the Opponent*. Paper presented at the Workshop on Cooperative Information Agents (CIA IV), Boston.
- Oprea, M. (2002). An Adaptive Negotiation Model for Agent-based Electronic Commerce. *Studies in Informatics and Control*, 11(3), 271-279.
- Vahidov, R., Chen, E., & Feng, Z. (2005, June 10-13, 2005). *Experimental Evaluation of Agent-Supported e-Negotiations*. Paper presented at the Group Decision and Negotiation'2005 (CD-ROM Proceedings), Vienna, Austria.
- Zeng, D., & Sycara, K. (1998). Bayesian Learning in Negotiation. *International Journal of Human-Computer Studies* 48, 125-141.