Predictive Automated Negotiators Employing Risk-Seeking and Risk-Averse Strategies

Marisa Masvoula, Constantine Halatsis, Drakoulis Martakos

Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, University Campus, Athens 15771, Greece
{marisa, halatsis, martakos}@di.uoa.gr

Abstract. Intelligent agents that seek to automate various stages of the negotiation process are often enhanced with models of computational intelligence extending the cognitive abilities of the parties they represent. This paper is focused on predictive strategies employed by automated negotiators, and particularly those based on forecasting the counterpart's responses. In this context a strategy supporting negotiations over multiple issues is presented and assessed. Various behaviors emerge with respect to negotiator's attitude towards risk, resulting to different utility gains. Forecasting is conducted with the use of Multilayer Perceptrons (MLPs) and the training set is extracted online during the negotiation session. Two cases are examined: in the first separate MLPs are used for the estimations of each negotiable attribute, whereas in the second a single MLP is used to estimate the counterpart's response. Experiments are conducted to search the architecture of the MLPs.

Keywords: forecasting opponent's offers, negotiation strategy, multilayer perceptrons for time-series forecasting, predictive negotiating agents

1 Introduction

Negotiation is defined as an exchange mechanism between two or more parties that jointly determine outcomes of mutual interest. The field has attracted the interest of researchers from several scientific fields, providing different viewpoints and approaches [1-6]. Computer science contributes to the development of negotiation theory and examines its applied nature with the construction of negotiation tables, decision and negotiation support systems, software agents and software platforms [4]. The use of AI-based techniques to support various stages of the negotiation process has been viewed as a step towards extending the negotiator's cognitive abilities. In this paperwork focus is set on the implementation of automated agents that use predictive decision-making strategies to attain more beneficial negotiation outcomes. In this respect a negotiation strategy employed by predictive agents that engage in multi-issue negotiations is presented. Different types of behaviors with respect to the level of risk emerge, and negotiation outcomes assessed in terms of utility gain (measures of individual satisfaction) are studied. In the second section definitions and terminologies along with related work are presented, citing systems enhanced with

AI-based techniques to improve the negotiation outcome. In the third section a predictive decision making mechanism supporting automated negotiations over multiple issues is proposed, and in the fourth section two cases of conducting predictions with the use of neural networks are presented. Finally in the fifth section experimental results are illustrated and in the sixth section future research issues are discussed.

2 Definitions, Terminology and Related Work

In a negotiation process participants exchange offers and counter-offers in the search of an agreement. The outcome can be a compromise or a failure and satisfaction of each participant α is measured in terms of a utility function $U_\alpha(X): R^n -> [0,1]$, where X is an offer vector. The negotiable objects may consist of multiple attributes (*issues*). For each issue participants specify a range of permissible (reservation) values (a minimum and a maximum) which they are not willing to exceed. Additionally in many cases participants set a deadline indicating the maximum time they can spend in a negotiation encounter. The specific rules of communication that guide the interaction constitute the negotiation *protocol* and determine the way messages are exchanged. In this paper the negotiation protocol used to support automated negotiations is based on the one described in [7]. The decision making rules or *strategies* are used to determine, select and analyze the decision alternatives. In the simple case where negotiation is conducted between two non learning agents, alternatives are generated with the use of formal decision functions [7]. More sophisticated agents enhance their strategies with AI-based techniques and develop particular skills with the scope to maximize the incurred utility to the party they represent. In our previous work [8] a categorization of such agents is given. This research is focused on the prediction of the counterpart's future offers. Predictions can be generally grouped to single and multi lag. In cases of single–lag predictions agents estimate the very next offer of their counterpart, while in cases of multi-lag predictions they foresee future offers of their counterpart several time steps ahead. Reference [9] depicts the development of a neural network predictive model in order to facilitate "What-if" analysis and generate optimal offers in each round. A similar negotiation support tool is applied by [10] in a supplier selection auction market, where the demander benefits from the suppliers' forecasts, by selecting the most appropriate alternative in each round. In references [11],[12] an agent applies the predictive mechanism only at the pre-final step of the process, in order to increase the likelihood of achieving an agreement, and to produce an outcome of maximal utility. An older work concerning single-lag predictions in agents' strategy can be found in [13]. Trading scenarios via an internet platform are facilitated with the use of SmartAgent, enhanced with predictive decision making. The estimation of the counterparts' next move is used at each negotiation round to adjust the agents' proposal and leads to increased individual gain of the final outcome. As far as multi-lag predictions are concerned, interesting approaches can be found in [14-16], where prediction of counterpart's future offers and strategy, has been used to effectively detect and withdraw from pointless negotiations. In

references [17-19] a negotiating agent enhanced with predictive ability in order to determine the sequence of optimal offers "knowing" the sequence of opponent's responses, has been implemented.

This paper is focused on single-lag predictions and builds on Oprea 's earlier work [13], extending the strategy of the predictive agents to support negotiations over multiple issues (and not just single-issued as described in Oprea). Additionally a variation of the strategic rule is provided with the scope to generate different types of behaviors with respect to the agent's attitude towards risk. In the following section the extended strategy is illustrated.

3 Extending the strategy of predictive agents to support multi-issue (automated) negotiations incorporating different risk attitudes

The negotiation environment considered is tied to bilateral or one-on-one multi-issue negotiations, where all issues are bundled and discussed together (package deal). The formal model of negotiation is comprised by the set of agents A = { α, b }, a finite set of quantitative issues under negotiation I = {i_1, i_2, . . ., i_n}, the domain of reservation values D_i^a : $[\min_i^a, \max_i^a]$ attributed by each agent α in A for each issue i in I, and the deadline $T_{\max}^a$ of each agent α. In the cases studied time variable t is discrete and expresses the interaction step (negotiation round). The possible outcomes of a negotiation can be understood in terms of utility $U^a(X_{(a \to b)}^t)$ where

$X_{(a \to b)}^t = \left(x_{1(a \to b)}^t, x_{2(a \to b)}^t, ..., x_{n(a \to b)}^t \right)^T$ is the negotiation offer sent from agent α

to b at time t, and each x_i denotes the offered value of negotiable issue i. Utility of the offered value is computed as the weighted sum of utilities attributed to each issue, thus:

$$U^a(X_{(a \to b)}^t) = \sum_{i=1}^{n} w_i^a * U^a(x_i)$$

(1)

where $U^a(x_i)$: $D_i^a \to [0,1]$ is the utility function specified by agent α to evaluate

the value x_i of issue i and w_i^a is a normalized weight signifying the relative importance of issue i to agent α.

Each agent α is configured with a default strategy $S^{\alpha,}$, which determines the level of concession in each round. Classification of strategic families to time dependent (TD), resource dependent (RD) and behavior dependent (BD) is used as in [7], reflecting the agent's behavior with respect to the elapsing time, resources and counterpart's responses respectively. In each time step t agent α estimates the next

offer of his counterpart, $\hat{X}_{b \to a}^{t+1} = \left(\hat{x}_{1(b \to a)}^{t+1}, \hat{x}_{2(b \to a)}^{t+1}, ..., \hat{x}_{n(b \to a)}^{t+1} \right)^T$. The proposed

decision rule makes use of the default strategy (S^a) of the predictive agent to generate offers until the detection of a "meeting point" (MP) with the "opponent". MP is a point which would result an established agreement if the agent was guided solely by

his default strategy. When such point is detected, and according to the agent's attitude towards risk, agent risks staying in the negotiation in order to maximize the utility of the final agreement. In this respect two extreme attitudes can be generated: risk-seeking and risk-averse. The risk-seeking agent is willing to spend all the remaining time until expiration of his deadline engaging in an adaptive behavior to turn the estimations of his counterpart's responses to profit. This risk-seeking behavior is based on the one discussed [13], and is extended to support multiple issues. More specifically:

Risk-seeking Behavior:

For each issue i

If issue value is increasing with time

$$x^t_{i(a \to b)} = \hat{x}^{t+1}_{i(b \to a)} - \varepsilon$$

Else

$$x^t_{i(a \to b)} = \hat{x}^{t+1}_{i(b \to a)} + \varepsilon$$

End For

Generate Offer $X^t_{a \to b} = \left(x^t_{1(a \to b)}, x^t_{2(a \to b)}, ..., x^t_{n(a \to b)} \right)^T$

where ε is a domain dependent parameter.

On the other hand risk-averse agents follow a more conservative behavior when they detect an MP. They do not make any further concessions and insist on sending their previous offer, waiting for the opponent to establish an agreement.

Risk-Averse Behavior:

For each issue i

$$x^t_{i(a \to b)} = x^{t-2}_{i(a \to b)}$$

End For

Generate Offer $X^t_{a \to b} = \left(x^t_{1(a \to b)}, x^t_{2(a \to b)}, ..., x^t_{n(a \to b)} \right)^T$

Fusions of the two extreme attitudes have led to the specification of risk portions (RPs) which characterize the predictive agent's behavior after the detection of MP. RP_α determines the percentage of the distance between MP and deadline T^a_{max} that agent α is willing to adopt the risk-seeking behavior. After RP_α is consumed agent adopts the risk-averse behavior. For a predictive agent who is not willing to take any risks RP_α is set to 0%, while for an agent who is willing to risk until expiration of his deadline RP_α is set to 100%. The decision making rule repeated in each step is thus formulated as follows:

If $U^a\left(\hat{X}^{t+1}_{b \to a} \right) > U^a\left(X^t_{a \to b} \right)_{default}$ *(detection of MP)*

    If RP_α is not consumed
   Generate Offer adopting Risk-Seeking Behavior
    Else
   Generate Offer adopting Risk-Averse Behavior
    Else
        Generate Offer $\left(X^t_{a \to b} \right)_{default}$

Where $\left(X_{a \to b}^{t}\right)_{default}$ is the offer generated by agent α at time t based on his default strategy.

Figure 1(a) illustrates how the predictive agent with RP 100% may "tease" his opponent until an agreement is established. At the other end, Figure 1(b) illustrates a risk-averse agent with RP 0% who reaches an agreement faster than the risk-seeking agent, but incurs smaller increase in utility. Negotiation is conducted between a provider and a consumer agent, over service terms of electricity trade, characterized by four negotiable attributes (number of Kwh, Price per Kwh, Penalty terms (returns), and Duration of service provision). The latter agent uses the predictive decision rule.
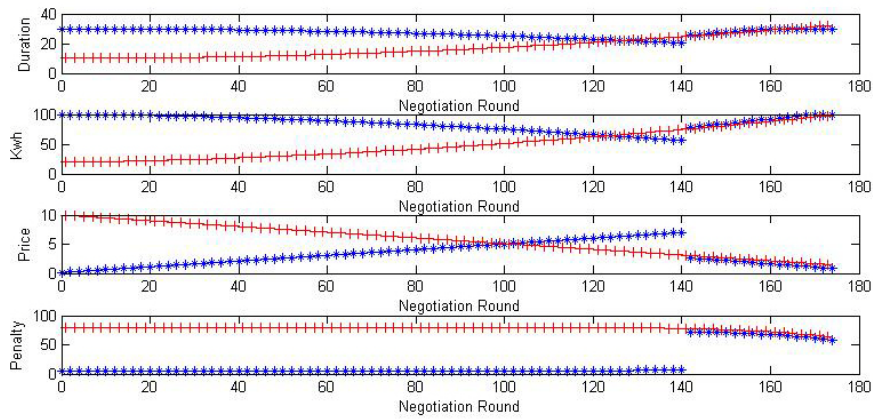
(a)

(b)

Fig. 1. (a) a risk-seeking consumer adjusts his offer with respect to the counterpart response and (b) a risk-averse consumer "freezes" his offer after MP is detected.

## 4 Forecasting Tools Employed by Predictive Agents

Forecasting techniques employed by predictive agents are mainly summarized into statistical approaches (particularly non-linear regression), mathematical models based on differences, and neural networks. As discussed in [20], what is most appealing in neural networks is that they are data-driven, self adaptive, capable of modeling non-linear relations and do not require specific assumptions of the underlying data distributions. Current trend on providing offer forecasts lies on neural networks. In this research, a Multi-layer Perceptron (MLP) with one hidden layer using sigmoid activation functions and one output layer using linear activation functions is used for the prediction of counterpart's next offer, as it is shown that such network is capable of approximating any continuous function [21]. The negotiation thread which comprises of the subsequent offers and counteroffers of the engaged parties formulates the data set which is used to train the network. As far as design issues are concerned, two cases are examined; in the first an MLP is considered for each issue, thus for a negotiation over n negotiable attributes n individual MLPs are constructed. Each MLP comprises of J_1 input nodes representing the counterpart's J_1 previously offered values of the particular issue, J_2 nodes in the hidden layer, and one node in the output layer representing the predicted response. The values of J_1 and J_2 are selected after empirical evaluation. Training using the Levenberg and Marquardt (LM) method is conducted during the negotiation session, giving rise to session-long learning agents. Each network is initialized with random weights and in every negotiation round the network is re-trained with data extracted from the current thread. In the second case a single MLP undertakes the task of prediction. If the J_1 previous offers of the opponent are considered for negotiations over n attributes, then an MLP with $n*J_1$ input nodes, J_2 nodes in the hidden layer and n nodes in the output layer is constructed. As in the first case, the network is initialized with random weights and is trained in each round of the predicting agent with data extracted from the current negotiation thread using the LM method. Values of J_1 and J_2 are also empirically evaluated.

5 Experiments: Measuring Outcomes with respect to the different types of risk attitudes

This section is divided in two parts. The first part is focused on assessing negotiation results with respect to agent's attitude towards risk. The second part is focused on searching optimal number of input and hidden nodes for each case discussed in section 4. A simulation of agent interactions has been developed in Java, and the Java classes have been imported and extended in Matlab (R2008a).

5.1 Evaluation of negotiation results with respect to agent behavior after detection of MP

The proposed strategy is tested with consumer agents with perfect predicting skills (yielding zero errors) and providers following TD strategies. The experimental workbench issues nine different scenarios with respect to deadline, and overlap of agreement zones of the two negotiators ({ $T_{\max}^{Con} = T_{\max}^{Pr}$, $T_{\max}^{Con} < T_{\max}^{Pr}$, $T_{\max}^{Con} > T_{\max}^{Pr}$ } $\times$ {Φ=0, Φ=0.33, Φ=0.66}), where $T_{\max}^{a}$ $\in$ [50:100:350], a={Con,Pr}, and Φ is a parameter which indicates overlap of the agreement zones [7]. In each scenario a variety of concession curves is considered in order to build the default strategies of the opposing agents. For each of the 2352 generated negotiation environments different RPs are set to predictive agent (consumer) (RP$_{Con}$ $\in$ [0:5:100]), leading to an overall of 49392 experiments. The objective is to measure the gain of consumer agent with respect to his attitude towards risk, and highlight the value of forecasting counterpart's next offer in multi-issued negotiations. Results are summarized in Figure 2 where the average gain of the agent adopting the learning strategy illustrated in section 3 is depicted with respect to Risk Portion (RP).

Fig. 2. (%) Average gain in Utility with respect to RP

As it is shown in Figure 2 an agent with RP 0% incurs an average increase of 2.4% in utility, while an agent with RP 100% incurs an average increase of 26.2% in utility. Figure 2 also illustrates the smooth increase of the percentage of utility gain with the increase of RP. Similar experiments are conducted for single-issued negotiations and in cases where an agreement is feasible average utility of the learning agent is increased up to 2% when RP is set to 0%, and up to 30.6% when RP is set to 100%.

5.2 Empirical Search of (sub)optimal network architectures considering the two cases

Results of section 5.1 are extracted when agents with perfect predicting skills are employed. In this section, agents enhanced with neural networks are considered. More specifically the two cases discussed in section 4 are implemented and assessed. Taking the first case, the predicting agent constructs an MLP for each negotiable issue. Negotiations with RP set to 100% are conducted and the average error of the predictive mechanism is computed. In each decision making step the agent makes an estimation of the counterpart's next offer. This estimation is compared to the true offer vector of the counterpart and the absolute error is computed in terms of

Euclidean distance. The subset of input features J_1 expressing the past offers of the opponent for a particular issue, as well as the number of hidden nodes are empirically searched in the space $\{2,3,4,5\}x\{2,3,4,5\}$. The search space comprises of 16 neural networks and is selected to be small since only a few patterns extracted from the current thread will be available for training. Preprocessing, in terms of normalization, is applied to the input data set. At the end of each negotiation, the mean of the absolute errors is computed for each network. The same procedure is also repeated in the second case, where a single neural network is used to predict the counterpart's next offer vector. For an offer which consists of n=4 attributes and for the case where the $J_1 \in \{2,3,4,5\}$ previously sent offers of the opponent are considered, the optimal number of input and hidden nodes is searched in the space $\{8,12,16,20\}x\{2,3,4,5\}$. For each case 192 negotiation environments are generated and 16 ANNs are tested, leading to a total of 3072 experiments. The overall mean of the absolute errors is used to assess the predictive models. Results show that the neural network yielding the smallest error and smallest standard deviation comprises of 5 input and 4 hidden nodes, when an MLP is constructed separately for each issue (first case). For this ANN the average increase in utility attained by the predictive agent is 10.78%. On the other hand, in the second case where a single MLP is employed for the prediction of the counterpart's response, the smallest average error is yielded when 8 input nodes (stemming from the counterpart's 2 previous offer vectors) and 5 hidden nodes are used. This model returns an average increase in utility of 10.5%. The smallest average standard deviation is yielded when 20 input and 5 hidden nodes are used. The last ANN yields 10.34% average increase in utility. The low value of the average standard deviation signifies smoother predictive curves, as estimations do not deviate much from the mean. Table 1 summarizes the results with respect to the combination of input-hidden nodes.

Table 1. Mean errors and mean standard deviations for each combination of (input,hidden) nodes in the MLP, are illustrated for each case. Minimum values are depicted in bold style.

Input \ Hidden	Mean Error				Mean Std Deviation			
	2	3	4	5	2	3	4	5
Case 1								
2	1.62	1.36	1.6	1.81	3.49	1.67	1.97	2.07
3	1.60	1.20	1.17	1.26	3.03	1.72	1.38	1.47
4	1.69	1.11	1.03	1.08	3.44	1.37	1.19	1.29
5	1.57	1.07	**0.98**	1.07	2.91	1.32	**1.16**	1.26
Case 2								
8	1.19	0.63	0.53	**0.49**	3.63	1.38	1.03	0.97
12	1.14	0.64	0.50	0.51	3.21	1.18	0.88	0.92
16	1.00	0.61	0.55	0.51	2.27	1.14	0.98	0.89
20	1.10	0.69	0.53	0.52	3.02	1.35	0.97	**0.83**

6 Conclusions and Future Research

In this paper a predictive strategy adopted by agents who engage in multi-issue negotiations is presented and assessed. Different behaviors emerge from the specifications of risk portions (RP) after the detection of a point of agreement, also termed a meeting point (MP), with the opponent. At one extreme an agent may be guided by a highly risk seeking behavior, where he risks staying in the negotiation until expiration of his deadline so as to maximize the utility of the final outcome. Maximization may be accomplished through adaptation of the predictive agent's offers based on the estimations of the counterpart's responses. At the other extreme an agent may be guided by a risk-averse behavior "freezing" his final offer and not making any further concessions. It is shown that intermediate behaviors, controlled by the RP parameter, lead to smooth increases of the average utility gain. Negotiation forecasts are undertaken by Multilayer Perceptrons (MLPs). Training data are extracted online and are composed of the previously sent offers of the opponent agent. The MLPs are retrained in each round as the data set augments with the counterpart's incoming offer. Two cases are illustrated. In the first an MLP is used to guide the prediction of each individual attribute, while in the second case a single MLP is used to estimate the vector of the counterpart's response. MLPs prove capable of capturing the negotiation dynamics if retrained in each round. However retraining is often time consuming, and time is crucial in most negotiation cases. For this reason as a future research issue we plan on focusing on the resources required, as well as on examining more flexible, evolving structures which engage in online, life-long learning, with one-pass propagation of the training patterns. These characteristics are met by Evolving Connectionist Structures (ECoS), discussed in [22].

References

1. Raiffa, H.: Contributions of Applied Systems Analysis to International Negotiation. In: Kremenyuk's, V.A. (ed.) International Negotiation: Analysis, Approaches, Issues, 2nd Edition, pp. 5--21. Jossey-Bass, San Francisco CA USA (2002)
2. Goh, K.-Y., Teo, H.-H., Wu, H., and Wei, K.-K.: Computer Supported Negotiations: An experimental study of bargaining in electronic commerce. In: Proc. of the 21st Int. Conf. on Information Systems, pp. 104--116. (2000)
3. Bichler, M., Kersten, G., and Strecker, S.: Towards a Structured Design of Electronic Negotiations. Group Decision and Negotiation. 12(4), 311--335 (2003)
4. Kersten, G.: E-negotiations: Towards Engineering of Technology-based Social Processes. In: Proc. of the 11th Eu. Conf. on Information Systems, ECIS 2003, Naples Italy (2003)
5. Moor, A. De and Weigand, H.: Business Negotiation Support: Theory and Practice. International Negotiation. 9(1),31--57 (2004)

6. Schoop, M.: The Worlds of Negotiation. In Proc. of the 9th Int. Working Conf. on the Language-Action Perspective on Communication Modeling, pp. 179--196 (2004)

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

8. Masvoula, M., Kanellis, P., Martakos, D.: A Review of Learning Methods Enhanced in Strategies of Negotiating Agents. In: Proc. of 12th ICEIS, pp. 212--219, (2010)

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

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

11. Papaioannou, I. V., Roussaki, I. G., & Anagnostou, M. E.: Comparing the Performance of MLP and RBF Neural Networks Employed by Negotiating Intelligent Agents. In: Proc. of the IEEE/WIC/ACM Int. Conf. on intelligent Agent Technology, pp. 602--612 (2006a)

12. Papaioannou, I. V., Roussaki, I. G., & Anagnostou, M. E.: Towards Successful Automated Negotiations based on Neural Networks. In: Proc. of the 5th IEEE/ACIS Int. Conf. on Computer and Information Science, pp. 464--472 (2006b)

13. Oprea, M.: The Use of Adaptive Negotiation by a Shopping Agent in Agent-Mediated Electronic Commerce. In: Proc. of the 3rd Int. Central and Eastern European Conf. on Multi-Agent Systems, pp. 594--605. Springer Verlag, Heidelberg, Berlin, (2003)

14. Hou, C.: Predicting Agents Tactics in Automated Negotiation. In: Proc. of the IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology, pp. 127--133. (2004)

15. Roussaki, I., Papaioannou, I., & Anagnostou, M.: Building Automated Negotiation Strategies Enhanced by MLP and GR Neural Networks for Opponent Agent Behaviour Prognosis. In IWANN 2007. pp. 152--161. Springer-Verlag, Heidelberg, Berlin (2007)

16. Papaioannou, I., Roussaki, I., & Anagnostou, M.: Detecting Unsuccessful Automated Negotiation Threads When Opponents Employ Hybrid Strategies. In Proc. of the 4th int. Conf. on intelligent Computing. pp. 27--39. Springer-Verlag, Heidelberg, Berlin (2008)

17. Brzostowski, J., & Kowalczyk, R.: Modelling Partner's Behaviour in Agent Negotiation. In: AI 2005. Lecture Notes in Computer Science 3809, pp. 653--663. Springer-Verlag, Berlin, Heidelberg (2005)

18. Brzostowski, J., & Kowalczyk, R.: Adaptive Negotiation with On-Line Prediction of Opponent Behaviour in Agent-Based Negotiations. In: Proc. of the IEEE/WIC/ACM int. Conf. on intelligent Agent Technology. pp. 263—269. (2006a)

19. Brzostowski, J., & Kowalczyk, R.: Predicting Partner's Behaviour in Agent Negotiation. In: Proc. of the 5th int. joint conf. AAMAS. pp. 355--361. ACM, New York USA (2006b)

20. Zhang, P.: Neural Networks in Business Forecasting. Idea Group Publishing, US and United Kingdom (2004)

21. Hornik, K., Stinchombe, M., & White, H.: Multilayer feed-forward networks are universal approximators. Neural Networks. 2(5), 359--366 (1989)

22. Kasabov, N.: Evolving Connectionist Systems: The Knowledge Engineering Approach. Springer-Verlag, London (2007)