

BEHAVIOR VISUALIZATION OF AUTONOMOUS TRADING AGENTS

Tomoharu Nakashima, Hiroko Kitano, Hisao Ishibuchi

College of Engineering
Osaka Prefecture University
Gakuen-cho 1-1, Sakai, Osaka 599-8531, Japan
{nakashi, kitano, hisaoi}@ie.osakafu-u.ac.jp

KEYWORDS

Fuzzy rule, black-box analysis, futures trading

ABSTRACT

In this paper, we visualize the behavior of a futures trading agent by using fuzzy if-then rules. The main aim of this paper is to graphically interpret how trading agents make decisions such as to buy or to sell a futures stock. Fuzzy if-then rules are used for this aim because their antecedent part specifies the features of time series. The procedure for visualizing the behavior of trading agent is first to train a trading agent that consists of a set of fuzzy if-then rules. By carefully examining, we select a small number of prominent fuzzy if-then rules that most represent the behavior of trading agents. Finally those selected fuzzy if-then rules are represented in a graphically understandable manner.

INTRODUCTION

Fuzzy rule-based systems have been successfully applied to various problems such as control problems (Sugeno 1985, Lee 1990). The advantage of the fuzzy rule-based systems is its interpretability. Since fuzzy if-then rules in a fuzzy rule-based system are written using linguistic values, human users can linguistically understand the meaning of fuzzy if-then rules. The validity of fuzzy if-then rules can be measured using two well-known criteria in the field of data mining: confidence and support (Agrawal et al. 1996, Agrawal and Srikant 1994). The fuzzy version of these measures are also introduced in (Hong et al. 2001, Ishibuchi et al. 2001).

Although there are numerous researches on extracting patterns or rules from a large data base, the number of researches on how such extracted rules are effective for practical use in a particular domain is not large. The purpose of this paper is to examine the effectiveness of the extracted rules for supporting decision making of human users.

In this paper, we consider a virtual futures market as a problem domain. The virtual market allows a number of autonomous agents to take part in the futures market. Human beings are also allowed to trade in the futures stock index in the virtual market. Autonomous agents and human beings are required to determine whether they buy the futures stock index or sell, the limit price, and the quantity of the futures trade.

We have developed an autonomous agent that trades in futures stock index in a virtual market (Nakashima et al. 2002). An adaptive fuzzy rule-based system was used in the autonomous agent. The adaptive fuzzy rule-based system

consists of a number of fuzzy if-then rules that linguistically provide the decision making on the trading action (i.e., buy the futures index or sell) for different conditions. The evaluation (i.e., successful trade or not) of the trade action is performed after the new spot price and the new futures price are obtained. According to the evaluation, the agent adjusts the weights of fuzzy if-then rules in the adaptive fuzzy rule-based system. That is, we increase the weights of fuzzy if-then rules if the agent's decision making in the previous time step is successful. On the other hand, the weights of fuzzy if-then rules are decreased if the decision making is not successful. Since the weight update can be performed on-line, it is expected that the performance of the autonomous agent is gradually improved during the course of the trade. Thus the resultant fuzzy rule-based system after the enough number of trade can be viewed as a knowledge base for the virtual futures trade.

In this paper, we try to linguistically interpret the behavior of the autonomous agent by examining the weights of fuzzy if-then rules in the adaptive fuzzy rule-based system. We select a small number of fuzzy if-then rules that have a contrast between the weights associated with *Buy* the futures index and *Sell*. After the selected fuzzy if-then rules are graphically transformed into the trading knowledge base, the trading knowledge base is shown to a human trader who participates in the virtual futures trade. The human being can consult the trading knowledge throughout the futures trade. Statistical evidence shows that the trading knowledge base extracted from the adaptive fuzzy rule-based system improves the performance of human traders.

A VIRTUAL FUTURES MARKET U-MART

Recently, virtual economic markets have attracted a great deal of attention for analyzing economic systems and developing autonomous agents. From the view point of the economics, the advantages of the virtual economic markets are as follows. First, one can analyze patterns of humans' trading behavior with respect to the trading action. Secondly, one can examine how to avoid the speculative action such as violent fluctuations of stocks. On the other hand, from the engineering point of view, we can examine the actual effectiveness of the use of learning methods, evolutionary methods, and multi-agent techniques in economic systems. The U-MART (Unreal Market as Artificial Research Test-bed) project is one of such virtual markets where multiple players including human beings can simultaneously trade in a futures stock index (Fig. 1).

In the U-MART, a machine (i.e., an autonomous agent) or a human being is called a U-MART client and is given market information such as time series data of spot prices and futures prices. Clients also have its own current information such as its position (i.e., a balancing amount of

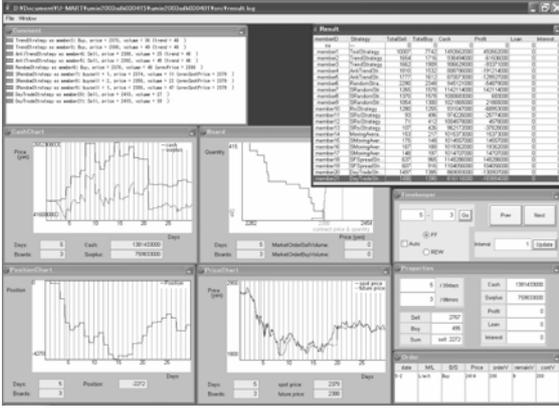


Fig. 1 Snapshot of U-Mart

the futures index trade), remaining cash, and time to the final settlement. Based on the above information, each client has to make a decision on whether it buys or sells the futures index, the limit price of the futures trading, and the quantity of the futures trade. Thus, a client in U-MART can be viewed as an input-output system A as follows:

$$A(\mathbf{S}, \mathbf{F}, Pos, Cash, t) = (BS, P, Q), \quad (1)$$

where \mathbf{S} and \mathbf{F} is the time series of the spot prices and the futures prices (called the U-MART prices), respectively, Pos is the position of a client, $Cash$ is the remaining cash, t is the remaining time to the final settlement, BS represents the client's trading decision on the futures stock index, P is the limit price, and Q is the quantity of the trade. Each U-MART client interacts with the U-MART server for trading in a futures stock index through the TCP/IP protocol. Among the input variables to the U-MART client, the two time series \mathbf{S} and \mathbf{F} are externally provided by the U-MART server and the other input variables are internally held by the U-MART client. In Fig. 2, we show a general view of the trade between a U-MART client and the U-MART server.

The U-MART server determines the futures index price by a method called *Itayose*. In *Itayose*, the U-MART server first collects an order from each U-MART client such as buy or sell of the futures index, the limit price, and the quantity of the trade. Then it compares the buy orders with the sell orders. The futures price is determined at the point where the price and quantity of buy orders are matched by those of sell orders.

The goal of the U-MART clients is to maximize the profit caused by the difference between the selling futures prices and the buying futures prices.

ADAPTIVE FUZZY RULE-BASED SYSTEM

We have already developed a learning U-MART client for the virtual futures market (Nakashima et al. 2002). In (Nakashima et al. 2002), the U-MART client maintains an adaptive fuzzy rule-based system that determines whether the client should buy or sell a futures stock index. The weights of fuzzy if-then rules correspond to the support for buying or selling the futures stock index. During the course of the futures trades, weights of fuzzy if-then rules are updated on-line according to the evaluation of the trade at the previous time step. The trade is evaluated at each time

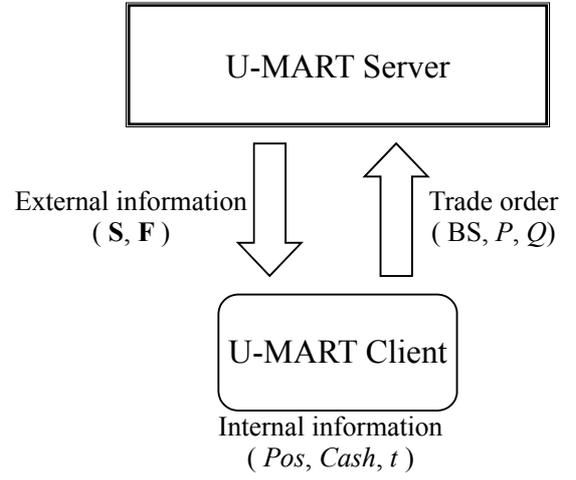


Fig. 2 The general view of the trade in U-MART

step by the high-and-low relation between the spot prices and the futures prices (i.e., U-MART prices). The following subsections describe the adaptive fuzzy rule-based system used for the autonomous agent in detail.

Problem Formulation and Fuzzy If-Then Rules

In this subsection, we explain the fuzzy rule-based system that was applied to the learning U-MART client in (Nakashima et al. 2002). In the futures market, time series data of both the spot and the futures prices are available to the adaptive fuzzy rule-based system. Let us assume that n pieces of information are used by the agent for determining whether to buy or sell the futures stock index. In this case, the problem of the futures market for our futures trade agent can be viewed as a two class pattern classification problem with n -dimensional inputs. A fuzzy rule-based system is applied to this n -dimensional two-class pattern classification problem. The adaptive fuzzy rule-based system in the learning U-MART client consists of fuzzy if-then rules of the following type:

$$R_j: \text{ If } x_1 \text{ is } A_{j1} \text{ and } \dots \text{ and } x_n \text{ is } A_{jn} \\ \text{ then Buy with } b_j \text{ and Sell with } s_j, \quad j = 1, \dots, N, \quad (2)$$

where R_j is a label of j -th fuzzy if-then rule, $\mathbf{x} = (x_1, \dots, x_n)$ is an input vector to the fuzzy rule-based system, A_{j1}, \dots, A_{jn} are antecedent fuzzy sets, and b_j and s_j are real values of the fuzzy if-then rule R_j corresponding to buying and selling the futures stock index, respectively.

In the implementation of the learning U-MART client in this paper, we use the difference between the spot price at the current time step and those at the three different time steps as three input values x_1, x_2, x_3 to the fuzzy rule-based system (Fig. 3). That is, the learning client determines whether it buys or sells the futures stock index from an input vector $\mathbf{x} = (x_1, x_2, x_3)$. Thus, in this paper we deal with the decision making problem as an three-dimensional two-class pattern classification problem. The fuzzy if-then rules can be written as follows:

$$R_j: \text{ If } x_1 \text{ is } A_{j1} \text{ and } x_2 \text{ is } A_{j2} \text{ and } x_3 \text{ is } A_{j3} \\ \text{ then Buy with } b_j \text{ and Sell with } s_j, \quad j = 1, \dots, N, \quad (3)$$

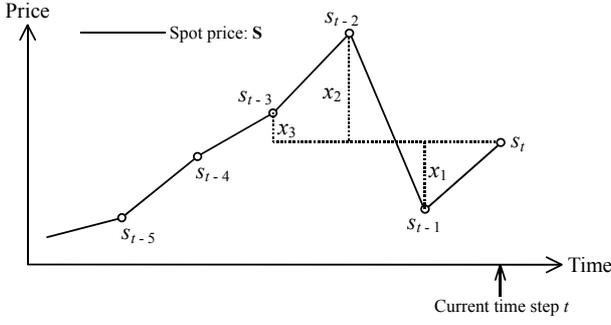


Fig. 3 Input variables in our fuzzy rule-based system

In Fig. 3, $\mathbf{S} = (s_1, \dots, s_t)$ is the time series of spot prices from the beginning of the trade (i.e., s_1) until time step t (i.e., s_t) where s_k is a spot price at time step k . We use the following three pieces of information as input variables for the fuzzy rule-based system:

$$x_1 = s_t - s_{t-1}, \quad (4)$$

$$x_2 = s_t - s_{t-2}, \quad (5)$$

$$x_3 = s_t - s_{t-3}. \quad (6)$$

From the above explanation, we can see that the fuzzy rule-based system performs a mapping from a three dimensional state vector $\mathbf{x} = (x_1, x_2, x_3)$ to a single binary value corresponding to either *Buy* or *Sell*.

Fuzzy Inference and Decision Making

Let us consider that at a particular time we have already calculated three input variables x_1 , x_2 , and x_3 for the fuzzy rule-based system. In this subsection, we show how an agent makes a decision on whether it buys or sells the futures stock index. Assume that there are N fuzzy if-then rules in the fuzzy rule-based system. In this paper, we divide each axis of input variables into three fuzzy sets as in Fig. 4. In Fig. 4, N, Z, and P represent linguistic terms *negative*, *zero*, and *positive*, respectively. Since there are three input variables in our fuzzy rule-based system and three fuzzy sets for each input variable, the total number N of fuzzy if-then rules involved in the fuzzy rule-based system is $N = 3^3 = 27$.

Note that each fuzzy if-then rule has two weight values associated with buying and selling the futures index, respectively. After the calculation of three input values x_1 ,

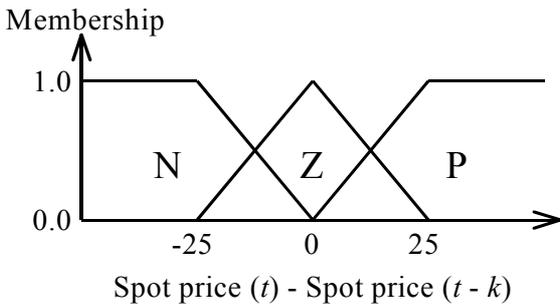


Fig. 4 Membership functions

x_2 , and x_3 , Q_{Buy} and Q_{Sell} are calculated using fuzzy inference as follows:

$$Q_{Buy} = \frac{\sum_{j=1}^N \mu_j(\mathbf{x}) \cdot b_j}{\sum_{j=1}^N \mu_j(\mathbf{x})}, \quad (7)$$

$$Q_{Sell} = \frac{\sum_{j=1}^N \mu_j(\mathbf{x}) \cdot s_j}{\sum_{j=1}^N \mu_j(\mathbf{x})}, \quad (8)$$

where $\mathbf{x} = (x_1, x_2, x_3)$ is an input vector to the fuzzy rule-based system, $\mu_j(\cdot)$ is the compatibility of the input vector \mathbf{x} with the fuzzy if-then rule R_j . The compatibility μ_j of the input vector \mathbf{x} with the fuzzy if-then rule R_j is calculated by the following product operator:

$$\mu_j(\mathbf{x}) = \mu_{j1}(x_1) \cdot \mu_{j2}(x_2) \cdot \mu_{j3}(x_3), \quad (9)$$

where $\mu_{ji}(\cdot)$ is the membership function of an antecedent fuzzy set A_{ji} in the j -th fuzzy if-then rule R_j (see Fig. 4).

After calculating Q_{Buy} and Q_{Sell} , the agent makes a decision on whether the agent buys or sells the futures index based on the following decision rule:

[Decision Rule]

If $Q_{Buy} > Q_{Sell}$, the agent buys the futures index,

Else if $Q_{Buy} < Q_{Sell}$, the agent sells the futures index,

Otherwise, the agent's trade is the same as the decision at the previous time step.

Note that the autonomous agent does not make a decision for the first three time steps since there are not enough information to calculate an input vector for the first three time steps. That is, the agent only collects the information for the decision making for the first three time steps.

In this paper, we assume that the agent tries to optimize the decision making only on whether the agent buys or sells the futures index under various conditions of the time series data of spot prices. Thus fuzzy if-then rules in the adaptive fuzzy rule-based system have only two weights corresponding to *Buy* and *Sell*. There are, however, two more things to be determined in order to trade the futures index. One is the limit price and the other is the quantity of the trade.

The other two pieces of information such as the limit price and the quantity of the trade order are determined as follows. First, the limit price P in (1) is determined as follows:

$$P = \begin{cases} s_m - 5, & \text{if decision is Buy,} \\ s_m + 5, & \text{otherwise.} \end{cases} \quad (10)$$

That is, the limit price is deterministically decided based on the spot price at the current time step. On the other hand, the quantity Q of the trade in (1) is specified as $Q = 200$. That is, we don't adaptively determine the quantity of the trade according to any information but fixed to a prespecified value.

On-Line Learning of Fuzzy If-Then Rules

In this subsection, we show how weights of fuzzy if-then rules in our fuzzy rule-based system are adjusted so that the agent can maximize the profit through the trading.

Let us assume that the agent has already made a decision on whether the agent buys or sells the futures index. At the next time step, the agent is given another information on the time series data of spot prices (i.e., s_{m+1}) and the futures prices (i.e., f_{m+1}). We evaluate the agent's trade decision (*Buy* or *Sell*) according to the high-and-low relation between s_{m+1} and f_{m+1} as follows:

[Evaluation Criterion]

If $Q_{Buy} < Q_{Sell}$ and $s_{m+1} > f_{m+1}$, then the decision is evaluated as successful,

Else If $Q_{Buy} < Q_{Sell}$ and $s_{m+1} < f_{m+1}$, then the decision is evaluated as successful,

Otherwise the decision is evaluated as unsuccessful.

That is, we evaluate the decision making based on the absolute price difference between the spot price and the futures price at the following time step. If the agent's decision is *Buy*, the evaluation for the decision is successful only if the spot price is higher than the futures price. On the other hand, if *Sell* is chosen, the evaluation for the decision is successful only if the spot price is lower than the futures prices. This criterion is derived from the observation that the spot price and the futures price must be coincide at the final settlement (see Fig. 5). This evaluation is used for updating the weights of fuzzy if-then rules.

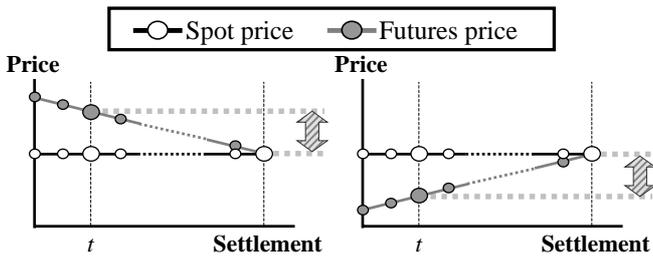


Fig. 5 Final settlement.

The main idea of the weight update is that the weights of the fuzzy if-then rules that contribute to the successful decision making are increased while we decrease the weights of those fuzzy if-then rules that are responsible for unsuccessful decision making. Thus the update rule of the weights is described as follows:

$$b_j^{new} = \begin{cases} b_j^{old} + \alpha \cdot (1 - b_j^{old}) \cdot \mu_j(\mathbf{x}), & \text{if successful,} \\ b_j^{old} - \alpha \cdot b_j^{old} \cdot \mu_j(\mathbf{x}), & \text{otherwise,} \end{cases} \quad (11)$$

$$s_j^{new} = \begin{cases} s_j^{old} + \alpha \cdot (1 - s_j^{old}) \cdot \mu_j(\mathbf{x}), & \text{if successful,} \\ s_j^{old} - \alpha \cdot s_j^{old} \cdot \mu_j(\mathbf{x}), & \text{otherwise,} \end{cases} \quad (12)$$

where α is a positive learning rate and b_j and s_j are weight values of the j -th fuzzy if-then rule R_j that are associated with buying and selling the futures stock index, respectively. Note that only the weights corresponding to the selected action are updated by the above equations. We do not update the weights corresponding to the action that is not selected. For example, when we select to *Buy* the futures stock index, the weights b_j , $j = 1, 2, \dots, N$, are updated and we do not modify the weight s_j that are associated with the action *Sell*.

To summarize, the procedure of our learning client for the futures trade is described as follows:

[Procedure of on-line learning for futures trading]

- Step 1: *Initialization*. Set initial weights of the fuzzy if-then rules to either some prespecified values or random values.
- Step 2: *Decision making*. Using the time series of spot prices, calculate the value corresponding to *Buy* or *Sell* as in (7) and (8). Make a decision of *Buy* or *Sell* according to the decision rule described in the previous subsection. The limit price and the quantity of the trade are also determined.
- Step 3: *Evaluation*. Given the futures price and the spot price at the following time step, evaluate the decision making of *Buy* or *Sell* as successful or not according to the evaluation criterion described in this subsection.
- Step 4: *Weight update*. Update the weights of fuzzy if-then rules involved in the fuzzy rule-based system. Note that only those weights corresponding to the decision making (i.e., *Buy* or *Sell*) at the current time step are updated.

These above steps are iterated until the contract month (i.e., the final settlement) is reached in the futures trading.

VISUALIZATION

Procedure of Knowledge Extraction

The learning U-MART client in the last section can be used as a knowledge acquisition tool since the adaptive fuzzy rule-based system in the U-MART client can be seen as a knowledge base for the virtual futures trade. In this section, we examine such possibility through laboratory experiments. The knowledge extraction procedure consists of two phases: tuning of fuzzy if-then rules in the adaptive fuzzy rule-based system and selecting a small number of fuzzy if-then rules with a large contrast between consequent weights. In the following subsections, each phase of the knowledge extraction is explained.

Tuning and Interpreting the Fuzzy Rule-Based System

First, the learning U-MART client with the adaptive fuzzy rule-based system is iteratively applied to the virtual futures market. Since the learning client needs a number of iterations for learning the weights of fuzzy if-then rules, we repeated the virtual futures trade several times. After the futures trade, it is expected that those fuzzy if-then rules that are related to the critical input states have a contrast between weights for *Buy* and *Sell*. For example, the weight of *Buy* is larger than *Sell* for a fuzzy if-then rule if the U-MART client has made a number of successful decision making of *Buy* in a situation compatible to the antecedent part of the fuzzy if-then rule. Such a fuzzy if-then rule is likely to suggest *Buy* in the corresponding situation. Another example is that if the U-MART client has made a number of unsuccessful decision making in a situation compatible to the antecedent part of a fuzzy if-then rule, the weight corresponding to the decision making becomes smaller than the weight corresponding to the other decision making. In this case, the suggestion by such a fuzzy if-then rules is not to perform the trade action (either *Buy* or *Sell*) with the smaller weight.

Selecting a Small Number of Fuzzy If-Then Rules

We examined the weights of each fuzzy if-then rule to select a small number of fuzzy if-then rules with a strong contrast between weights for *Buy* and *Sell*. From the total number of 27 ($=3^3$) fuzzy if-then rules, we manually selected five such fuzzy if-then rules. Table 1 shows the selected fuzzy if-then rules with a strong contrast in the weights.

Table 1 Selected Fuzzy If-Then Rules

No	x_3	x_2	x_1	q_{j1}	q_{j2}
1	N	N	N	0.598	0.086
2	N	Z	P	0.778	0.318
3	N	P	N	0.546	0.870
4	P	Z	Z	0.573	0.800
5	P	P	P	0.924	0.141

Note that these fuzzy if-then rules were selected manually and subjectively according to the difference in the weights of fuzzy if-then rules. Although it is possible to systematically select a small number of fuzzy if-then rules using some statistical technique, it is beyond our scope of this paper. It will be investigated in our future research.

Experiments with Human Users

In this subsection, we show the experimental results where human users are provided with a small number of the selected fuzzy if-then rules when they participate in the virtual futures trade. In order to make the selected fuzzy if-then rule more understandable, we visualized the selected fuzzy if-then rules as shown in Fig. 6. In Fig. 6, graphs (a)-

(e) correspond to the selected rules No. 1-5 in Table 1, respectively. The visualization is done such that the antecedent linguistic values (N, Z, and P) are interpreted as the relative position between the present spot price s_t and the previous spot prices s_{t-3} , s_{t-2} , and s_{t-1} , and the recommended action is determined according to the extreme value of the consequent value (0 or 1) for each selected fuzzy if-then rule (Note that the value 0 means that the corresponding action is not recommended and it is recommended when the value is 1).

In our experiments, six human users separately participated in the virtual futures trade twice. One experiment is done with the presentation of the selected fuzzy if-then rules, and the other without the presentation of the selected fuzzy if-then rules. We performed this experiment for six different human users. For three human users, the first experiment was done with the presentation of the selected fuzzy if-then rules and the second experiment without the presentation. On the other hand, the experiments for the other three human users were done with the selected fuzzy if-then rules presented in the first experiment and without the presentation in the second experiment. This is because we need to minimize the effect of the ordering condition in the experiments. That is, the bias of presenting the selected fuzzy if-then rules in the first experiment for the first three human users is offset by the bias of presenting them in the second experiment for the other three.

We show the experimental results in Table 2. Table 2 shows the remaining assets after the final settlement.

Table 2 Experimental results

Human ID	Without guide	With guide
A	258,687,000	1,449,355,000
B	878,280,000	2,093,095,000
C	961,393,000	2,335,901,000
D	*	1,675,616,000
E	926,983,000	1,220,221,000
F	740,654,000	508,363,000

* shows the human user went bankruptcy.

From Table 2, we can see that almost all the human users could performed better with the presentation of the selected fuzzy if-then rules than without them. Thus we can expect that our learning client could become a human decision support tool. In order to confirm this observation statistically, we perform the Wilcoxon's rank-sum test. The Wilcoxon's rank-sum test is a nonparametric test that is a sample t -test based solely on the order in which the observations from the two samples fall. In the Wilcoxon's test, we order the results of human users in the descending

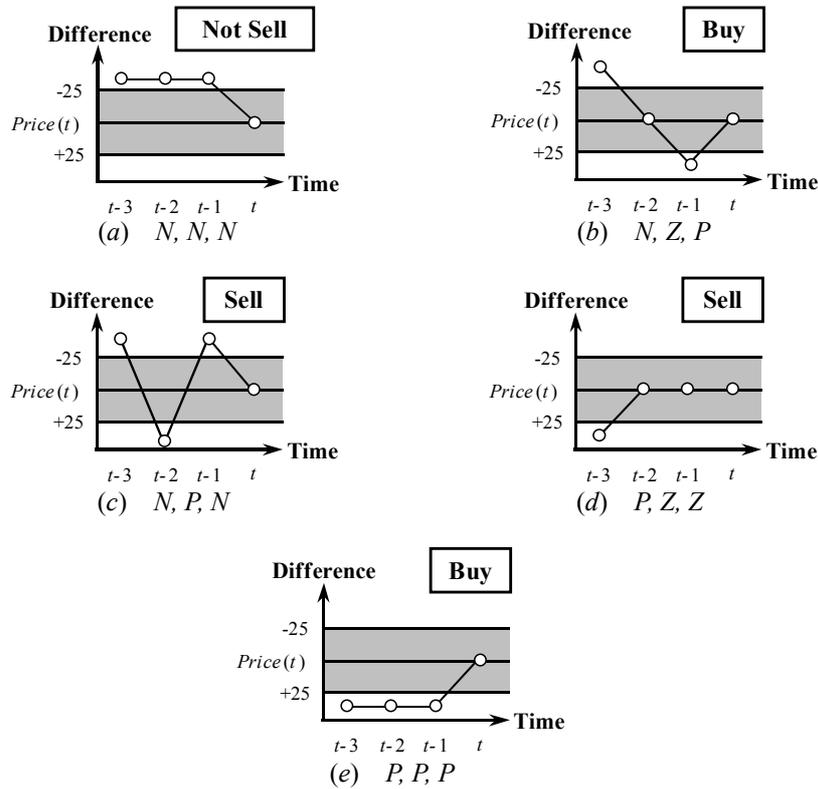


Fig. 6 Visualization of the selected fuzzy if-then rules

order of the final remaining assets in Table 2. The order is used as a rank for each human user. The sum of the rank is used as the statistic for one-sided test. In the test that compares the null hypothesis (there is no difference between the result with and without the presentation of the selected rules) against the alternative hypothesis (there is difference), the null hypothesis is rejected with a 0.05 level. Thus, we can statistically say that the human users can perform better with the help of the selected fuzzy if-then rules.

CONCLUSIONS

In this paper, we presented how an autonomous agent in a virtual futures market is graphically analyzed. The autonomous trading agent has a learning mechanism during the course of the trade. A set of fuzzy if-then rules were used whose antecedent part represents the ups-and-downs of time series data. The consequent part of fuzzy if-then rules is the trading decision of the agent, i.e., *Buy* or *Sell*. First we applied a learning method to a virtual futures market in order to obtain the possible futures trading knowledge. Then we manually selected a small number of fuzzy if-then rules with a strong contrast in weights for decision making options. The selected fuzzy if-then rules were presented to human users after we visualized those fuzzy if-then rules in order for human users to easily understand them. Experimental results showed that there was a positive effect of presenting the selected fuzzy if-then rules. That is, human users could achieve better remaining

assets by using the extracted knowledge for the futures trade.

REFERENCES

- M. Sugeno. 1985, "An Introductory Survey of Fuzzy Control", *Information Science*, Vol. 30, No. 1/2:59-83.
- C. C. Lee. 1990, "Fuzzy logic in control systems: fuzzy logic controller Part I and Part II", *IEEE Trans. Syst., Man, Cybern.*, Vol. 20: 404—435.
- R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo, 1996, "Fast discovery of association rules", *Advances in Knowledge Discovery & Data Mining*, 307-328, AAAI Press, Menlo Park, CA.
- R. Agrawal and R. Srikant, 1994, "Fast algorithms for mining association rules", *Proc. of 20th International Conference on Very Large Data Bases*, 487-499. Expanded version is available as IBM Research Report RJ9839, 1994.
- T.-P. Hong, C.-S. Kuo, and S.-C. Chi, 2001, "Trade-off between computation time and number of rules for fuzzy mining from quantitative data", *International Journal of Uncertainty Fuzziness and Knowledge-Based Systems*, vol. 9, no. 5: 587-604.
- H. Ishibuchi, T. Yamamoto, and T. Nakashima, 2001, "Fuzzy data mining: Effect of fuzzy discretization", *Proc. of 1st IEEE International Conference on Data Mining*, 241-248.
- T. Nakashima, T. Ariyama, and H. Ishibuchi, 2002, "On-line learning of a fuzzy system for a futures market", *2002 International Conference on Fuzzy Systems and Knowledge Discovery*, 54-58.