COMBINED NUMERIC/KNOWLEDGE BASED HIERARCHICAL CONTROL

Part II. On-line Fault Diagnosis With Deep Knowledge Based Self-learning of Heuristic Rules

FLORIN GHEORGHE FILIP

Research Institute for Informatics
8-10 Averescu Avenue,
71316 Bucharest
ROMANIA

PETER DAVID ROBERTS

Control Engineering Centre
City University
Northampton Square, London EC1 OHB
UNITED KINGDOM

JIE ZHANG

Department of Chemical & Process Engineering
University of Newcastle upon Tyne
Newcastle upon Tyne, NE1 7RU
UNITED KINGDOM

ABSTRACT

This is a new part of our series dedicated to presenting the results obtained in applying the combinations of numeric and knowledge based techniques to various layers of the control hierarchy such as regulation, supervision and process co-ordination. This paper describes a diagnosis system (developed by dr.J.Zhang and Professor P.D.Roberts, while working together at the Control Engineering Centre, City University, London), which detects and diagnoses faults by the combined use of both shallow knowledge, in the form of heuristic rules, and deep knowledge, in the form of qualitative models, and exhibits learning characteristics in that the rules can be learnt from deep knowledge based diagnosis. Since heuristic rules can usually give efficient diagnosis, they are used to propose a hypothesis when abnormalities occur in the measurements of the monitored process. The proposed hypothesis is then tested on a deep qualitative model of the process to enhance diagnostic reliability. The heuristic rules used are learnt from the deep knowledge based diagnosis using the ideas of explanation based learning (EBL). After a successful diagnosis, which is not proposed by a heuristic rule, the qualitative deviations in on- line information are calculated to form a training example which is generalised based on the deep model of the process. By such means, heuristic diagnostic rules can be accumulated automatically and, therefore, ease the task of knowledge acquisition. The proposed learning method has been tested on the fault diagnosis of a pilot scale mixing process.

KEYWORDS: Fault diagnosis, expert systems, qualitative reasoning, self-learning of diagnostic rules.

1. INTRODUCTION

The first part of the series was "A Survey of Reported Results" and was published in a previous issue of the journal.It was a review of several world-wide results regarding the use of KBS or numeric/knowledge based techniques in various off-line control applications (design of control problems, process identification) and on-line applications (heuristic control, expert control, fault detection and diagnosis). This second part addresses the specific problem of on-line fault diagnosis and includes original solutions advanced by the Control Engineering Centre, City University, London.

Failures in an industrial process can be understood as the deviations of some process units from their normal functions. Failures may cause poor product quality, damage equipment, cause process shut down, or even result in hazardous conditions. Therefore, it is necessary to continuously monitor process conditions in order to detect and diagnose faults promptly. This task is traditionally carried out by process operators.

Recently developed expert systems techniques show considerable potential in the application of process condition monitoring, fault detection and diagnosis. Expert systems for industrial process fault diagnosis can generally be divided into two categories: a shallow knowledge based approach and a deep knowledge based approach. In the first category the knowledge base contains heuristic rules which encode the experiences of process operators. This type of expert system can usually diagnose faults very efficiently because heuristics can provide valuable short cuts [5,8]. Lapointe et al [5] developed an expert system for waste water treatment process diagnosis - BIOEXPERT, in which shallow knowledge is used for diagnosing the more common faults. Since the knowledge base does not contain any knowledge about system structure and component functions, it may have difficulties when dealing with novel faults and infrequently occurred faults. In contrast, in the deep knowledge based approach, the knowledge base contains information on system structures and unit functions as well as physical laws governing the process. With such a knowledge base, fault diagnosis can be carried out with greater reliability. However, the diagnostic efficiency is affected by its detailed knowledge base, because the diagnosis system needs to explore the entire causal path from a failed component to the observed abnormalities.

To enhance both efficiency and reliability, a combination of the two approaches should be considered. There is a trend towards building fault diagnosis systems using both shallow and deep knowledge [5,8,14]. Rich and Venkatasubramanian [14] discuss a fault diagnosis system for a chemical process using both types of knowledge. They propose a two-tier architecture for integrating compiled and deep level knowledge in that the process specific compiled knowledge is stored at the top tier, while the lower tier holds deep knowledge. During diagnosis, the compiled knowledge is invoked first. If a diagnosis result cannot be obtained from the compiled knowledge, the diagnosis will drop down to the deep level knowledge.

To reduce the effort of encoding and debugging diagnostic heuristics from diagnostic experts, machine learning techniques can be used to automatically acquire the required knowledge. Recently several researchers have attempted to incorporate a learning mechanism into process fault diagnosis systems to make them more intelligent [3,4,9,10,11,13]. In Pazzani's approach [9,10], a set of initially developed heuristic rules are used to propose a hypothesis when an abnormal condition is encountered, and a deep model is then used to confirm this hypothesis. If it cannot be confirmed, then the heuristic rule which proposed this hypothesis is considered to have failed and it is revised by adding additional terms to its condition part to limit its applicability. This is called failure-driven learning since learning is initiated when a hypothesis failure occurs. Through this failure-driven learning, the existing heuristic rules can be refined but there may exist situations where there are no heuristic rules corresponding to some failures, especially failures which occur infrequently. In such situations, it would be desirable that the system can still diagnose the fault and learn a new heuristic rule. This is not addressed in Pazzani's approach [9,10]. Rich and Venkatasubramanian [11] discuss a causality-based failure-driven learning approach. In their approach, when a heuristic rule fails to propose the right hypothesis, the rule is revised and the system will drop down to deep knowledge based diagnosis, and it could learn a new heuristic rule. This method is developed for off-line diagnosis as can be seen from the context of [11,14]. The condition parts of some learnt heuristic rules include the negation of the failures of some other components, and this information is obtained from the operator. The aim of these failure-driven learning approaches is mainly to refine the existing heuristic rules.

Self-learning of diagnostic rules through inductive learning [2,6] are discussed in [3,4,13]. In this approach, a set of training data covering a large variety of potential faults is generated from simulation or obtained from past operating experience. The raw training data obtained are first transformed into symbolic form, such as shapes and curvatures. The training data are divided into positive examples and negative examples. Positive examples are the training data describing the behaviour of the process under a particular fault for which a rule is being learnt, while the rest of the training data are negative examples. Through inductive learning, it is required to find the characteristic descriptors, which are present in all the positive examples and may also appear in some of the negative examples, and discriminate descriptors, which occur in all the positive examples and none of the negative examples. This approach does not use any deep knowledge about the process and a large variety of training data are needed in order to produce desired rules.

One of the recently developed approaches in machine learning, Explanation Based Learning (EBL) [7], seems more suitable for integration with deep knowledge based diagnosis systems. In EBL, a model of the concerned domain is required and only one training example is needed to learn a concept. In deep knowledge based diagnosis, a deep model of the diagnosed process is available and, through explanation based learning, the requirement of a large variety of training data is avoided. EBL could also avoid the long training process which occurs in some other approaches of machine learning.

In this paper we describe an on-line fault diagnosis system which uses both deep knowledge and heuristics. During diagnosis, the system will first invoke the heuristic rules to propose a hypothesis. If a hypothesis can be proposed, then a deep model is used to discriminate this hypothesis. Otherwise, the diagnosis is based entirely on the deep model. The fault diagnosis system will test a set of candidate faults by inserting each fault as a disturbance to the qualitative model. The candidate which can explain the observed abnormalities is taken as the diagnosis result. If the successful diagnosis is not proposed by a heuristic rule, then the learning system will learn a rule corresponding to the diagnosed fault utilising the ideas of EBL. By such means, a diagnostic rule base could be gradually assembled and the diagnostic efficiency could be improved.

In the next section, diagnosis using both heuristics and deep knowledge is described. Section 3 describes the procedure of learning heuristic rules from deep knowledge based diagnosis. An illustrative application to the on-line fault diagnosis of a mixing process is presented in Section 4. The last section contains some concluding remarks and suggestions for further work.

2. FAULT DIAGNOSIS USING BOTH HEURISTICS AND DEEP KNOWLEDGE

The diagnosis system described here is based on a deep qualitative model of the monitored process and also uses heuristic rules which are learnt from successful diagnoses. The diagnosis system contains two parts: fault detection and fault diagnosis.

Fault detection is performed by checking if any constraints imposed by the qualitative model are violated and the fault diagnosis is initiated once the presence of a fault is detected.

When abnormalities are present in on-line measurements, the fault diagnosis system will predict the behaviour of the process, in the form of qualitative increments (increase, steady, and decrease) of certain measured variables over a period, through qualitative simulation and compare this with the observed behaviour and, if they are identical, then the process is considered to be at a normal condition. Otherwise, the measurements of several successive samples are taken to eliminate the effects of measurement noise. The expected behaviour of the process at these successive samples is predicted and compared with the observations. The presence of a fault is detected if, in the majority of these samples, the predicted behaviour does not coincide with the observations.

During the fault diagnosis phase, the diagnosis system will first try to use heuristic rules to propose a hypothesis which is to be confirmed by the deep model. The behaviour of the process under this hypothesis is predicted through qualitative simulation, where the qualitative model represents the process under the hypothesis, and is compared with the observation and, if they match, the hypothesis is confirmed. By the combined use of deep knowledge and heuristic rules the diagnosis system will enhance both efficiency and reliability.

The heuristic rules may be incomplete and, therefore, they may not propose a hypothesis in some cases. When such a situation is encountered, the diagnosis will be based entirely on the deep knowledge. A set of hypotheses will be formulated based on the patterns of model violations in the fault detection phase and tested on the qualitative models. The expected behaviour of the process under each hypothesis is predicted from the corresponding qualitative model and is compared with the observations. The hypothesis which can explain the observed abnormalities, in that the predicted behaviour under it coincides with the observations, will be taken as the diagnosis result. This diagnosis strategy is known as "hypothesis-test strategy" [8]. The qualitative models representing the process under different hypotheses should have sufficient resolution to distinguish among the hypotheses. After a successful diagnosis, which is not proposed by a heuristic rule, the diagnosis system will learn a rule corresponding to the diagnosed fault. This is described in the next section.

3. LEARNING DIAGNOSTIC HEURISTIC RULES

3.1 Explanation based learning

Explanation based learning is a recently developed approach to machine learning [7]. In this type of learning, the solution to a sample problem is generalised into a form that can later be used to solve conceptually similar problems. The generalisation process is driven by the explanation of why the solution worked. Knowledge about the domain allows the explanation to be developed and then generalised, thereby producing a general concept from the solution to a specific problem. One advantage of EBL systems is that they only require a small number of training examples to learn a concept. This is due to the fact that EBL systems possess the ability to explain what is relevant in these examples. The strength of EBL arises largely from its use of prior knowledge to quide its learning.

3.2 Learning diagnostic rules

If a successful diagnosis is not proposed by a heuristic rule, the learning system will learn a rule corresponding to this diagnosis. It will recognise the essential symptoms of a fault and construct a rule for this fault. The rules used are in the form:

$$\text{If } S_1 \,\&\, S_2 \,\&\ldots\&\, S_n$$
$$\text{THEN Fault}_i$$

which states that if symptoms S_1 to S_n are present then the ith fault occurs. The symptoms S_1 to S_n used here correspond to n different on-line information sources, which could be on-line measurements and controller outputs, and each symptom is considered to take one of the following values: increase, steady, decrease, and * which is a wild card and means that the corresponding symptom is not important. When applying a rule, * can match with any values. The main function of this wild card is to generalise a rule.

After a successful diagnosis, which is not proposed by a heuristic rule, the deviations in all the on-line information sources are calculated and converted into qualitative forms. These form a set of training data for the diagnosed fault. Since the fault may only affect certain on-line information and, therefore, the training example should be generalised to form a diagnostic rule. The generalisation process is performed by utilising the ideas of EBL.

As mentioned in the previous section, fault detection is performed by comparing the observed behaviour of the process with its prediction and a fault is detected when there exist discrepancies between observation and prediction. Fault diagnosis is performed by predicting the behaviour of the process under various hypotheses and the hypothesis which can eliminate the above discrepancies will be taken as a diagnosis result since it could explain the observed abnormalities. This deep knowledge based diagnosis could be used to explain which features in the training example are relevant to or important for the fault and which are not, and, hence, could be used to generalise the training example.

Suppose O is the set representing the observed behaviour of the process and P_n is the set representing the predicted behaviour under a normal operating condition. Then the discrepancies ($O - O \cap P_n$) are caused by a fault and represent the essential features of the fault, whereas the features represented in $O \cap P_n$ are considered to be irrelevant or not important in detecting this fault. In the explanation based generalisation, the symptoms corresponding to the on-line information involved in ($O - O \cap P_n$) are reserved and are termed "characteristic descriptors" of the fault, and the other symptoms, corresponding to the on-line information involved in $O \cap P_n$, are eliminated by assigning the value *. By such a means, the training example can be generalised to form a diagnostic rule.

The rules learnt above may be incorrect unless the characteristic descriptors for all the faults are mutually exclusive. For example, if the characteristic descriptors for the fault f_1 are S_1, S_2, and S_3, and those for the fault f_2 are S_1 and S_2, then the rules for f_1 and f_2 are

$$\text{IF } S_1 \,\&\, S_2 \,\&\, S_3 \text{ THEN } f_1$$

and

$$\text{IF } S_1 \,\&\, S_2 \text{ THEN } f_2$$

respectively. It can be seen that the second rule is more general than the first and, if S_1, S_2, and S_3 are present both rules are applicable. However, if the rules are used in a specific order and the more specific rule is put before the general rule, this problem could then be avoided. Another way to solve this problem is to put an additional term, which is the negation of the difference of the characteristic descriptors of the two faults, in the condition part of the more general rule. For example, the second rule above can be modified to

$$\text{IF } S_1 \,\&\, S_2 \,\&\, \text{not } (S_3) \text{ THEN } f_2$$

4. APPLICATION TO THE ON-LINE FAULT DIAGNOSIS OF A MIXING PROCESS

4.1 Fault diagnosis of a mixing process based on its qualitative model

The on-line learning method described above has been incorporated into a previously developed fault diagnosis system applied to a mixing process [15]. The pilot scale mixing process is shown in Figure 1, where two tanks in cascade receive hot and cold water supplies. Both streams enter tank 1 where mixing takes place, and the contents of tank 1 passes to tank 2 and subsequently out to the pool tank. Dïagnosis is performed based on measurements of level and temperature in both tanks. The level and temperature in tank 2 are controlled by a controller resident in a BBC-B microcomputer which also communicates with and is supervised by a DEC Microvax II host computer. The main supervision task is to find abnormal behaviour and diagnose the associated fault.

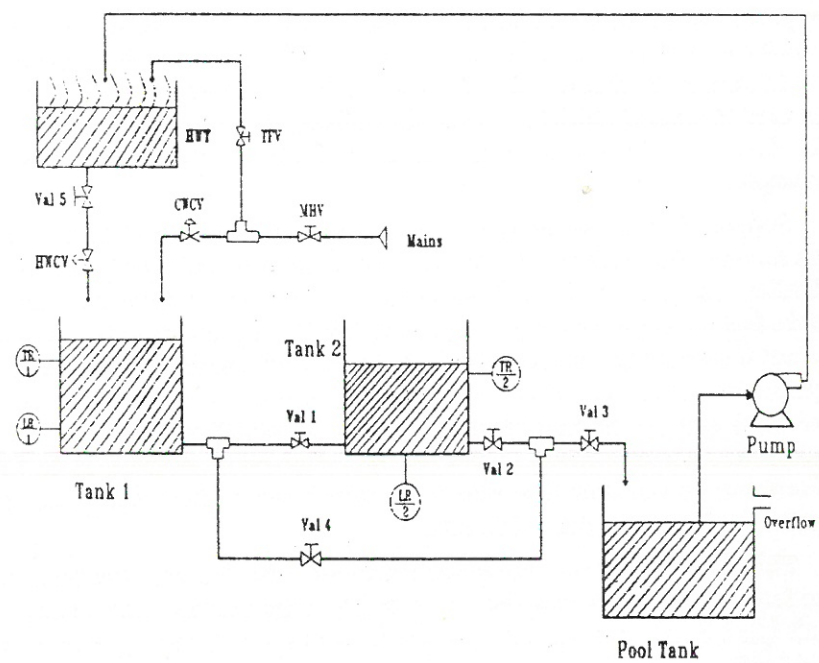

Figure 1. The Mixing Process

An on-line fault diagnosis system [15], which detects and diagnoses faults based on a qualitative model of the mixing process, has been developed. The model is constructed mainly based on de Kleer and Brown's confluence based qualitative reasoning method [1] and the model representing the process under normal operating conditions is listed below

$$\left[\frac{dH_1}{dt}\Big|_{t2}\right] = \delta_{1,2}Q_c + \delta_{1,2}Q_h \tag{1}$$

$$\left[\frac{dH_2}{dt}\Big|_{t2}\right] = \delta_{1,2}Q_{01} \tag{2}$$

$$\left[\frac{dT_1}{dt}\Big|_{t2}\right] = \delta_{1,2}Q_h - \delta_{1,2}Q_c \tag{3}$$

$$\left[\frac{dT_2}{dt}\Big|_{t2}\right] = \left[T_1 - T_2\right]\delta_{1,2}Q_{01} + \delta_{1,2}T_1 \tag{4}$$

$$\delta_{1,2}Q_{01} = \delta_{1,2}(H_1 - H_2) \tag{5}$$

where H1, H2, T1 and T2 are the levels and temperatures in tanks 1 and 2 respectively, Q_c and Q_h are the controlled flow rates of cold and hot water input streams respectively, Q_{01} is the output flow rate from tank 1 to tank2, $[dX/dt|_{t2}]$ denotes the qualitative value of dX/dt at time t2, $\delta_{1,2}$ X denotes the qualitative increment of X between the current time t2 and a previous time t1 when the process is at its steady state. This model may be modified in different ways to represent the process under different faults.

If abnormalities are present in on-line measurements, then the qualitative values of dH_1/dt, dH_2/dt, dT_1/dt, and dT_2/dt are predicted from Eqs(1) to (5) and are compared with the on-line measurements. If discrepancies occur then the measurements of another three sets of successive samples are taken to reduce the effects of noise. If discrepancies are detected in at least three out of the four sets of measurements, then it is considered that a fault occurred in the process. Otherwise, the process is considered to be at a normal condition.

Based on the patterns of model violations in the fault detection phase, a set of hypotheses, each assuming that a different fault occurred, can be formulated and discriminated through qualitative simulation. Given a hypothesis, the qualitative values of dH_1/dt, dH_2/dt, dT_1/dt, and dT_2/dt at the five successive samples are predicted from the corresponding qualitative model and are compared with on-line measurements. If the above discrepancies can be eliminated in at least three out of the four samples, then the hypothesis is considered to be able to explain the observed abnormalities and is taken as the diagnosis result.

4.2 Learning heuristic rules from deep knowledge based diagnosis

After a successful diagnosis, which was not proposed by a heuristic rule, the learning process is initiated. The deviations in the measurements of levels and temperatures and controller outputs to the cold and hot water control valves between the time when abnormalities are detected and a previous time when the process was at a normal state are calculated and converted into qualitative form (increase, steady, and decrease).

These are considered as symptoms and form a training example for the corresponding fault. This example will be generalised based on the deep model of the process to form a diagnostic rule.

Since the qualitative increments of levels and temperatures are predicted from the qualitative model and compared with the actual measurements in fault detection and diagnosis, the symptoms corresponding to levels and temperatures can be generalised by utilising the results in the fault detection phase. If a process variable's predicted qualitative increments are identical to those converted from the measurements in the fault detection phase, then it is believed that this variable is not important in detecting the fault and the symptom corresponding to this variable is generalised by assigning the value *. Controller outputs to valves are affected by the controlled variables which may be affected by the faults. In generalising the symptoms corresponding to controller outputs, it is considered that a controller output could not be affected by a fault if the symptom corresponding to this output has the value "steady", and the symptom is generalised by assigning it the value *.

4.3 Experiments

The learning system is developed and incorporated with the qualitative model based diagnosis system. Initially there is no heuristic rules and the diagnosis is performed based entirely on the deep model. The following failures, cold water control valve fails low, cold water control valve fails high, hot water control valve fails low, hot water control valve fails high, and partial blockages of hand valve 1 and hand valve 2, were initiated separately. After each diagnosis, a diagnostic rule was learnt for the corresponding fault.

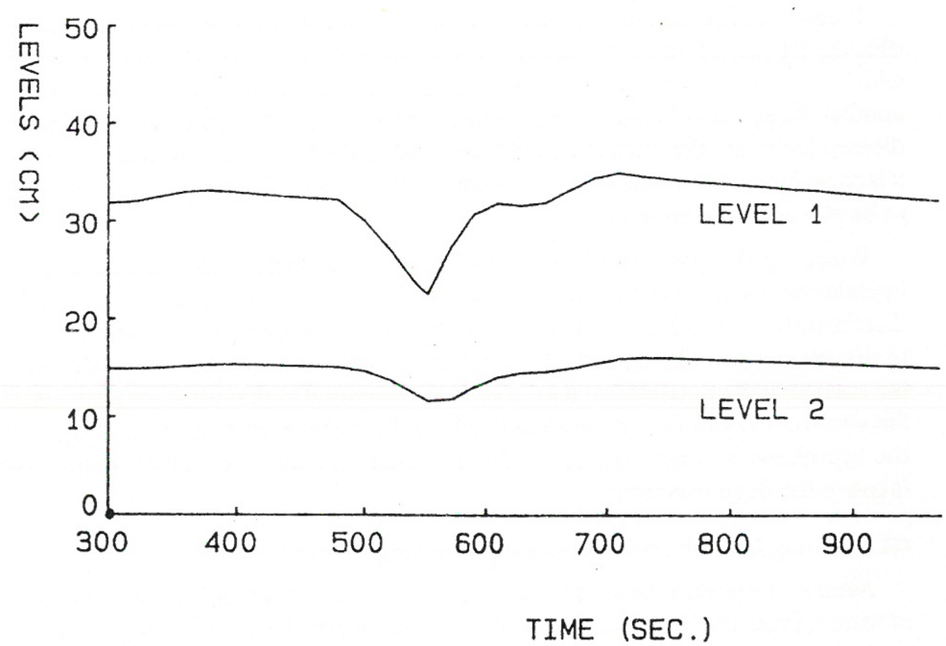

Figure 2. On-line Level Measurements

The on-line measurements and controller outputs covering the event when the fault "cold water control valve fails low" are provided in Figures 2 to 4. At 540 seconds, the diagnosis system detected abnormal measurements in levels and it then swiftly collected

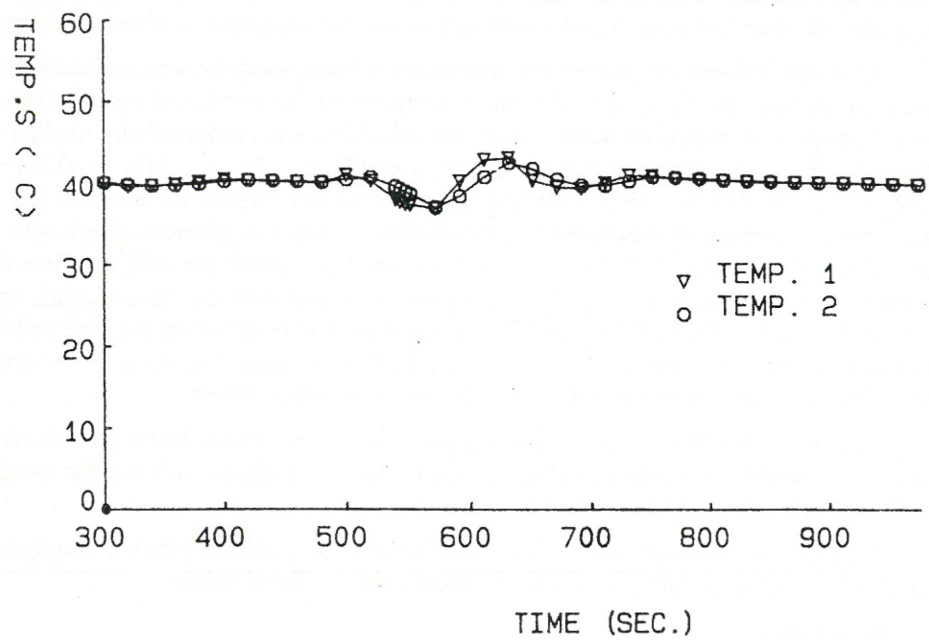

Figure 3. On-line Temperature Measurements

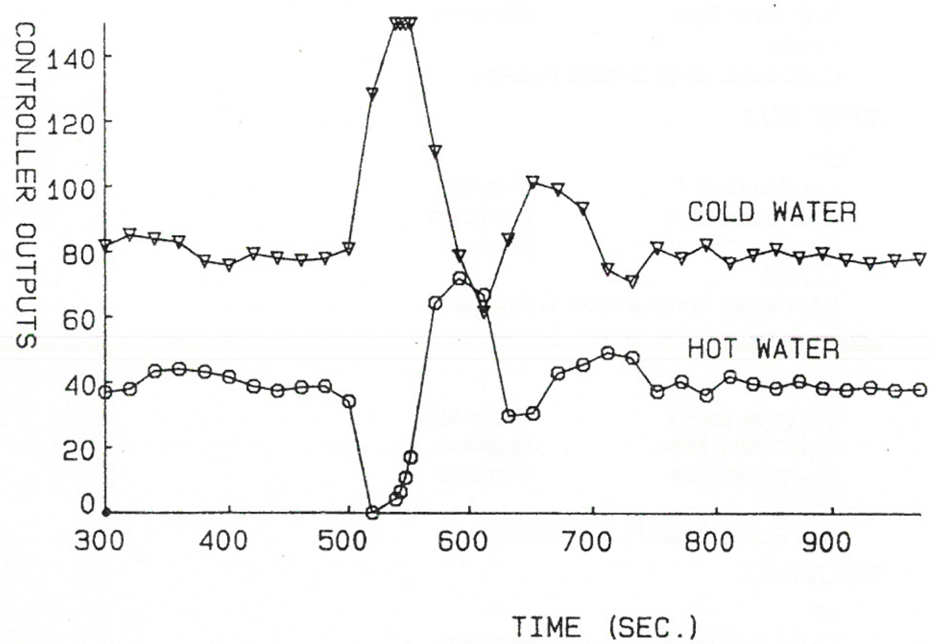

Figure 4. Controller Outputs to Control Valves

another three sets of measurements to confirm the fault detection. The fault detection was confirmed and the diagnosis system began to diagnose the fault. It first tried to propose a hypothesis from heuristic rules but no hypotheses could be proposed since there were no rules corresponding to this fault. It then diagnosed the fault by testing a set of hypotheses. Based on the patterns of violations of the qualitative model, a set of hypotheses were formulated and tested and the correct diagnosis result was obtained.

After the fault was diagnosed, the increments in levels, temperatures, and controller outputs between the time, when abnormal behaviour was detected, and a previous time, when the process was at its normal state, are calculated and converted into qualitative form. In this case, the qualitative increments are: $\delta T_1 = -$, $\delta T_2 = -$, $\delta H_1 = -$, $\delta H_2 = -$, $\delta Q_C = +$, and $\delta Q_h = -$, where + stands for increase and - stands for decrease. These form a training example for the fault. In the qualitative reasoning phase during diagnosis, it is found that the predictions for H_1 under normal conditions are different from the observations and these for T_1, T_2, and H_2 are consistent with the observations, and, hence, it is believed that T_1, T_2, and H_2 are not important in detecting the fault and the symptoms corresponding to them are eliminated by assigning the value *. A rule is therefore formulated and it is Rule No.1 in the list provided below.

After a rule has been learnt, it is compared with the previously learnt rules to see if it is more specific than others. If it is, then the rules are re-ordered such that the specific rule will be put before the general one.

After the above-mentioned six faults were all separately initiated, the learning system learned a diagnostic rule for each fault. These rules are listed below:

RULE NO.1

IF
Level in tank 1 Decrease
Cold water flow Increase
Hot water flow Decrease
THEN
Cold water control valve fails low

RULE NO.2

IF
Level in tank 1 Increase
Cold water flow Decrease
Hot water flow Increase
THEN
Cold water control valve fails high

RULE NO.3

IF
Temp. in tank 1 Decrease
Cold water flow Increase
Hot water flow Increase
THEN
Hot water control valve fails low

RULE NO.4

IF
Temp. in tank 1 Increase
Cold water flow Decrease

Hot water flow	Decrease
THEN	
Hot water control valve fails high	

RULE NO.5

IF	
Level in tank 1	Increase
Level in tank 2	Decrease
Cold water flow	Increase
THEN	
Hand valve 1 is blocked	

RULE NO.6

IF	
Level in tank 2	Increase
Cold water flow	Decrease
Hot water flow	Decrease
THEN	
Hand valve 2 is blocked	

It can be noticed that the rules are generalised in that their condition parts contain fewer symptoms than the training examples. After the rules had been learnt, the failures were initiated again with similar severities as their previous initiations and they were all proposed by the corresponding diagnostic rules. A fault may not be proposed by the corresponding heuristic rule if it is initiated with a quite different severity from that when the rule was learnt. In such situations, the learning system could learn a different rule for this fault.

5. CONCLUSIONS

Diagnosis using both deep knowledge and heuristic rules would be a desirable way to enhance diagnostic efficiency and reliability. Valuable shortcuts for diagnosis may be available in the form of heuristic rules. We have presented a method for learning heuristic rules from deep knowledge based diagnosis using the ideas of explanation based learning. Explanation based learning seems to be a suitable approach to be integrated with any deep knowledge based systems where the deep knowledge can be used to guide the learning process and this avoids the requirements of a large variety of training data. This may be suitable for developing an on-line fault diagnosis system for a new process where heuristic rules for diagnosis may not be available or for a complex process where the rules cannot easily be obtained. For such applications, a deep knowledge based diagnosis system is first developed and, after each diagnosis, the significant patterns in the on-line measurements are recognised and are generalised to form a heuristic rule. By this means the heuristic rule base can be automatically assembled.

To produce correct rules, the deep knowledge, upon which the learning system is based, should be correct. Some further research could be conducted to combine EBL with inductive learning or neural learning to address the problem that the deep knowledge is not completely correct. An initial investigation in combining EBL and neural learning is reported in [12]. The development of such an approach to the self-learning of diagnostic rules could be a future research topic.

The third part of the series will take a step forward in the control hierarchy by describing the application of combined numeric/knowledge based techniques in process co-ordination, an overall optimization of the operation of several process units in industrial complexes.

REFERENCES

1. DE KLEER, J. and BROWN, J. S., **A Qualitative Physics Based on Confluences,** in ARTIFICIAL INTELLIGENCE, Vol.24, December 1984, pp.7-84.
2. FORSYTH, R. and RADA, R., **Machine Learning: Applications in Expert Systems and Information Retrieval,** ELLIS HORWOOD LIMITED, 1986.
3. GUPTA, U. K. and ALI, M., **Hierarchical Representation and Machine Learning from Faulty Jet Engine Behavioural Examples to Detect Real Time Abnormal Conditions,** The First International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems, Tennessee, June 1-3, 1988, pp.710- 720.
4. KE, M. and ALI, M., **MLS, A Machine Learning System for Engine Fault Diagnosis,** The First International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems, Tennessee, June 1-3, 1988, pp.721-727.
5. LAPOINTE, J., MARCOS, B., VEILLETTE, M., LAFLAMME, G. and DUMONTIER, M., **BIOEXPERT - An Expert System for Waste Water Treatment Process Diagnosis,** in COMPUTERS AND CHEMICAL ENGINEERING, Vol.13, No.6, pp.619-630, 1989.
6. In R. S. MICHALSKI, J. G. CARBONEL and T. M. MITCHELL (Eds.), **Machine Learning: An Artificial Intelligence Approach,** TIOGA PUBLISHING COMPANY, 1983.
7. MOONEY, R.J. and BENNETT, S. W., **A Domain Independent Explanation-based Generalizer,** AAAI-86, Philadelphia, August 11- 15, 1986, pp.551-555.
8. MOORE, R. L. and KRAMER, M. A., **Expert Systems in On-line Process Control,** in Chemical Control III, Asilomer, CA. 1986, pp.839-867.
9. PAZZANI, M. J., **Refining the Knowledge Base of a Diagnosis Expert System: An Application of Failure-driven Learning,** AAAI-86, Philadelphia, August 11-15, 1986, pp.1029-1035.
10. PAZZANI, M. J., **Failure-driven Learning of Diagnosis Heuristics,** IEEE TRANSACTIONS ON SYSTEMS, MAN AND CYBERNETICS, Vol.17, No.3, May/June, 1987.
11. RICH, S. H. and VENKATASUBRAMANIAN, V., **Causality-Based Failure-driven Learning in Diagnostic Expert Systems,** AIChE JOURNAL, Vol.35, No.6, June 1989, pp.943-950.
12. SHAVLIK, J. W. and TOWELL, G. G., **An Approach to Combining Explanation-based and Neural Learning Algorithms,** CONNECTION SCIENCE, Vol.1, No.3, 1989, pp.231-253.
13. SUBRAMANIAN, C. and ALI, M., **Detecting Abnormal Situations from Real Time Power Plant Data Using Machine Learning,** The First International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems, Tennessee, June 1-3, 1988, pp.738-746.
14. VENKATASUBRAMANIAN, V. and RICH, S. H., **An Object-Oriented Two-Tier Architecture for Integrating Compiled and Deep-Level Knowledge for Process Diagnosis,** COMPUTERS AND CHEMICAL ENGINEERING, Vol.12, Nos.9/10, 1988, pp.903-921.
15. ZHANG, J., ROBERTS, P. D. and ELLIS, J. E., **Fault Diagnosis of a Mixing Process Using Deep Qualitative Representation of Physical Behaviour,** JOURNAL OF INTELLIGENT AND ROBOTIC SYSTEMS, Vol.3, No.2, 1990, pp.103-115.