Multiple Views of Knowledge in Diagnosis
Abu-Hakima, Suhayya

Publisher’s version / Version de l'éditeur:

Proceedings of the Knowledge-Based Systems and Robotics Workshop, 1993

Multiple Views of Knowledge in Diagnosis

Suhayya Abu-Hakima
Knowledge Systems Laboratory
Institute for Information Technology
Building M-50, Montreal Road
National Research Council of Canada
Ottawa, Canada
K1A 0R6
email: suhayya@ai.iit.nrc.ca
tel: (613) 993-8551
fax: (613) 952-7151
NRC #37097

Abstract

This paper argues that automated knowledge acquisition for diagnosis has had limited success in both failure-driven diagnosis and model-based diagnosis. The paper describes fault-based and model-based reasoning for diagnosis and surveys some of the approaches to knowledge acquisition in both areas. The Diagnostic Remodeller (DR) algorithm I am currently implementing for the automated generation of behavioural models from fault-based knowledge is presented. An example of fault-based knowledge from the Jet Engine Troubleshooting Assistant (JETA) is used to demonstrate how a behavioural model can be extracted with DR.

Fault-Based Diagnosis

Fault-based reasoning (FBR) is used in many diagnostic systems. Knowledge in FBR is largely based on maintenance manuals and interviews with experts intended to capture heuristic knowledge about the maintenance and repair of a device or process. The knowledge in these systems is often represented as hand-coded rules or frames which are organized into troubleshooting hierarchies. At the top level of the hierarchy is the general knowledge representing a problem with the device. This general problem is refined systematically until the leaf nodes of the hierarchy which represent physical repairs to the device are reached. Once these repairs are achieved by a human technician some diagnostic systems re-test to confirm that the symptoms and diagnosed faults are cleared through backtracking in the hierarchy.

FBR systems have evolved considerably since the development of MYCIN [Scott et al. 77]. MYCIN was developed to provide advice treatment for microbial infections. The MYCIN programs started with hand-coded rules which later evolved into meta-rules in NEOMYCIN to provide some structure to an otherwise flat knowledge base [Clancey 86]. The MYCIN approach remains a very widely used approach in FBR systems as described in the literature review of [Abu-Hakima 93].

Diagnosis is often referred to as a classification problem. Chandrasekaran and his colleagues developed MDX, a system that diagnosis a form of liver disease, cholestasis [Chanrasekaran et al. 79]. MDX has a diagnostic hierarchy which is referred to as a conceptual hierarchy since it guides the reasoner globally through diagnoses clustered as concepts that establish local contexts. Local uncertainties and hand-coded knowledge represented in frames are used to guide the diagnosis [Chandrasekaran and Tanner 86]. MDX has served as a model for many well-structured diagnostic systems including RATIONALE [Abu-Hakima 88] and JETA [Halasz et al. 92].

RATIONALE is a workstation diagnosis system that reasons explicitly so that it may support the user with sophisticated explanations of diagnoses that help justify system behaviour and clarify reasoning. In it many of the ideas advocated by Chandrasekaran for structuring FBR systems and handling uncertainty are applied. This approach was found to be ideal for explicitly representing knowledge so that it may be explained [Abu-Hakima and Oppacher 90]. RATIONALE diagnoses faults with Xerox workstations. It generates dynamic and static template-based explanations that include why, how and what-if responses. Explanation remains a major objective of FBR systems and most systems have why and how explanation but do not necessarily generate hypothetical (what-if) ones. RATIONALE's knowledge is in hand-coded frames.

The Jet Engine Troubleshooting Assistant (JETA) is a tool developed to assist a technician in diagnosing aircraft engines using a hypermedia interface which provides contextual help. For a diagnostic application to properly support hypermedia, one requires a structured manner by which to represent the knowledge, reason about it interactively, display it dynamically and explain it to the user (see [Abu-Hakima et al. 93] for a thorough description of JETA's hypermedia interface). JETA's knowledge representation and reasoning strategies are more flexible than those of other diagnostic systems including RATIONALE's.

JETA's troubleshooting knowledge is represented as a diagnostic network that is hierarchical in nature. Each node in the network corresponds to a decision point in the troubleshooting process and the links represent relations directing the flow of control between nodes. The overall network is much broader than it is deep since there are many components and associated symptoms. The number of nodes along a network path varies from four to twelve in a network of approximately 200 nodes. Possible next moves in the network are represented as children of a node. Any node can have multiple parents since a component malfunction may be due to many causes. The troubleshooting knowledge is hand-coded at each diagnostic node as a frame using a custom command language. In JETA as in RATIONALE, advice generating slots are included in the frame and their contents are output to the user as diagnoses or procedures to follow to find a fault. In JETA, advice is supported with a schematic or a graph. An indexed database of schematics and graphs is kept so that only pointers to the database are kept in the frame. The current implementation of JETA links text, graphs and schematics.

Model-Based Diagnosis

Model-based reasoning (MBR) for diagnosis concentrates on reasoning about the expected and correct functioning of a device. A device is modelled based on its components and their expected behaviour [Hamscher and Struss 90]. Such models range from quantitative ones to qualitative ones and all attempt to approximate device behaviour as accurately as possible. Once a device model is stabilized then a device's observed behaviour can be predicted from the model. If a discrepancy in behaviour is detected then possible candidates based on assumed component faults are generated. These candidates are generated based on assumptions that describe correct model behaviour. Sequential diagnosis is used to choose observations, augment a prediction for the candidate faults and update the list of candidates until a dominant candidate is found.

In MBR there are many conflicting definitions for models. They range from causal models represented as semantic networks with links specifying the relations between component nodes to full blown numerical simulations for complex systems and processes that have taken decades to perfect. Generating models is a key problem in MBR. Some researchers generate causal models, others generate models with structure and behaviour while others generate functional models for devices. Knowledge in models has thus far been hand-coded by experts that understand device component behaviour and function.

Davis was one of the earlier proponents of MBR. In [Davis 84] he describes a theory to exploit reasoning on the basis of device structure and behaviour. He defines paths of causal interpretation. He also describes constraint suspension used to identify which components are responsible for which faults. He argues that we need to balance complexity versus model completeness in diagnosis thus we need to enumerate and layer categories of failure. Quite a bit of work has followed Davis' examples and theories.

De Kleer and Williams published a key paper on MBR for diagnosis describing GDE, the General Diagnostic Engine [de Kleer and Williams 87]. GDE infers behaviour from device structure and functionality. It is applied to digital circuits and makes use of an ATMS (Assumption-Based Truth Maintenance System). This work forms the cornerstone of ATMS-based model-based reasoning systems. It was followed by many papers that criticized the approach as not computationally practical in diagnosing faults with large complex systems. Some of the papers criticizing GDE propose the use of hierarchical fault-based reasoning to reduce the computational complexity of de Kleer and Williams' approach. Struss has developed GDE+ which handles: simple dynamic aspects, multiple tests, hierarchical knowledge and unreliable observations [Struss 89]. GDE+ is a migration back to heuristic or empirical diagnoses using fault-based reasoning. Struss points out that neither GDE nor GDE+ address: changing device structures, complex temporal behaviour (feedback), uncertainty or the use of qualitative models in reasoning. In [Struss and Dressler 89] the authors advocate the representation of a fault view for each component. They point out that a fault and a healthy view (state) for a component cannot be true in the same time instant (consistent belief rule). They also give the 'no good inference rule' where the node and its opposite which represents a fault cannot be true at the same instant. The ATMS is then modified to reason with the fault as well as the no-fault behaviour of a device. Their work gives excellent insight into combining model and fault-based diagnosis to deal with GDE's shortcomings.

Other MBR authors have argued about the definition of device functionality versus behaviour. Sticklen in [Sticklen et al. 88] describes modelling a device's functionality by:

- decomposing the device into sub-devices,
- stating abstractly the functions, goals and purpose of the device and
- representing the manner of achieving the device functions, goals and purpose.

A good definition of functionality is one which argues that function is the set of goals the device is intended or designed to achieve [Malin and Liefker 91].

Automatic Diagnostic Knowledge Acquisition

Machine learning is a key approach in knowledge acquisition for diagnosis. Machine learning includes empirical and analytic learning. Empirical learning focuses on learning for classification (including learning rules from real or simulated data for diagnosis). Analytic learning addresses learning for problem solving tasks which include planning, diagnosis, design, natural language understanding, control and execution. There has been an explosion of work in machine learning in recent years. It is viewed as one of the key approaches of reducing the knowledge acquisition bottleneck [Boose 91; Gaines and Shaw 91].

The MOLTKE (MOdels, Learning and Temporal Knowledge in Expert systems) testbed for diagnosis under development at the University of Kaiserslautern, Germany is described in [Althoff et al. 90]. The system is designed to acquire device knowledge for diagnosis. It has an MBR mechanism for acquiring device models based on their components. A *component* of the model includes a name, ports to other components (with optional test costs), possible internal states (with optional test costs), behaviour of the component (either in state tables or rules that represent the constraints the component sets up between its ports and states), subparts and their interconnections (if the component is non-atomic), typical malfunctions with name and effects (model typical behavior when the component fails) and a priori probability of failure). MOLTKE uses case-based reasoning to acquire and refine knowledge that is generalized to a fault-based hierarchy. It also uses explanation-based learning to refine the rules in the fault-based hierarchy to get the minimum reasoning paths for a solution. MOLTKE has been applied to a Computerized Numerical Control (CNC) machining center. It is also under investigation for the problem of driving mining machines.

Carnes and Fischer describe the use of supervised learning for the placement of sensors in a thermal model [Carnes and Fisher 92]. They use cluster analysis using COBWEB to group together observations for diagnosis. An explicit model of diagnosis (a model-based simulation) is used to place the sensors by generating the various states that need to be measured for design or diagnosis. The approach is novel however sensor placement is not a new issue tackled by machine learning [Abu-Hakima 93].

Scotty is an expert system under development at Rocketdyne for the Space Shuttle Main Engine (SSME) [Modesitt 90]. It is aimed at automating the analysis carried out on firing the shuttle engines. The FBR expert system is based on the experience of propulsion engineers and it consists of 125 manually derived rules and 1400 automatically induced rules from example runs which are embedded in a distributed software environment. The rules are poorly structured and will be moved to a relational database management system (RDBMS) to improve the knowledge representation. The author is also planning to enhance Scotty's graphics and generalize the approach to other Rocketdyne engines. Rule induction from data as a technique is not new but this integration with FBR knowledge and its application is somewhat novel.

ACES (Attitude Control Expert System) diagnoses anomalies in the attitude control system of the DSCS-III satellite [Pazzani 90]. ACES is fault-based (rules represented as Prolog predicates). A fault is confirmed or denied by comparing the observed behavior to that predicted with a simulator. In the case where the simulation denies the fault, the heuristic that proposed the fault is expanded to include the tests that the simulator performed to rule out the fault. ACES uses explanation-based learning to identify the conditions under which the heuristic will propose a fault that is denied. The author concludes that failure-driven learning finds sufficient conditions for ruling out a fault and success-driven learning finds sufficient conditions for establishing a fault (but not necessarily ruling others out). Pazzani's work is novel and very relevant to the refinement of fault-based knowledge using model-based reasoning and explanation-based learning.

Pearce describes two parallel approaches to acquiring knowledge for FBR diagnosis [Pearce 88]. The first is generated by a knowledge engineer who hand codes rules form an expert. The second uses examples of failure from a simulation and applies AQ, Quinlan's induction algorithm, to the data. The author concludes that the second approach gives much better results. The only drawback to the use of induction of rules from simulated data is that one needs a good simulation of a device to generate the examples from which the expert system is induced. It is not always possible to find or generate such simulations.

The collection of jet engine sensor data and its interpretation while comparing multivariate linear regression and instance-based learning is described in [Turney and Halasz 92]. The results of the study indicate that instance-based learning can be the basis for a useful diagnostic tool for aircraft engine technicians. The application of instance-based learning for generating libraries useful in jet engine diagnosis is novel. Such an algorithm could be integrated with the Jet Engine Troubleshooting Assistant [Halasz et al. 92] to provide a technician with an on-line sensor monitoring and jet engine diagnostic tool.

There has been tremendous activity in machine learning in recent years. In empirical learning classification algorithms such as ID3 and AQ have been used to induce diagnostic rules from real or simulated data. Classification learning extracts rules from positive and negative examples. In analytic learning explanation-based learning (EBL) has been used in the form of speedup learning to generalize diagnostic rules and shorten reasoning chains. I believe that neither classification nor EBL addresses the problem of *knowledge-rich* learning where structured knowledge is learned. Such rich knowledge would result from learning to produce hypothesis hierarchies such as those described in fault-based reasoning. In addition, learning from structured knowledge to produce new knowledge, such as learning a device model from its fault hierarchy has not been addressed. Learning complex structures especially for diagnosis is by no means an easy problem but it is one that needs to be further addressed by a combination of researchers in both the machine learning and diagnosis fields. Some researchers which have combined learning (empirical or analytic) with FBR and MBR have met with more success as exemplified by some of the complex problems above. I believe that the key to resolving the knowledge acquisition bottleneck in diagnosis lies in the integration of various machine learning algorithms with partially hand-coded knowledge bases used in FBR or MBR.

Automatic Generation of Behaviour from Fault Knowledge: the *DR* algorithm

Hypothesis

Humans use failure-driven reasoning for successful device diagnosis and repair. As humans reason they build primitive mental models of the devices they diagnose and repair. The hypothesis for the *DR* algorithm is that knowledge of failure and repair embodied in most structured diagnostic knowledge-based systems can be used to derive rudimentary device models. The *DR* algorithm will extract rudimentary device models from fault knowledge. The device model represents structure and behaviour and is driven by device functionality.

Motivation

A great deal of effort is expended hand-coding complex knowledge bases for diagnostic FBR. The artifacts these diagnostic systems are developed for are often expensive machines which have been designed and continuously modified so that no existing accurate schematic or design of their behaviour remains. The J85-CAN-15 is a jet engine which is the first application of JETA. The J85-CAN-15 engine was designed in the 1950's and has easily had at least one modification a year since its launch. As a result of modifications and stresses of daily use (flying in the arctic and flying in desert heat) the jet engine is a very different device than was originally designed and sometimes displays inexplicable behaviour. No existing design schematics can completely capture the engines's behaviour. It is also a very difficult device to diagnose. For these reasons a tool such as JETA was developed. As is typical with FBR systems, JETA does not diagnose novel faults. Learning the device component model, its behaviour and functionality using the FBR knowledge provides the technician with a tool that can achieve model-based diagnosis. For these reasons it was concluded that the *DR* algorithm should be implemented.

Background

If we follow the de Kleer [de Kleer and Williams 87] approach which represents a device with functionality as a set of components with behaviour. The device can be diagnosed by assuming a faulty component and enumerating the behavioural states that the fault propagates in the remainder of the device. This is compared to the behaviour that a technician is observing in attempting to isolate a problem. Model-based diagnosis can detect novel faults since the behaviour of the device is the basis of its knowledge representation and reasoning. Fault-based reasoning uses the faults of a device rather than its actual behaviour, hence it cannot detect novel faults. However, model-based reasoning can lead to a combinatorial explosion in producing a diagnosis for complex systems (for example, an aircraft engine) and it does not lend itself to causal explanation.

I am currently implementing the *DR* algorithm intended to address the automatic generation of a functional model of a device from its fault knowledge. That implies the automatic generation of MBR knowledge from FBR knowledge. By extracting a functional model both fault and model-based diagnosis can be pursued in a single system gaining from the advantages of the two approaches while minimizing the disadvantages. The *DR* algorithm is being applied in the area of complex electromechanical devices, specifically jet engines.

Objectives of DR Algorithm

DR is an algorithm that would take as input the fault knowledge of a device. It may also be necessary to take as input some background knowledge related to the device to attempt to learn its full component structure and connectivity. *DR* initially extracts from the fault knowledge base all references to device components and subsystems. Given these components the algorithm backtracks through a diagnostic hierarchy of nodes to generate hypotheses for component connectivity. To further establish component connectivity, *DR* examines symptomatic or parametric knowledge that activates the diagnostic nodes. Symptomatic knowledge is knowledge of device failure which can be used to generate hypotheses about correct device function. This knowledge will be used to derive behavioural knowledge between components.

Approach to the *DR* Algorithm

The top level design of the *DR* algorithm is shown in Figure 1. Two phases clearly divide the design of the algorithm. In the first phase, I propose the use of an existing knowledge base that diagnoses a complex electromechanical system as input to *DR*. The Jet Engine Troubleshooting Assistant (JETA) is a system implemented to diagnose faults with aircraft engines [Halasz et al. 92]. Background knowledge referred to as a zero-based model of the device will be the second input. This background knowledge is general in nature, for example it could include knowledge that a pump delivers some liquid from a sink to a source and needs a pressure to increase or decrease the flow of liquid. It could also include some knowledge about feedback control in moderating the flow of a liquid to a source based on the level of the liquid at the source. It may also include some relational knowledge about Pressures, Temperatures and Volumes as defined in Physics.

To achieve the knowledge-rich learning proposed as the output for *DR* one requires the use of a structured and explicit knowledge representation that can adequately represent diagnostic causality. I propose to use the RATIONALE [Abu-Hakima 88] knowledge representation which is explicit and provides diagnostic causality for the representation of the learned model. The second phase of the *DR* algorithm compares the functional knowledge of a device exemplified in its device model of components and their behaviours with the original fault knowledge. The purpose of the second phase is to look for inconsistencies and gaps in the knowledge. The gaps discovered in the fault knowledge could then be used to diagnose novel faults.

The objective of the *DR* algorithm is to discover and refine a behavioural model of a device whose fault knowledge is represented. In the most general sense the algorithm must identify the components of the device, generate links between those components and generate hypotheses for the behaviours between the components. To achieve this the *DR* algorithm must:

1. identify the end nodes (components) in hierarchy
   - -these are nodes that have no child or sibling refinements
   - -these are nodes explicitly labelled as replace or repair-type in JETA as they explicitly refer to device components
   - -tag appropriate nodes as physical components based on semantic information
2. identify the end nodes in the diagnostic hierarchy related to subsystem to be modelled
   - -perform a match with the name or any acronym that could match the subsystem
3. identify the inheritance (direct) links between nodes
   - -backtrack from end node to parent node and tag
   - -tag shared parents of a node
   - -tag siblings of a parent
4. identify the indirect links between nodes
   - -tag nodes that activate other nodes without a familial link
   - -tag nodes that activate other nodes through the evaluation of parameters
5. hypothesize about the relations (behaviours) between the nodes
   - -cluster nodes related by direct or indirect links
   - -use semantic information (node names or parameter ranges/values) to relate nodes
6. output the device model for verification to the user
   - -map out the identified components of the subsystem
   - -relate the components through the hypotheses

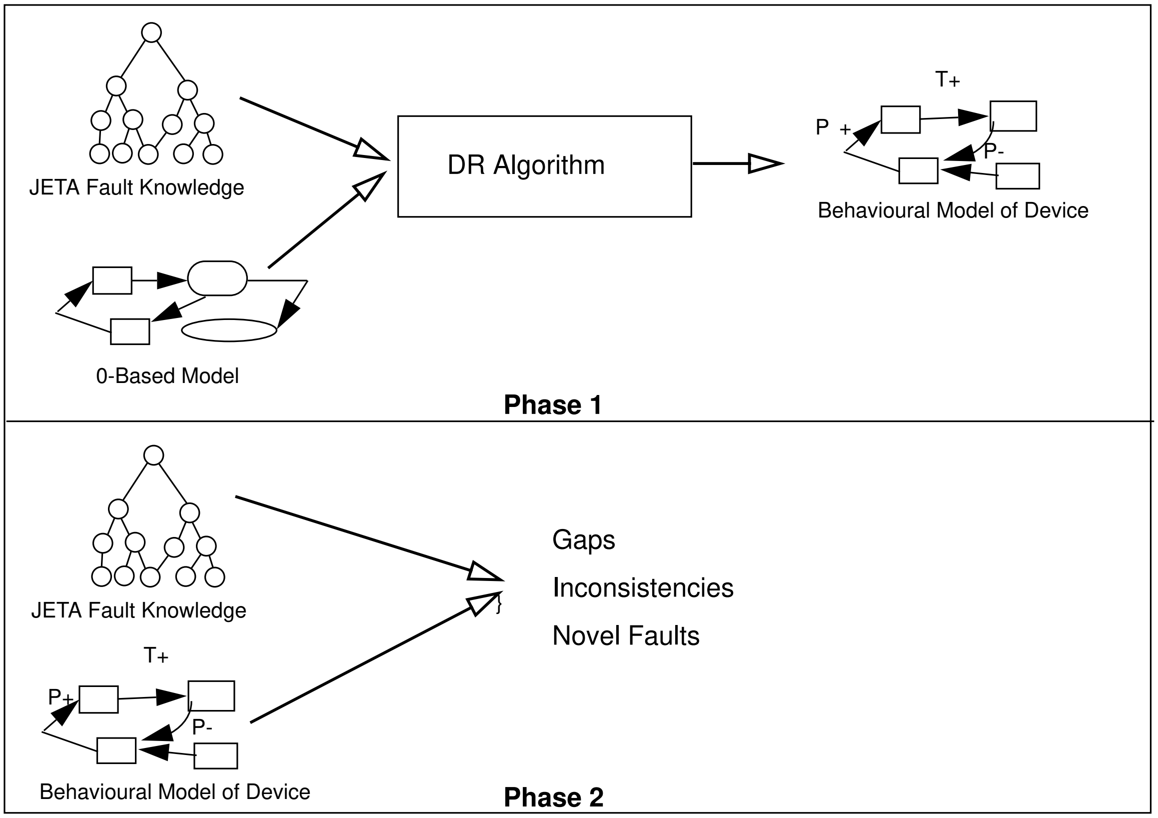

FIGURE 1. *Diagnostic Remodeller (DR) Algorithm Design*

Extracting a Model from JETA's Fault Knowledge: An Example

An analysis of the JETA fault knowledge shows layers of knowledge which can be visualized as diagnostic trees. The topmost layer is an entry point to jet engine faults and subsequent layers organize the faults into various branches. The phases of operation branches lead to various symptomatic nodes labelled as snags. These snags in turn are refinable down to repair and replacement nodes which form the bottom layer of the diagnostic hierarchy[1]. If one examines the knowledge encoded in these bottom nodes more closely one discovers that they represent faults directly on physical engine components. These physical component fault nodes can be grouped into those affecting one of thirteen subsystems by their nomenclature. One can follow the six steps of the *DR* algorithm introduced above to discover the behavioural network for the main fuel system of the jet engine:

1. The diagnostic hierarchy is referred to as a network since it includes relations that are not directly inherited that allow the JETA reasoner to jump around between nodes thus forming more of a network than a hierarchy.

Step 1

One can identify 9 replace nodes through the JETA node frame slot 'node-type'.

Step 2

If one takes a specific subsystem, the MFS (Main Fuel System), one can extract the names of 5 fuel system replacement nodes by pattern matching with the node nomenclature *N-MFS-XXX (this is an internal representation that was used by the knowledge engineer to distinguish between nodes):

1. main fuel control (MFC)
2. overspeed governor for MFC (OSG)
3. MFC/ABFC signal line
4. main fuel pump supplying MFC
5. a pressurizing and drain (P+D) fuel valve

Step 3

For each of the 5 replacement nodes direct single and shared inheritance links can be traced, for example:

- the MFC and MFC/ABFC signal line nodes share the parent node no hesitation at acceleration area

- the P+D valve and main fuel pump share the engine flameout parent node
- the main fuel control (MFC) delivers fuel to the engine fuel nozzles

Step 4

Indirect links between the subsystem nodes can be identified by activation rules and evaluation of parameters.

Step 5

A causal topological network can be the basis for hypothesized component-behaviour relations. Nodes are clustered based on direct or indirect links. They can also be clustered based on semantic knowledge such as names and ranges of parameter values.

Step 6

The hypothesized device model that includes component enumeration and behaviour can be output to the user for verification. It can also be tested by 'breaking' the correct modes of operation as described for Phase 2 of the *DR* algorithm.

As a device model of hypotheses that link JETA's components is formed it is important to explain *DR*'s results. This can be achieved by mapping the newly learned components and behaviours into a set of RATIONALE hypotheses so that the explanation strategies can be output. This can be achieved by mapping JETA's knowledge representation for the causal component nodes to RATIONALE's.

Issues

There are five main issues that I will need to answer in the implementation of the *DR* algorithm. The first is what is the exact form of the learned model when some or no background knowledge is used. If no background knowledge is used is the model much more than a causal model rather than a component-behaviour model? The latter model is required to diagnose novel faults. Also, what is the minimum background knowledge required to learn a device model from FBR knowledge. That can be extended to explore whether background knowledge alone provides one with a device model. Can this model in turn be used to derive FBR diagnoses?

A traditional simulation of a device or process relates to a device model. A key question is what is needed in addition to the simulation to generate a device model or a FBR model which in turn can be used for diagnosis?

The *DR* algorithm will be implemented as a general algorithm useful in generating models for devices other than jet engines. It is not obvious that its background knowledge will make it specific for generating a particular model. However the FBR knowledge used as input will make it specific to generating a model for a particular device. One question could be whether or not the algorithm can be generalized further so that it diagnosis abstract (e.g. software) versus physical (e.g. jet engine) systems.

Finally, the *DR* algorithm requires a highly structured FBR knowledge base. One key question is what criteria will allow it to extract a causal model from a rule- versus a frame-based fault model.

Conclusion

This paper argues that automated knowledge acquisition for diagnosis has had limited success in both failure-driven diagnosis and model-based diagnosis. The *DR* algorithm for the automated generation of causal models from fault-based knowledge is introduced. An example of fault-based knowledge from the Jet Engine Troubleshooting Assistant (JETA) is used to demonstrate how a causal model of the main fuel system of a jet engine can be extracted with *DR* from the fault knowledge.

References

Abu-Hakima, S. [1988], *RATIONALE: A Tool for Developing Knowledge-Based Systems that Explain by Reasoning Explicitly.* Masters thesis, Carleton University, Ottawa, Canada. May 1988.

Abu-Hakima, S., and Oppacher, F. [1990], Improving explanations in knowledge-based systems: RATIONALE. Knowledge Acquisition journal (2). 301-343. December 1990.

Abu-Hakima, S. [1993], Diagnostic techniques in knowledge-based systems: a review of approaches, applications and issues. 100 p. To be published as an NRC ERA (1993).

Abu-Hakima, S., Halasz, M., and Phan, S. [1993], An Approach to Hypermedia in Diagnostic Systems. To appear as Chapter 10 in Intelligent Multimedia Interfaces, Edited by Mark Maybury, AAAI Press (1993).

Althoff, K-D., Maurer, F., and Rehbold, R. [1990], "Multiple knowledge acquisition strategies in MOLTKE." *Current trends in knowledge acquisition*, Published by IOS, Amsterdam, Netherlands. pp. 21-40, (1990).

Boose, J.H. [1991],
"Knowledge acquisition tools, methods and mediating representations", Knowledge Acquisition for Knowledge-Based Systems, Edited by H. Motoda, R. Mizoguchi, J. Boose and B. Gaines. Published by IOS Press, the Netherlands. pp. 25-62, (1991).

Carnes, J.R. and Fisher, D. [1992],
"Inductive learning approaches to fault isolation and sensor placement", *Notes of the Third International Workshop on the Principles of Diagnosis*, October 12-14, Rosario, Washington pp. 132-138, (1992).

Chandrasekaran, B., Gomez, F., Mittal, S. and Smith, J.W. [1979],
"An approach to medical diagnosis based on conceptual structures". Proceedings of the Joint Conference on Artificial Intelligence, 134-142.(1979).

Chandrasekaran, B. and Tanner, M.C. [1986],
"Uncertainty handling in expert systems: uniform vs. task-specific formalisms". In Kanal L.N. and Lemmer J., editors, Uncertainty in Artificial Intelligence, North Holland Publishing.(1986).

Clancey, W.J. [1986],
"From GUIDON to NEOMYCIN and HERACLES in twenty short lessons: ORN final report 1979-1985". The AI Magazine, pp. 40-60,(1986).

Davis, R. [1984],
"Diagnostic reasoning based on structure and behaviour", Artificial Intelligence, Vol. 24, pp. 347-410 (1984).

de Kleer, J., Williams, B.C. [1987],
"Diagnosing Multiple Faults", *Artificial Intelligence*, vol. 32, (1987).

Gaines, B. and Shaw, M. [1991],
"Foundations of Knowledge Acquisition", *Knowledge Acquisition for Knowledge-Based Systems*, Edited by H. Motoda, R. Mizoguchi, J. Boose and B. Gaines. Published by IOS Press, the Netherlands. pp. 3-23, (1991).

Halasz, M., Davidson, P., Abu-Hakima, S., and Phan, S. [1992],
JETA: A Knowledge-based Approach to Aircraft Gas Turbine Engine Maintenance. Journal of Applied Intelligence, 2, pp. 25-46, Kluwer Academic Publishers, NRC 31832, (1992).

Hamscher, W. and Struss, P. [1990],
"Model-Based Diagnosis", *AAAI-90 Tutorial Notes*, Received at the eighth national conference of artificial intelligence, Boston, Massachusetts. July 29, 1990. pp. 1-179, (1990).

Malin, J.T., and Liefker, D.B. [1991],
"Representing functions/procedures and processes/structures for analysis of effects of failures on functions and operations", Proceedings of 1991 NASA Goddard Conference on Space Applications of Artificial Intelligence, Greenbelt, Maryland, pp. 141-151, (May 13-15 1991).

Modesitt, K.L.[1990],
"Inductive knowledge acquisition experience with commercial tools for space shuttle main engine testing", *Proceedings of the fifth conference on Artificial Intelligence for space applications*, NASA Marshall Flight Center, pp. 487-496, (May 1990).

Pazzani, M.J. [1990],
"Learning in order to avoid search in logic programming", *Computers & Mathematics with Applications*, 20(9-10), pp. 101-10, (1990).

Pearce, D.A. [1988],
"The Induction of Fault Diagnosis Systems from Qualitative Models", *AAAI-88, Seventh National Conference on Artificial Intelligence*, Saint Paul, MN, (21-26 August 1988).

Scott, A.C., Clancey, W.J., Davis, R. and Shortliffe, E.H. [1977],
"Explanation capabilities of production-based consultation systems". Stanford University, Dept. of Computer Science, Report No. STAN-CS-77-593,(1977).

Sticklen, J., Chandrasekaran, B., and Bond, W.E. [1988],
"Distributed causal reasoning for knowledge acquisition: a functional approach to device understanding", 3rd AAAI sponsored Knowledge Acquisition Workshop for Knowledge-Based Systems, Banff, Canada, pp. 34-1 to 34-18, (1988).

Struss, P. [1989],
"New Techniques in Model-based Diagnosis", *Proceedings of Knowledge-based Computer Systems*, Bombay, India, (11-13 December 1989).

Struss, P., Dressler, O. [1989],
"Physical Negation - Integrating Fault Models into the General Diagnostic Engine", *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI-89)*, Detroit, MI, (20-25 August 1989).

Turney, P. and Halasz, M. [1992],
"Instance-based learning and multivariate linear regression applied to J85 transient performance data", *Internal Report, Knowledge-Systems Laboratory, IIT, NRC*, pp. 1-25. (1992).