Automated Decision-Analytic Diagnosis of Thermal Performance in Gas Turbines

To be presented at 1992 ASME Turbo Expo, Cologne, Germany, June 1992.

John S. Breese
Eric J.Horvitz
Mark A. Peot
Knowledge Industries
125 California Ave.
Palo Alto, CA 94306

Rodney Gay
Enter Software
855 Oak Grove Ave.
Menlo Park, CA 94025

George H. Quentin
Electric Power Research Institute
3412 Hillview Ave.
Palo Alto, CA 94304

Abstract

We have developed an expert system for diagnosis of efficiency problems for large gas turbines. The system relies on a model-based approach that combines an expert's probabilistic assessments with statistical data and thermodynamic analysis. The system employs a causal probabilistic graph, called a belief network, to update the likelihoods of alternative faults given information about diverse classes of information. In response to any subset of findings or reported observations, the system suggests the most cost-effective tests to perform to determine the source of a performance problem. We discuss the decision-analytic methodology that underlies the development of the system and present results of an initial version of the system. Finally, we discuss future planned development and evaluation, toward the ultimate goal of applying the system in the day-to-day maintenance of gas-turbine power plants.

A Overview

For several years, there has been considerable interest in computer-based tools for assisting people to diagnose difficult problems with power generation, medicine, aerospace, and manufacturing. In 1991, the Electric Power Research Institute (EPRI) initiated a project with the goal of developing expert systems for diagnosing efficiency-related problems in large gas and oil-fired turbines. The management of operating efficiency for turbines has traditionally received relatively little attention. Nevertheless, the appropriate diagnosis and maintenance of turbine efficiency can provide significant cost savings under many operating regimes. Moreover, the effective diagnosis and treatment of efficiency problems can serve to curtail the development of more serious operating problems requiring expensive overhauls and prolonged down time.

There have been other efforts to build expert systems for the diagnosis of problems with power-generating facilities. The project is innovative in that we have developed an expert system for diagnosing turbine efficiency problems based on a synthesis of automated probabilistic reasoning and a more traditional thermodynamic model. The probabilistic reasoning component of the system relies a belief-network representation. A *belief network* is a causal network that enables the representation of general probabilistic dependencies among disorders and relevant observations. The belief network is central to two main phases of the expert system's operation. First, we use the belief network to assign uncertainty to alternative explanations of input data. Second, we consider the uncertainty in alternate hypotheses in value-of-information calculations to determine the next best tests to perform.

Interest has been growing in belief networks in the community of artificial-intelligence (AI) investigators. Belief networks have been applied recently to capture the subtleties of expertise in several application areas (Henrion et al., 1991). For example, large belief-network knowledge bases have been developed for several medical diagnosis specialities. However, there has been little previous work on the use of belief networks for diagnosing turbine efficiency problems. Moreover, there have been no previous attempts to employ belief networks for analysis of problems that hinge, in part, on the solution of complex mathematical models to predict how key input variables affect output variables. Our emphasis on diagnosis of efficiency problems, as opposed to start-up or operating malfunctions as in other expert systems, leads to the need for the system to have considerable analytic knowledge about the thermal cycle of the turbine. Thus, we could not avoid the task of integrating detailed thermodynamic analytical models with probabilistic reasoning.

We shall review issues surrounding the handling of uncertainty in expert systems. We then introduce the use of belief networks, an expressive representation for capturing probabilistic dependencies among observations and disorders. Finally, we shall describe the current status of the project and review the operation of the system.

B Uncertainty and Expert Systems

With few exceptions (Klempner et al., 1991), most projects to date on the construction of expert systems for problems with power generation have relied on rule-based reasoning methods. Rule-based expert systems date back to the earliest expert systems. These systems are referred to as *rule-based* because inference is performed through the logical chaining of *IF–THEN* rules, acquired from an expert. These systems typically assume *certainty* in the relationships among observations and disorders. For analyzing complex systems, assumptions of determinism between observations and disorders are typically invalid. Uncertainty is ubiquitous in reasoning about faults in complex systems given a small set of initial observations. Typically, we cannot directly inspect an internal fault without incurring great costs. Thus, we cannot easily determine with certainty the presence or nature of a fundamental cause. Rather, we must reason with uncertainty about the associations between observables, and alternative disorders that can cause those observations.

Considering probabilistic relationships is crucial for considering the relative likelihood of alternative hypotheses, for identifying the next best tests to make, and for determining ultimately what action should be taken (i.e., when the expected cost of an unscheduled shutdown is dominated by the benefits of exploratory or preventative maintenance). Conversations with expert diagnosticians quickly reveal the importance of uncertain relationships among faults and observations in the power generation systems. For example, a damaged injector may lead to a out-of-bounds reading for the average exhaust gas temperature. However, the problem may also be caused by some other primary disorder. Indeed, the problem may even be nonexistent; that is, an erroneous temperature may be reported if there is failure of a thermocouple sensor or of the sensor electronics. If we are to represent expertise, we need to have a representation that allows us to capture probabilistic dependencies.

In the early-1970s, computer-science investigators attempted to extend rule-based systems to allow experts to express uncertain relationships that seemed important in medical diagnosis (Buchanan and Shortliffe, 1984). An extension to rule-based reasoning, named the *certainty factor* (CF) model, was developed to introduce uncertainty to the logical links (Shortliffe and Buchanan, 1975). To date, the CF model, as well as straightforward logical inference, have dominated the attention expert system developers for electric power utilities and other application areas.

Unfortunately, several theoretical and empirical analyses have demonstrated that logical reasoning—even when extended with the CF model—is inadequate for representing, reasoning, and maintaining knowledge bases that contain general probabilistic dependencies (Heckerman, 1986; Heckerman and Horvitz, 1987; Pearl, 1988). Problems noted to date with rule-based methods for reasoning under uncertainty include limitations in the expressiveness of rules, difficulties with maintenance, and suboptimality of information-acquisition. Rule-based methodologies for reasoning under uncertainty typically impose invalid independence assumptions (Heckerman, 1986; Heckerman and Horvitz, 1987; Pearl, 1988) and have been shown to be an inefficient representation for expressing independence among observations and disorders (Shachter and Heckerman, 1987; Heckerman, 1990). In addition, updating and maintaining a set of rules can pose an enormous challenge. There are few tools for verifying the consistency of a set of rules after changes are made to some portion of a rule base. Finally, the traditional rule-based architecture does not allow for the dynamic custom-tailoring of requests for additional tests and information so that the requests are tailored to the current state of information. Rather, they require users to evaluate stereotypical comprehensive checklists.

C Belief Networks

Recently, there have been significant advancements in the area of *probabilistic expert systems*. The underlying representation in these systems is a probability model encoded as a belief network (Pearl, 1988; Henrion et al., 1991). The belief network provides a mechanism for modeling dependencies between faults and symptoms and for performing diagnostic reasoning. The network is a representation that allows the model builder, working with an expert, to efficiently encode expert knowledge about probabilistic dependencies among important distinctions in a domain. Belief networks allow for the efficient maintenance of large knowledge bases by providing a language for effectively specifying thousands of consistent rules. This approach can be used with a value-of-information analysis to custom-tailor questions based on the uncertainty and the cost-effectiveness of alternative tests in specific contexts.

A belief network is a directed, acyclic graph containing *nodes* representing relevant distinctions or propositions (e.g., hypotheses and observations), and *arcs*, representing probabilistic dependencies among nodes. Each node is associated with a set of mutually exclusive and exhaustive values that represent alternative states of a proposition. Figure 1 displays a portion of a belief network for a gas turbine. Here, possible faults include failures of the oil-cooling system, failure of bearings, or failure of the bearing temperature sensor. As indicated by the arcs between the nodes, these faults can affect the probability of observing different values of findings. For example, a bearing failure or an oil cooling-system failure can change the probability of observing an abnormally high oil temperature. A bearing failure can also affect the reading of a vibration sensor, as well as the the value of measured bearing temperature. Note that a direct failure of the bearing temperature sensor can also be responsible for an increase in the measured bearing temperature, thus belief networks provide a mechanism to model sensor failure explicitly.

The actual numerical expression of conditional dependencies is not expressed in the structure of a belief network. After the general dependency structure of a belief network has been determined, we assess the probability distributions over descendant nodes, given specific values of ancestor nodes. An important component of the numerical assessment procedure is the determination of the prior probabilities, or prevalences of alternate failures. That is, each primary failure is assigned a failure probability obtained from analysis of operating experience or mean-time-

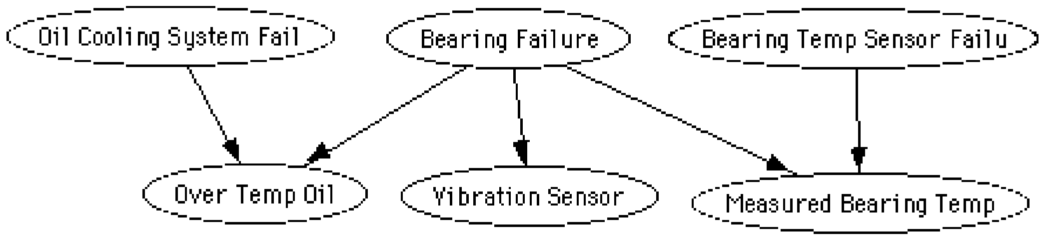

Figure 1: A small belief network that represents probabilistic dependencies among observables and disorders. Directed arcs capture information that the value of a predecessor node can change the probability distribution over an descendant node.

between-failures (MTBF) data.

The diagnostic inference problem is to update the probabilities of alternate possible faults as additional information about observations becomes available. A set of algorithms have been developed to compute the probabilities of values of nodes in a belief network coherently, given the specification of the values of some subset of nodes (Pearl, 1988). For example, let us suppose an operator observes a warning about high oil temperature, but the vibration sensor is reading normal. We wish to determine the probability of a bearing failure versus a sensor failure. To obtain such likelihood information, we apply a belief network algorithm to calculate the probability of a bearing or cooling system failure given our observations about the temperature and vibration sensors.

Figure 2 shows a fragment of a belief network that has been constructed for the GE–Model MS7001 Industrial Gas Turbine. The belief network enables us to model complex dependencies specified by a system engineer and to perform the necessary probabilistic updating. This network is structured as a causal model, that is, the arcs point in the causal direction. We can trace a path from a node representing that a foreign object in the compressor to nodes representing bent or missing compressor blades, which in turn can affect compressor efficiency.

To assess the probabilities of alternative states given predecessors, we assess in the causal direction. A feature of belief networks is the ability to assess the networks in the direction that is most comfortable for the expert (Shachter and Heckerman, 1987). Studies have demonstrated the people are more facile at assessing probabilities in the *causal direction* (Kahneman et al., 1982) than in the diagnostic direction. In assessing probabilities in the causal direction, we ask an expert to consider the likelihood that an observation will be one of several values, given that a fault is true. Cognitive psychologists have found that people have more trouble with assessing probabilities in the *diagnostic* direction, that is, assessing the probability that alternate faults will be true, given observations.

The assessment of rules for a rule-based system requires experts to specify rules in the diagnostic direction, by specifying links from observations to faults and actions. Herein lies the efficiency of maintaining belief networks versus knowledge bases in rule-based systems. With belief networks, we can change the behavior (e.g., the failure rate) of a particular component simply by changing a single node, link, or probability distribution over a link. Such simple changes typically lead to thousands of changes in the diagnostic output of the expert system. When we attempt to update the behavior of a component of a gas turbine, as represented by a rule-based system, we must carefully consider how the modification affects many diagnostic rules. As all engineers familiar with rule-based expert systems have discovered, this task is time-consuming and tedious. Belief network algorithms are used to compute the probability of alternative faults, given some

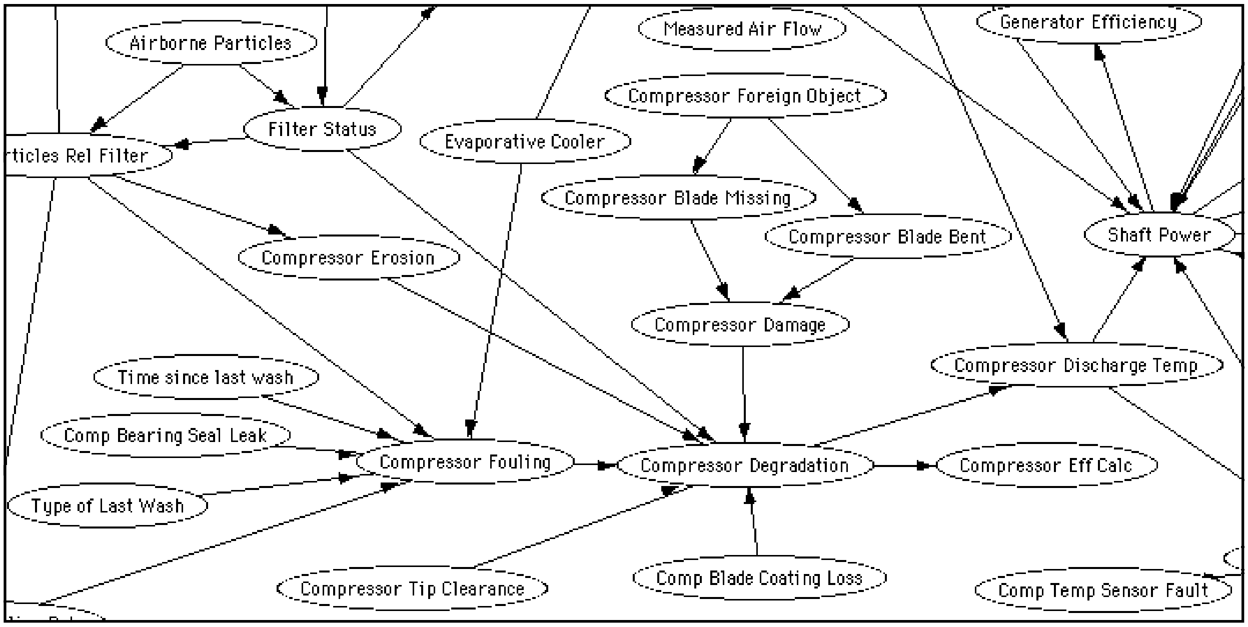

Figure 2: A portion of the GE–Model MS7001 Industrial Gas Turbine belief network showing the causal relationships among faults and observations.

set of findings. Thus, in essence, a belief network is used to generate diagnostic rules from a causal model.

D Incorporating Thermodynamic Analysis

In order to determine if there is an efficiency problem with a turbine, we must have some notion of what heat rates would be under normal or "no-fault" operating conditions. Determining that there is some deviation from design performance can be difficult because the efficiency varies with ambient temperature, pressure, and humidity among other factors. We have incorporated thermodynamic analysis of the turbine heat cycle into the diagnostic system. In contrast to the probabilistic model described in the previous section, the thermodynamic model consists of a set of deterministic equations relating airflows, fuel flow, temperatures, and pressures to measured and predicted power output levels. The analytic model, Efficiency Map, developed by Enter Software for EPRI has been integrated into the inference scheme. A fragment of a network describing the functional relationships in the thermodynamic model is shown in Figure 3. Note that these relationships are not characterized probabilistically. Efficiency Map operates by taking measurements for power output, fuel flow and heat value, water injection rates, and other design and operating parameters and using it to predict the values of such parameters as turbine and compressor efficiencies, firing temperature, and net power output. Comparing predicted to design or observed values for these variables (and more importantly, their trends over time) provides information regarding various faults.

We have designed a hybrid inference algorithm that uses regression analysis to process the analytical portion of the model and probabilistic reasoning algorithms to process the belief network. The technique for integrating the belief network with the thermodynamic model focuses on treating the *outputs* of the thermodynamic

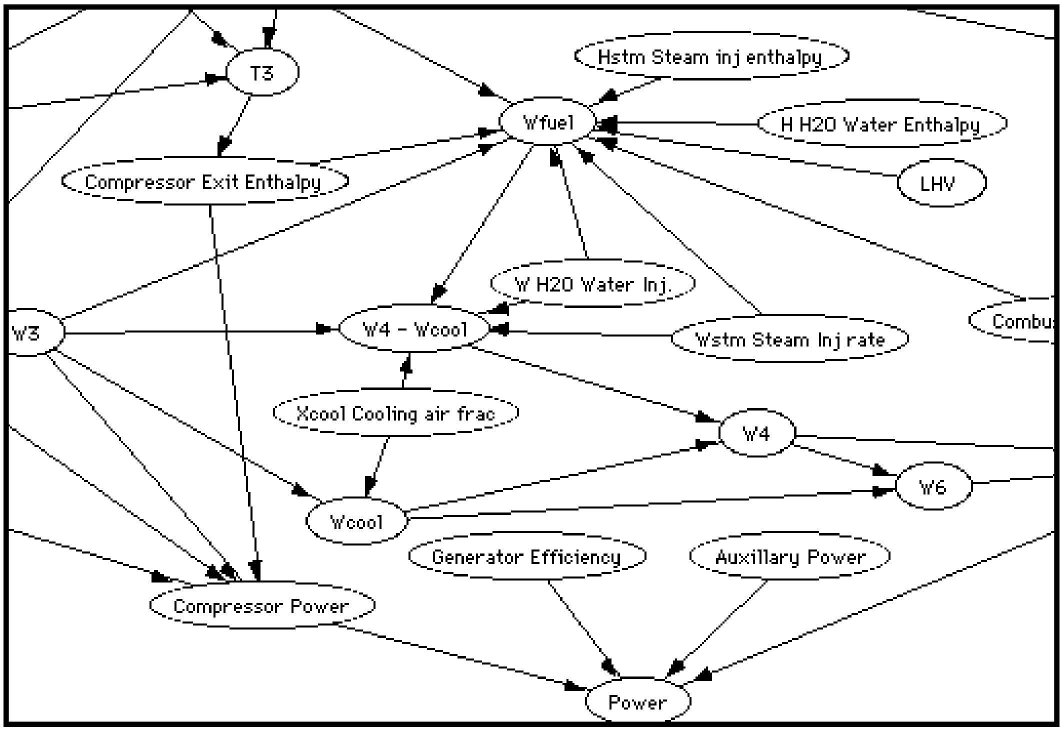

Figure 3: A portion of a network highlighting the thermodynamic analysis of turbines.

model as *inputs* into the belief network representation. The estimated efficiencies, temperatures, and pressures that are developed by Efficiency Map are used as findings for the belief network. The analytical findings are then fused with information about such things as service history, operating mode, and environmental conditions to generate an overall diagnosis.

The methodology for incorporating Efficiency Map estimates into the expert system addresses two distinct forms of uncertainty. The first is characterizing which faults in the system could cause various types of off-design behaviors that might be indicated by Efficiency Map. For example, various problems with turbine damage or erosion will cause a decrease in turbine efficiency. This causal relationship, while uncertain, is based on an understanding of turbine operation. The second facet of integration is characterizing the quality of the outputs from the thermodynamic model. The predictions of Efficiency Map are not perfect, due to uncertainty in the inputs to the model and approximations in the analytic methods. Uncertainty of this type is modeled by treating the outputs as as noisy estimates about the true value of the parameter. The belief net and associated updating algorithms are used to fuse these different types of uncertainty in the top level diagnosis.

E Delivering Advice to Users

The expertise and data contained in the belief network is incorporated into a shell for delivering diagnostic services to users in an easy-to-use fashion. The primary function of the system is to perform diagnosis: Given information about observations or findings, what is the likelihood of alternative failures or faults? The system can reason about the relevance of findings to a diagnosis, and it can suggest crucial findings for narrowing a diagnosis. The system also has an integrated facility for providing access to a videodisc library that enables the user to inspect an extensive base of diagnostic knowledge and slides.

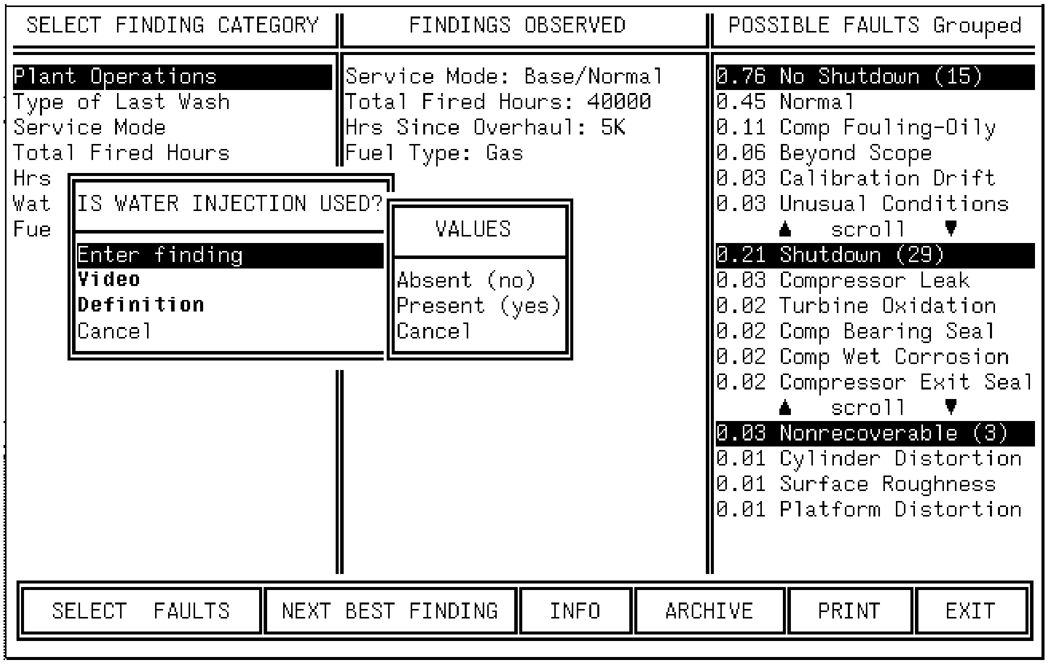

Figure 4: Selecting a finding and entering a finding value in the delivery system

The central function of the system is to form diagnoses for finding values that are entered. The set of faults diagnosed by the initial version of the expert system is listed in the Appendix. Figure 4 shows the sequence used to enter a finding. The findings entered are displayed in the FINDINGS OBSERVED column, and the diagnosis is formed in the column labeled POSSIBLE FAULTS at the right.

All faults that are consistent with the observed findings appear in the diagnosis column. At the left of each fault is its probability. The diagnosis is displayed initially in groups: Shutdown, No Shutdown, and Nonrecoverable. The faults in the "No Shutdown" group can be treated without shutting down the unit, while those in the "Shutdown" group require the unit to be brought down, either for an overhaul or servicing. The "Nonrecoverable" faults cannot be fixed in a standard overhaul.The faults are ranked according to their likelihood within each group. To the left of each group title, the probability for the entire group is listed. The number of faults within the group is listed on the right side of the group name.

Additional findings are entered by selecting them from the selection column. After entering each finding, the finding value is listed in the findings-observed column and the diagnosis is revised. In Figure 5, a new finding value (Low Power Output) has been added.

An extremely important capability of the system is its ability to suggest the best finding enter or test to perform. The system identifies that finding that can most effectively narrow the diagnosis to a single fault or fault group, based on a cost/benefit analysis. The system utilizes three diagnostic reasoning strategies to narrow the diagnosis:

1. Group-discrimination mode: Determines the best findings for distinguishing among the fault groups.
2. Fault-discrimination mode: Determines the best findings for distinguishing among all faults .
3. Pursual mode: Determines the best findings for distinguishing the fault with the highest probability (the leading fault) from the other faults .

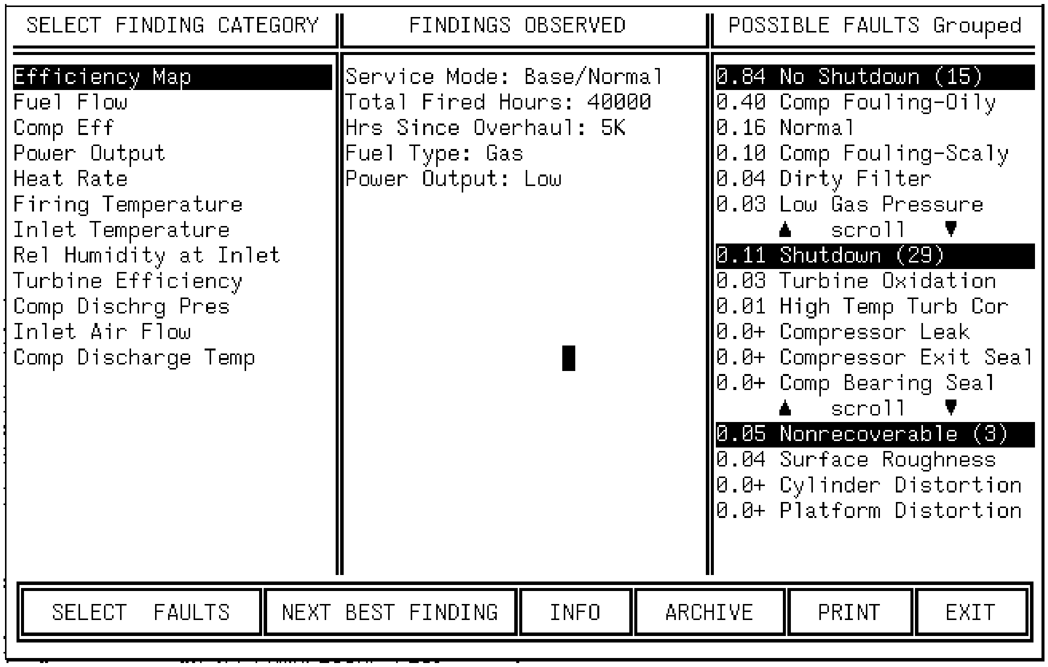

Figure 5: The diagnosis after entering another finding.

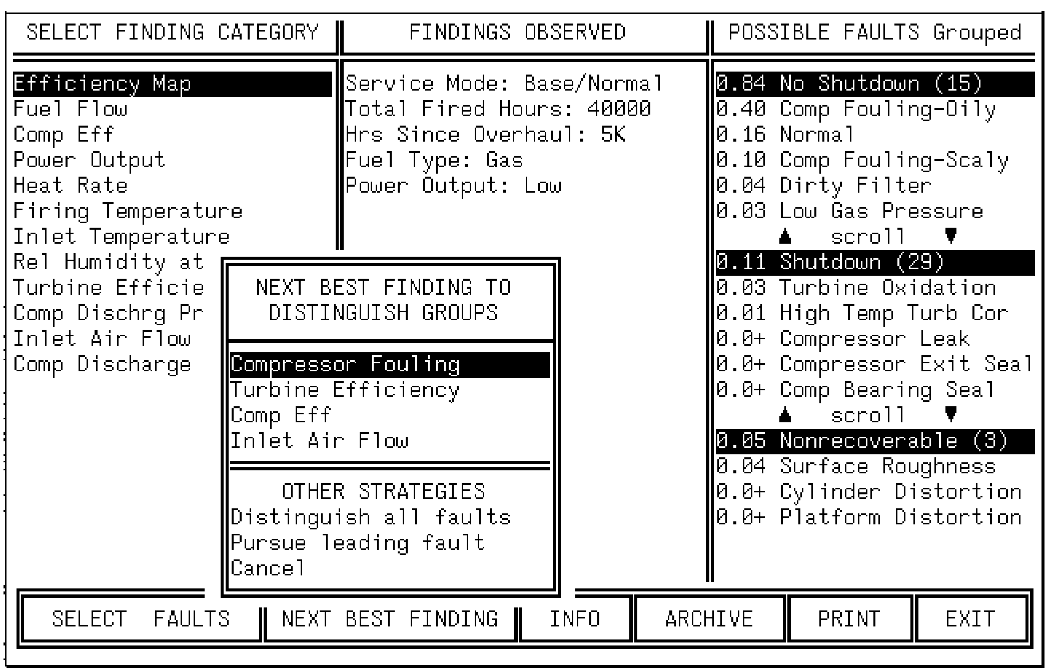

Figure 6: The crucial-findings pop-up window, showing the other strategies that can be selected.

Figure 6 shows the *next best-findings window* that distinguishes among groups. It is important to note that expensive tests will not be considered by the crucial-findings strategies until all relevant inexpensive findings have been entered. The system can generate a graphic explanation (justification) of the usefulness of a finding recommended by the system. An example explanation is shown in Figure 7. The effect that observing each value of the finding would have on the groups of faults or a single fault in the diagnosis is represented by the length and the direction of the arrows in the graph. The length of the arrow represents a logarithmic ratio. If the arrow associated with a value points towards a fault or group of faults, observing that value for the finding would raise the likelihood of that fault or group of faults, relative to the other findings.

In addition to diagnostic reasoning, the system can access a wide range of other information sources including such things as:

- Motion and Still Video
- Repair Techniques
- Reporting Requirements
- Parts and Tools

Each of these options can provide a link to video and/or textual information about that fault. For the engine modules, these information sources are developed on a customized basis. Video images include such things as color photographs, schematics, and repair charts.

F Summary and Conclusions

The probabilistic belief network approach has been applied to the diagnosis of thermal performance problems in operating industrial turbine engines. The use of a probabilistic causal model is particularly advantageous for such a complex system since there are few experts to rely on for diagnostic information and thermodynamic analysis plays a key role in diagnosis. The currently operating version of the system diagnoses approximately 50 faults using about 40 possible findings. These faults and findings are listed in the Appendix. Preliminary evaluation of the system indicates that it can distinguish well between various classes of faults and its recommendations for the next test to perform often agrees with that of experts in the field.

The methodology employed to develop and deliver the expertise in this system is innovative. We are able to fuse information from a deterministic, thermodynamic analysis with uncertain expert information from plant managers and operators. We anticipate that the model-based approach employed here will result in a more flexible and maintainable system than previous rule-based expert system efforts.

We are currently planning to install the expert system on-line at two GE–Model MS7001 Industrial Gas Turbine engines operated by the city of Santa Clara in Santa Clara County, California as a means of obtaining more complete validation information. Although no results of the application of the system to an operating engine are available at the time of publication, we expect some early results of this application may be available for discussion at the TURBO EXPO conference in June 1992.

Acknowledgment

This research is supported by the Electric Power Research Institute under EPRI Project RP3250-07, "A Thermal Performance Advisor for Gas Turbine Power Plants".

Appendix: Faults and Findings

The set of faults considered by the initial system are:

Fuel Nozzle Blockage
Fuel Oil Check Valve
Comp Temp Sensor

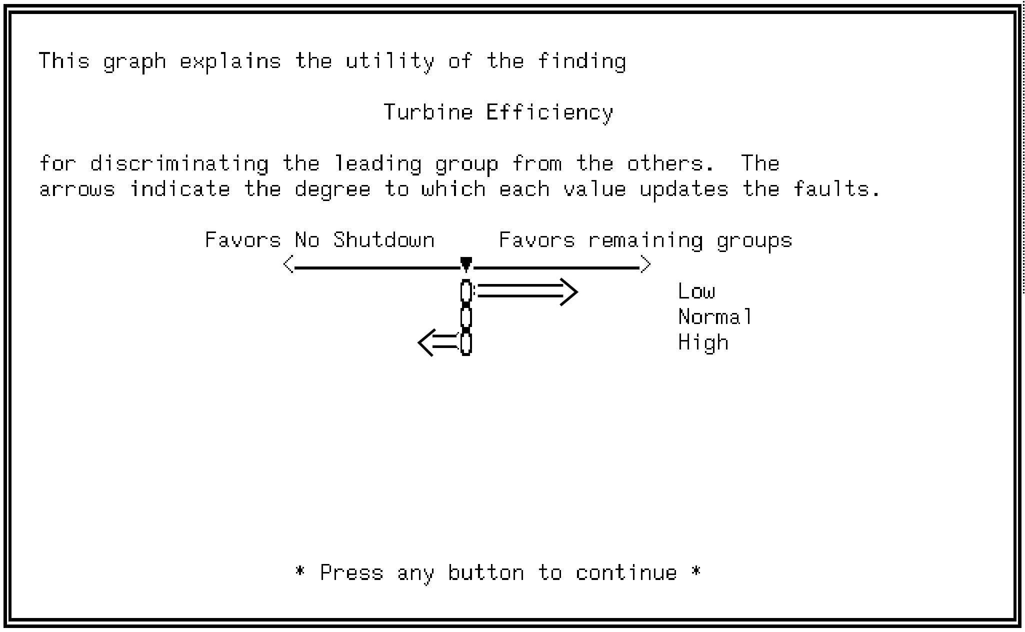

Figure 7: Graphic explanation of discriminatory power of crucial findings.

Compressor Erosion
Comp Bearing Seal
Compressor Fouling (Oily)
Compressor Tip Clearance
Compressor Leak
Ice Fouling
Compressor Exit Seal
Compressor Fouling (Scales)
Compressor Blade Bent
Comp Wet Corrosion
Compressor Blade Coat
Compressor Blade Missing
Comp Pres Transducer
Comp Sens Pres Tube
Dirty Filter
Improper GV Position
GV Indicator
Bearing Temp Sensor
Bearing Failure
Fuel Distribution Low Gas Pressure
Clogged Fuel Line
Gas Control Valve
Fuel Supply
Fuel Oil Pump
Oil Cooling System
Exhaust Thermocouple
Exhaust Blockage
Bleed Valve Open
Bleed Valve Leakage
Turbine Silica Deposits
Joint Leakage
Calibration Drift
Turbine Tip Clearance
Turbine Blade Damage
Transition Piece
Turbine + Transition
1st Stage Nozzle
Turbine Oxidation
High Temp Turb Corrosion
Inlet Wash Open
Cylinder Distortion
Surface Roughness
Platform Distortion
Generator Fault

The set of findings considered by the initial system are:

Type of Last Wash
Service Mode
Total Fired Hours

Time Since Overhaul
Water Injection
Fuel Type
Compressor Tip Clearance
Compressor Icing
Compressor Fouling
Comp Blade Coating Loss
Compressor Blade Missing
Compressor Erosion
Vibration
Turbine Tip Clearance
Turbine Corrosion
Turbine Blade Damage
Turbine Silica Deposit
Fuel Flow[1]
Compressor Efficiency [1]
Power Output[1]
Heat Rate[1]
Firing Temperature[1]
Comp Inlet Temperature[1]
Rel Humidity at Inlet[1]
Turbine Efficiency[1]
Comp Discharge Press[1]
Inlet Air Flow[1]
Comp Discharge Temp[1]
Particles from Filter
Airborne Particles
Guide Vane Angle
Guide Vane Stuck
Filter Type
Filter Particle
Cold Combustor Can
Pattern Factor
Hot Combustor Can
Exhaust Temp Spread
Exhaust Thermocouple
Generator Cooling
Measured Bearing Temp
Low Fuel Pressure
Bearing Vibration
Oil Temperature

[1]Efficiency Map Parameter

References

Buchanan, B. and Shortliffe, E., editors (1984). *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison-Wesley, Reading, Mass.

Heckerman, D. (1986). Probabilistic interpretations for MYCIN's certainty factors. In Kanal, L. and Lemmer, J., editors, *Uncertainty in Artificial Intelligence*, pages 167–196. North Holland, New York.

Heckerman, D. (1990). *Probabilistic Similarity Networks.* PhD thesis, Medical Information Sciences, Stanford University, Stanford, CA.

Heckerman, D. and Horvitz, E. (1987). On the expressiveness of rule-based systems for reasoning under uncertainty. In *Proceedings AAAI-87 Sixth National Conference on Artificial Intelligence*, pages 121–126, Seattle, Washington. American Association for Artificial Intelligence.

Henrion, M., Breese, J., and Horvitz, E. (1991). Decision analysis and expert systems. *AI Magozine*, 12(4):64–91.

Kahneman, D., Slovic, P., and Tversky, A., editors (1982). *Judgement under uncertainty: Heuristics and biases.* Cambridge University Press.

Klempner, G., A., K., and Lloyd, B. (1991). The EPRI generator monitoring system: Experience with the GEMS prototype. Presented at the American Power Conference.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufman, San Mateo, Ca.

Shachter, R. and Heckerman, D. (1987). Thinking backward for knowledge acquisition. *AI Magazine*, 8:55–63.

Shortliffe, E. and Buchanan, B. (1975). A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23:351–379.