Improved Explicability for Pump Diagnostics in Nuclear Power Plants

Andrew T. Young,* Graeme M. West,* Blair D. Brown,* Bruce Stephen,* and Stephen D.J. McArthur*

*Institute for Energy and Environment, Department of Electronic and Electrical Engineering, University of Strathclyde, 204 George Street, Glasgow, G1 1XW, U.K.

INTRODUCTION

To ensure the continued safe operation of many of the UK's fleet of advanced gas-cooled reactors (AGRs) effective and reliable monitoring of several of the key plant items is essential. Out of these key items, a significant portion of these are rotating plant assets, one asset in particular that is crucial to the operation of the station are the boiler feed pumps (BFPs). The BFPs in an AGR station move water from a condenser into a boiler, the water is then heated which produces steam and this steam turns the electricity-generating turbines. Currently, the operator of the AGR stations employs a time-based maintenance strategy for BFP assets: after a defined amount of time each asset is removed, replaced with a rotated spare, and a complete overhaul is then performed on the removed asset. This procedure can result in the removal of an asset before any significant wear has occurred, therefore increasing maintenance and generation costs. Conversely, this could result in an unplanned outage due to a component failure which leads to both a decrease in power output of the station and hence a decrease in revenue for the operator. Because these pumps are essential for the generation of electricity there are several pressure, temperature, vibration and speed parameters constantly monitored during the operation of this asset. Currently, data analysts have to manually analyse all this data by following a set diagnosis process, the consequential time burden on the analyst is therefore extremely high.

Data-driven approaches to solve this problem, and other similar problems, have the capability to produce accurate results similar to what the analysts can achieve in a fraction of the time. However, the majority of these techniques are black box techniques and lack explicability which is often a requirement for problems involving critical assets in the nuclear industry.

The main outcomes of this work are to address the time burden placed on the analysts by automating elements of the existing diagnosis process, through the implementation of an intelligent rule-based expert system, that provides adequate explicability to the user to satisfy requirements. Additionally, a recurring problem in the design of expert systems for industry is the cost involved with the knowledge elicitation process. Here we propose a questionnaire style approach, similar to what the domain experts currently use, to extract this knowledge without the need for a structured interview. By using this information a signal-to-symbol transformation algorithm is designed to assign time periods symbols that relate to the various rules defined by the domain experts. The final system combines the data-driven signal-to-symbol transformation algorithm and the rule-based expert system to produce a hybrid system that can be used to classify defects based on a set of rules and also explain to the user the reasoning behind this solution.

EXPERT SYSTEMS

Knowledge Based vs. Data Driven

While both knowledge-based, e.g. expert systems, and data-driven, e.g. machine learning can produce similar results they differ significantly in their implementation.

Machine learning techniques are primarily based around statistical modelling of the problem data. The model parameters are learned through a training process that attempts to produce the correct output given the input data for the majority of cases. In the case of a number of models, it is often very difficult to explain why or how these results have been produced.

Expert systems, unlike data-driven approaches, attempt to solve complex problems that comprise a significant amount of expert knowledge or human expertise. Encoded into the expert system is both the expert knowledge but also the reasoning behind this, therefore, they are able to explain and justify the reasoning behind any decision made.

For critical assets especially in the nuclear industry where safety is of the utmost importance, the lack of explicability of many data-driven approaches presents a problem. There is often a requirement to provide supporting evidence for applications such as this when classifying faults as there is a significant cost involved in the replacement or maintenance of these assets. This is one of the main advantages of expert systems over machine learning approaches. However, there is a significant cost involved in the development of an expert system through the man-hours required to extract the domain-specific knowledge from the engineers.

A typical expert system consists of five main components (Fig. 1):

1. **Knowledge Base:** This contains all the domain-specific digitized knowledge acquired from the engineers. For a rule-based expert system, this is represented in a number of **if...then...** rules. This can be considered as a fixed set of data, i.e. it remains the same throughout the decision making process.

2. **Real World View of Data:** The data, or facts relating to the data which is specific to the problem being solved is stored here. These facts are then matched with the **if...** conditions in the knowledge base to determine the solution or intermediate facts which are also stored in the real world view.

3. **Inference Engine:** This performs the processing of the knowledge by combining the rules in the knowledge base, the facts in the real world view and facts determined from the user to determine a solution.

4. **Explanation Mechanism:** Provides explanations and justifications for any decision made by the inference engine, this includes intermediate actions as well as conclusions. An adequate explanation mechanism is crucial for the acceptance of the system by the user.

5. **User Interface:** Provides communication between the expert system and the user. Inputs can be additional facts through the use of question and answers, natural language, graphics interface, or menu driven. The output is the solution with a detailed explanation of the reasoning behind this, the information can also be passed to an external program.
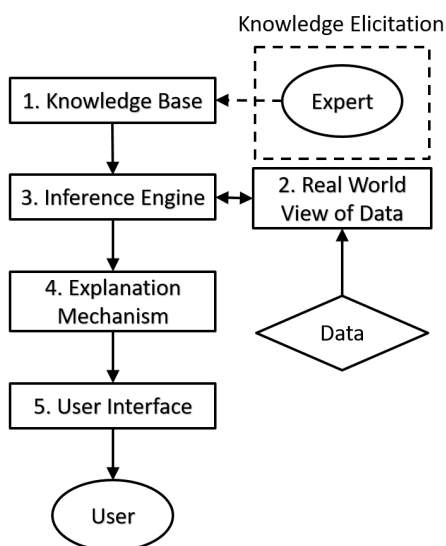

Fig. 1: Typical expert system architecture

Knowledge Acquisition

The knowledge acquisition phase is a critical step in the building of an expert system. This relates to the process of extracting the domain-specific knowledge from the engineers and organising this into a set of rules. A common approach to knowledge acquisition adopted for the implementation of an expert system follows a four-staged approach:

1. **Scoping:** This phase looks into gaining an understanding of the problem domain, identifying domain experts, and deciding an appropriate knowledge acquisition process, e.g. structured or unstructured interviewing, or questionnaires.

2. **Knowledge Elicitation:** Having decided on an appropriate knowledge acquisition process, this stage involves conducting a structured interview with domain experts.

3. **Knowledge Modelling:** Having extracted all the information from the experts in the previous step, the knowledge and information are collated and analysed. The output from this is a set of formal rules that can be used as a basis for the implementation of a prototype expert system.

4. **Knowledge Verification:** The prototype expert system is then applied to a specific problem by the experts to verify the correct capture and formulation of the knowledge. If the knowledge base is incorrect or incomplete it is necessary to understand the reason for this and what is missing, the entire process has to then be repeated until the knowledge base is sufficient for the given problem.

METHODOLOGY

For many industrial applications fault diagnosis involves the engineers following a set diagnosis process, therefore, the expert knowledge has already been acquired, although this is not always complete enough to be formalised into a set of rules for an expert system. For the BFP problem, the diagnostic rules for the asset were supplied by the original equipment manufacturer (OEM), see Table I. This related to each data stream necessary to diagnosis a given fault, and the rules were represented by a set of trends, i.e. stable, fluctuating, rising or falling. The data contained 37 faults and the associated trends for 10 specific data streams covering pressure, temperature and speed. While these rules are sufficient for the domain experts to analyse the problem, there is a significant amount of knowledge that is not captured. The diagnostic rules did not quantify what a specific rise, fall or fluctuation of a specific data stream was, and also whether these changes were consistent or rule specific. Additionally, due to the nature of the asset under investigation and the different types of data, (e.g. pressure and temperature), no information was supplied regarding the order these changes might occur. This is a common problem when formalising the rules for an expert system, as there is a considerable amount of acquired knowledge that engineers have

Cause	Pump Suction	Pump Discharge	Oil Supply	Gland CW Supply	Gland Drain / Outlet
Anti friction bearing failing	F		↓	↑	↑
Speed (rpm) too low		↑			
Higher sand concentration		F		↓	↓
Inadequate lubrication	F		↓		
Effects of high tides	↑				F

TABLE I: Example format of pressure datastream specific rules. (F is Fluctuating, ↑ is Rising and ↓ is Falling). Data representative of actual rules, however, specific rules have been changed.

Datastream	How much?	How long?	Order	Comments
Pump Suction	*5% dec*	*1 hour*	2	*As defined in O&M manual*
Pump Discharge	*10% dec*	*1 hour*	1	*this is the standard expected*
Gland Drain / Outlet	*5% dec*	*1 hour*	3	*response.*

TABLE II: Example of data trend qualification document response

to asses the problem. The preferred approach for extracting this additional information is to set up a structured interview, where a set of questions can be developed to address each point. However, in many applications, the engineers that asses these problems are time poor and therefore it is not always for them to allocate a sufficient amount of time to address all these issues. Alternative approaches aimed at reducing the time burden include observation of the expert or a questionnaire style approach. For this problem, a questionnaire style approach was adopted as this allowed the time impact on the experts to be spread over several weeks. The key information that had to be extracted from the experts about the rules previously supplied was a quantification of rising, falling and fluctuating, how long this trend occurs over, and any order to the rules. These questions were written up in a questionnaire to be sent to the domain experts. For each fault a table was produced, to answer these three questions, an example of this is shown in Table II with an example response.

After receiving this information from the experts it was then possible to begin to formulate this into a prototype expert system. Due to the individual rules being based on trends in the data it is necessary to have a pre-processing stage to calculate these trends. The approach adopted for this application was to use a signal-to-symbol transformation approach [1]. Splitting up the trace into various timesteps based on the information provided by the expert, each timestep is assigned a symbol which is either rising, falling, fluctuating or stable. By inputting these symbols as facts into the database, it is possible for the expert system to determine if there is a fault at that given moment in time. If a fault is detected the output will explicitly state the fault detected, the individual data streams used to determine this fault, and any associated comments provided by the experts as to why these trends relate to this fault.

FINAL SYSTEM

The complete architecture from the expert system is shown in Fig. 2. CLIPS (C Language Integrated Production System) is used for the main components of the expert system, i.e. the database, inference engine, and explanation mechanism, a shell program is written in MATLAB to interface both the user interface and the signal-to-symbol transformation algorithm with CLIPS. In addition to this, the original rules supplied by the OEM came in the form of a Microsoft Excel spreadsheet because of this and to keep the system as flexible as possible an algorithm was produced to convert this spreadsheet into a CLIPS format knowledge base ".clp" file. This will allow the user to add additional rules or alter existing rules as it becomes appropriate and also allows for a new set of rules for an alternate rotating asset.
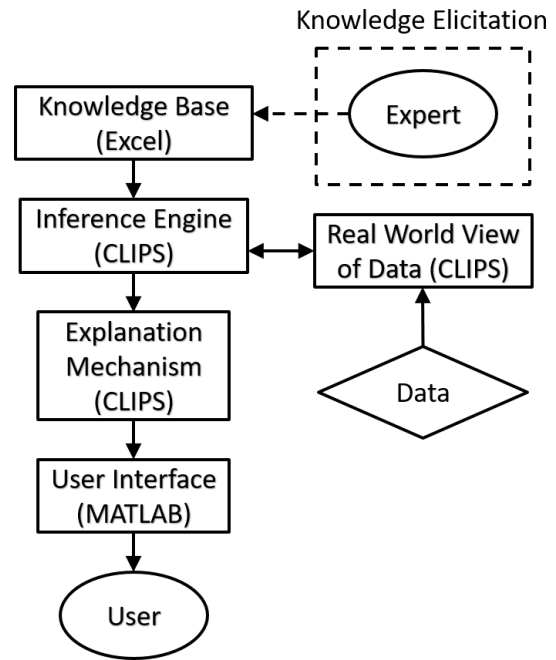

Fig. 2: Final expert system architecture

RESULTS

The output from the expert system will inform the user of the fault diagnosis conclusion, the datastreams used to reach this conclusion with the associated information relating to the quantification of any trends and also any additional comments provided by the experts. An example output of this is shown in Fig. 3.

```
Fault detected at 10th Feb 2021 between 1200 and 1300
Cause of fault was Anti friction bearing failing

Pump discharge decreased by over 10%
Pump suction decreased by over 5%
Gland drain/outlet decreased by over 5%

Comments: As defined in the O&M manual this is the
standard expected response for this type of fault.
```

Fig. 3: Example output to the user from the expert system

As well as providing this detailed information to the user about the fault detected, it is also necessary to supply the

analyst with additional information in a format that is familiar to them, so that they can verify the results. The way this was done was to show the output of the signal-to-symbol transformation for all the relevant signals. This allows the user to analyse each individual data stream to confirm that the trends have been classified correctly and this is the correct diagnosis of the fault. Example output for both the pump suction pressure and thrust bearing temperature is shown in Fig. 4. From this, the expert can not only see the relevant information but by displaying the historical information allows them to asses if this fault has happened in the past, and whether there is an underlying reason this fault is occurring again.
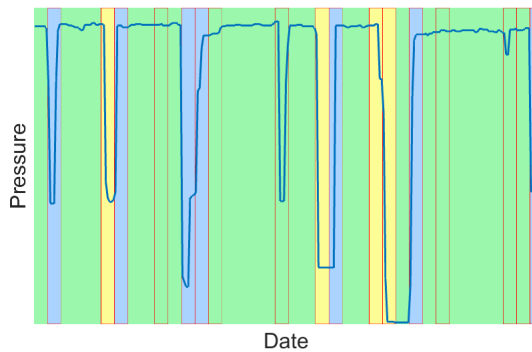
(a)
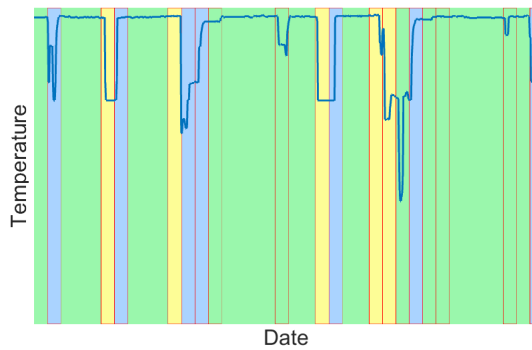
(b)

Fig. 4: Example of signal to symbol transformation for (a) pump suction pressure and (b) thrust bearing temperature. (Green - Stable, Blue - Rising, Yellow - Falling, and Red - Fluctuating)

To date all 37 rules from the operator's fault diagnosis process have been implemented in the expert system. However, this currently does not include information regarding the relative order of the features occurring in the data. While the 37 rules defined by the operator cover all known faults that can occur on the BFPs, this only covers a small region of the problem space. With 10 data streams and 4 possible states this produces a total of 1,048,576 (4^{10}) possible rules, although many of these states are likely to be physically impossible or related to normal behaviour.

CONCLUSIONS AND FUTURE WORK

Following the approach discussed in the paper, it has been possible to develop an expert system that can identify faults in near-realtime that occur in a BFP of an AGR. By following a formal knowledge acquisition process it was quickly possible to identify the missing information in the knowledge initially supplied by the OEM. Due to the problem of not being able to formally interview the domain experts it was necessary to develop a questionnaire that the experts could complete to acquire this missing information. The final system integrates several platforms, i.e. MATLAB, CLIPS and Microsoft Excel, to produce an end to end system for automatic fault detection in rotating assets.

As mentioned the current set of rules only covers a very small percentage of the total problem space therefore further work will involve the exploration of the remaining problem space to eliminate impossible states and attempt to uncover new rules based on historical states. Additionally, the temporal aspect of certain rules will have to be addressed. Finally, a complete system will have to be validated by the domain experts to ensure that all the knowledge was captured and formulated correctly. If not more work will be required to extract any additional information necessary.

ACKNOWLEDGEMENTS

REFERENCES

1. J. J. A. COSTELLO, G. M. WEST, S. D. J. MCARTHUR, and G. CAMPBELL, "Self-Tuning Routine Alarm Analysis of Vibration Signals in Steam Turbine Generators," *IEEE Transactions on Reliability*, **61**, *3*, 731–740 (Sep. 2012).