This is a peer-reviewed, accepted author manuscript of the following article: Young, A., West, G., Brown, B. D., Stephen, B., Michie, C., & McArthur, S. (Accepted/In press). Symbolic representation of knowledge for the development of industrial fault detection systems. Paper presented at International Congress and Workshop on Industrial AI 2021, Luleå, Sweden.

Symbolic Representation of Knowledge for the Development of Industrial Fault Detection Systems

Andrew Young
University of Strathclyde
204 George Street, Glasgow
G1 1XW, UK
+44 141 444 7241
andrew.young.101@strath.ac.uk

Graeme West
University of Strathclyde
204 George Street, Glasgow
G1 1XW, UK
+44 141 548 3542
graeme.west@strath.ac.uk

Blair Brown
University of Strathclyde
204 George Street, Glasgow
G1 1XW, UK
+44 141 444 7027
blair.brown@strath.ac.uk

Bruce Stephen
University of Strathclyde
204 George Street, Glasgow
G1 1XW, UK
+44 141 444 7260
bruce.stephen@strath.ac.uk

Craig Michie
University of Strathclyde
204 George Street, Glasgow
G1 1XW, UK
+44 141 548 2521
c.michie@strath.ac.uk

Stephen McArthur
University of Strathclyde
204 George Street, Glasgow
G1 1XW, UK
+44 141 548 4838
s.mcarthur@strath.ac.uk

ABSTRACT

In critical infrastructure, such as nuclear power generation, constituent assets are continually monitored to ensure reliable service delivery through pre-empting operational abnormalities. Currently, engineers analyse this condition monitoring data manually using a predefined diagnostic process, however, rules used by the engineers to perform this analysis are often subjective and therefore it can be difficult to implement these in a rule-based diagnostic system. Knowledge elicitation is a crucial component in the transfer of the engineer's expert knowledge into a format suitable to be encoded into a knowledge-based system. Methods currently used to perform this include structured interviews, observation of the domain expert, and questionnaires. However, these are extremely time-consuming approaches, therefore a significant amount of research has been undertaken in an attempt to reduce this. This paper presents an approach to reduce the time associated with the knowledge elicitation process for the development of industrial fault diagnostic systems. Symbolic representation of the engineer's knowledge is used to create a common language that can easily be communicated with the domain experts but also be formalised as the rules for a rule-based diagnostic system. This approach is then applied to a case study based on rotating plant fault diagnosis, specifically boiler feed pumps for a nuclear power station. The results show that using this approach it is possible to quickly develop a system that can accurately detect various types of faults in boiler feed pumps.

Keywords

Condition monitoring, nuclear power plants, expert systems, knowledge-based systems, automation.

1. INTRODUCTION

Fault detection and diagnostics is an active research area, especially in the nuclear industry for rotating machinery [1, 6, 10, 13]. There are two main approaches that can be adopted for the development of systems for fault detection or diagnostics. These are either data-driven approaches, e.g. machine learning, or knowledge-based approached, e.g. expert systems. While both of these techniques have similar aims and can provide similar results they differ quite significantly in their implementation.

The basis for data-driven approaches is centred around statistical models of the problem data. The individual parameters of the model are learned through a process called training where a large volume of data is input into the model and the model attempts to produce the correct output for the majority of cases. It should also be noted that related to many data-driven approaches a balanced dataset is required, i.e. there is an equal number of samples for both normal data and fault states. Due to the nature of these models and the lack of explicability for many data-driven approached; for critical assets (especially in the nuclear industry) that can present an issue. This is because supporting evidence is often required when making decisions on these assets as there is a significant amount of cost involved in the repair, replacement or downtime of these assets, another consideration is that "black box" techniques cannot support the writing of new safety cases.

Knowledge-based approaches are the second technique that can be used to solve this problem, they attempt to solve (or support the resolution of) complex problems where a significant amount of human expertise or expert knowledge is required. This knowledge is acquired from the engineers or domain experts through a process called knowledge elicitation, this is then formalised into a format that is compatible with the technique, e.g. as the rules for a rule-based expert system. The one main advantage of this type of approach over data-driven approaches is the ability to not only solve a problem but also to explain and justify the reasoning behind why a decision was made. However, this comes with the disadvantage that a significant time cost is associated with the capturing of the knowledge, then formalising this into a knowledge-based system. Because of this disadvantage, there his been a significant amount of research undertaken across numerous fields in order to streamline the knowledge elicitation process [9, 12], as this is the most time consuming part of the development of an expert system.

The next section of this paper provides background information into rule-based expert systems, a type of knowledge-based approaches. Section three proposes a new methodology for knowledge elicitation through the use of symbolic representation of the expert's knowledge and the parametrization of this knowledge. While section four presents a case study of this methodology applied to boiler feed pumps from an advanced gas-cooled reactor in the UK. Finally, the conclusions and future work are presented in section five.

2. RULE-BASED EXPERT SYSTEMS

Knowledge-based systems can be used for a variety of applications to provide not only accurate decisions but also the explanation and reasoning behind these decisions. One example of these is rule-based expert systems [5], they store the knowledge captured from the domain expert as a set of rules. These rules are formalised in a way that mimics the domain experts reasoning process and are mainly applied to knowledge or time-intensive problems.

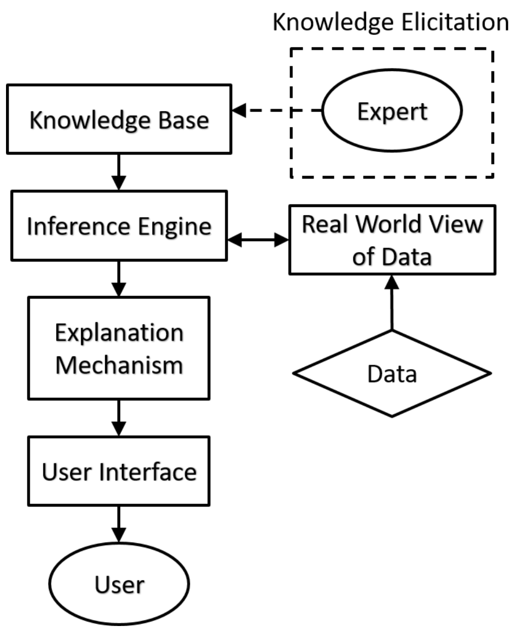

Figure 1: **Typical rule-based expert system architecture**

A typical rule-based expert system contains five main components, see Figure 1. The first of these is the **Knowledge Base** and this contains all the domain-specific captured knowledge from the experts. Figure 2 shows an example of how the rules are expressed in an expert system as a set of IF-THEN rules. This can be considered as a fixed set of data, therefore it remains the same throughout the decision making process. The **Real World View of Data** is the next component, this contains all the data, and facts relating to the asset under analysis. This can be considered the current state of the machine and hence is fluid and constantly changing. The facts relating to the asset are then compared with the IF condition in the knowledge-based to determine intermediate facts, which can then be stored in the Real World View of Data or a diagnostic conclusion. The third component is the **Inference Engine**, this performs the analysis by comparing the rules in the knowledge base to the facts stored in the real world view of data. The **Explanation Mechanism** provides justifications and an explanation as to why the inference engine has decided on a conclusion. This component is crucial for the system to be accepted by the user or by industry. Finally, a **User Interface** allows for communication between the user and the system, whether this is for the input of new facts relating to the data or the output of the diagnostic conclusions, this information can also be passed to external programs or systems.

***IF* :**
Datastream A = & *Datastream B* = &
Datastream C = & *Datastream D* =

***THEN* :**
Fault seven has occured

Figure 2: **Structure of rules stored in knowledge base**

3. SYMBOLIC REPRESENTATION OF KNOWLEDGE

For many industrial applications fault diagnosis involves the engineers following a predefined diagnosis process. Therefore, the expert knowledge has already been acquired to some extent, although this is not always complete enough to be formalised into a set of rules for a rule-based expert system. There is often a significant amount of subjectivity involved when the engineers asses the problem, due to their own experiences with the asset, rules of thumb, or different formal training. However, at a high level, they are often looking for standard data trends such as increases or decreases in specific data, or an increased noise or fluctuation. There is often no prescribed quantitative information relating to these trends that they analyse, such as how much increase or decrease relates to a specific rise or fall, or how much increase in fluctuation relates to a signal moving from stable to fluctuating, as these values will change based upon multiple factors, such as the type of machine, the age of the machine and the operational profile of the machine. Therefore, before the knowledge can be formalised into a rule-based expert system this additional knowledge must be acquired from the domain experts through the knowledge elicitation process. There are several different approaches for performing this knowledge elicitation [2], some of these include: structured and unstructured interviews: observation through active participation or focused observation; and task or decision analysis. For complex problems this is an extremely time-consuming task, this bottleneck in the development of a knowledge-based system has long been recognised and has hence been called the "knowledge elicitation bottleneck" [4].

The rest of this paper focuses on a new methodology to streamline this knowledge elicitation process by simplifying the knowledge into a set of symbols, or common language, that can be easily communicated between domain experts and data engineer. The proposed methodology is a three staged process that involves a minimum of two structured interviews.

3.1 DEFINITION OF SYMBOLS

These symbols were selected as low-level predicates that could be used to broadly describe a time series at any instant. The trends that were selected are shown in Figure 3, these symbols are a stable symbol which relates to normal behaviour, a rise and fall symbol for an increase or decrease over a specific time with a specified limit, and a fluctuating symbol for an increase in noise present in the signal.

(a) Stable (b) Rising (c) Falling (d) Fluctuating

Figure 3: Four selected symbols/trends

These symbols were selected as they are the most basic trends that can be present in time series data, and any complex trend can be constructed from these primitives. This allows for the domain experts to easily communicate the diagnostic process they follow using a common language.

3.2 DEFINITION OF RULE BASE

The next stage of the process is to set up a structured interview with the domain experts to agree on a definition of the rule base. This requires a definition of the individual faults that are being analysed, the specific datastreams necessary to determine those faults, and the associated trends for each of those datastreams. For each fault, a table can be produced that contains all the information discussed above, the example format of this table is shown in Table 1. Additionally, any comments that the engineers can provide at this stage will also prove to be extremely useful, this develops a rationale behind each piece of knowledge and for example could be: the physical reasoning behind the associated trends: or clarification on a subset of faults where a full data set or other operational influence is unavailable to fully diagnose a specific problem.

Following the meeting, each of the tables for the individual faults are combined to produce an overall rule-base for the asset being analysed, an example of this is shown in Table 2. Regarding system development, it is now possible to construct a prototype rule-based expert system using placeholder values for the quantitative parameters relating the each of the individual trends, which will be set in the next stage.

Table 1: **Example format for individual fault diagnostic rules.**

Datastream	Trend	Comments
Datastream A		
Datastream B		
Datastream C		
Datastream D		
Datastream E		

Table 2: **Example format for asset specific rule-base.**

Cause	Datastream A	Datastream B	Datastream C
Fault 1			
Fault 2			
Fault 3			
Fault 4			
Fault 5			

3.3 DEFINITION OF PARAMETERS

Having defined the necessary symbols to accurately interpret related data streams; agreed with the domain experts the individual faults and the associated trends used to asses these faults: the next stage is for all this information to be tabulated and parametrised. Subsequently, a second structured interview is arranged to determine the individual magnitudes for each specific trend associated with each specific rule. The previously mentioned symbols that are now shown in Figure 3 are regarded as the most basic trends that are present in the data. The expert knowledge that is required to qualify the diagnostic rules shown in Table 2 is the subtle differences in the trends in Figure 3. For each symbol various parameters must be assigned to them that accurately describe the possible variations in the symbols for different rules, this is shown in more detail in Figure 4 and Table 4.

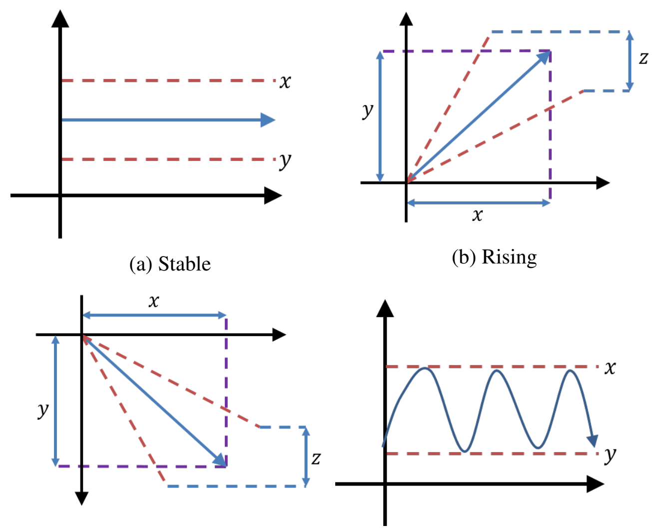

(a) Stable (b) Rising (c) Falling (d) Fluctuating

Figure 4: Definition of parameters for subtle difference in symbols/trends

This information and the corresponding parameters can be tabulated and presented to domain experts in a structured interview knowledge elicitation session. An example of this structure is shown in Table 3. The parametrisation of the knowledge allows for efficient and accurate capture of the domain specific knowledge by focusing the domain experts on a simplified version of the problem. This also facilitates the ease of formalising this knowledge into the rules for a rule-based expert system, without the need to listen to hours of audio recording or to interpret the engineer's answers to specific questions.

Table 3: **Example structure for rule specific table to be completed during knowledge elicitation session**

Rule	Parameters
Datastream A - Rising	$x =$, $y =$, $z =$
Datastream B - Falling	$x =$, $y =$, $z =$
Datastream C - Fluctuating	$x =$, $y =$
Datastream D - Stable	$x =$, $y =$

Table 4: **Description of parameters for quantifying the subtle difference in the trends**

Trend	Parameter	Description
Stable	x	The upper limit in variation for a signal to be consider stable
	y	The lower limit in variation for a signal to be consider stable
Rising	x	The time period for the rise to occur over
	y	The minimum change in the measurement
	z	Two values relating to the spread of the x and y parameters
Falling	x	The time period for the fall to occur over
	y	The minimum change in the measurement
	z	Two values relating to the spread of the x and y parameters
Fluctuating	x	The upper limit for the transition between stable to fluctuating
	y	The lower limit for the transition between stable to fluctuating

3.4 IMPLEMENTATION

After gathering all the expert knowledge from the knowledge elicitation meetings, the methodology proposed to evaluate the diagnostic rules on time series data is to first segment the data into specific time regions, see Figure 5. Splitting up the time series data into various timesteps, or segments, based on the information provided by the expert, each data stream/channel timestep is assigned a symbol which is either rising, falling, fluctuating or stable.

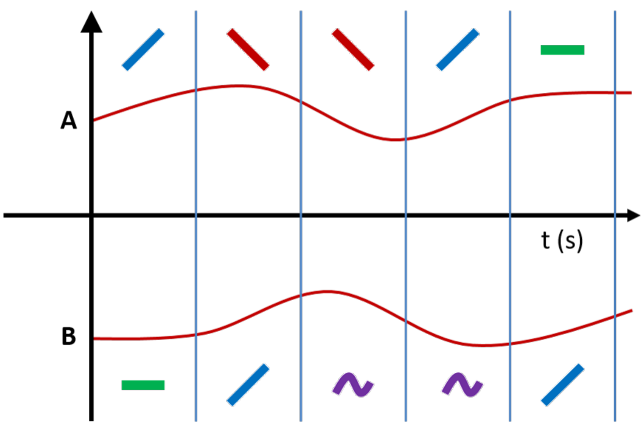

Figure 5: **This figure is an example obtained from. Captions must be below the figures.**

The assigning of the symbols is performed using a technique based on Signal to Symbol transformation [8] which has been successfully used for rotating plant in the nuclear industry previously [3]. For this application, the symbols are assessed by first calculating the average of the first and last 10% of the timestep, a comparison is then performed to determine which of the following four categories best describes the timestep. The categories are defined as: **Stable** less than 50% of the data is out with the thresholds set by x and y, **Fluctuating** more than 50% of the data is out with the thresholds set by x and y, **Rising** the mean value for the first 10% of the data is greater than y times the mean value of the last 10%, and **Falling** the mean value for the last 10% of the data is greater than y times the mean value of the first 10%, where x, y and z are defined in Figure 4. Algorithm 1 shows the pseudocode for this calculation and an example is shown in Figure 6 for two generic pressure and temperature datastreams.

Algorithm 1: **Signal to symbol transformation. Where x, y and z are defined in Figure 4**

if *50% of data (< x*mean(data) or > y*mean(data))* **then**
  Result: Stable
else if *50% of data (> x*mean(data) or < y*mean(data))* **then**
  Result: Fluctuating
else
  Calculate average of first and last 10% of data for x period of time;
  if *First < y*Last* **then**
    Result: Rising
  else if *First > y*Last* **then**
    Result: Falling

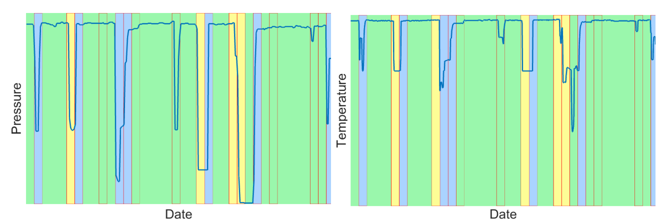

(a) (b)

Figure 6: Example of signal to symbol transformation for (a) pressure and (b) temperature. (Green - Stable, Blue - Rising, Yellow - Falling, and Red border - Fluctuating)

Having formalised the rules and implemented the Signal to Symbol Transformation as described above, it is possible to detect faults in near-realtime across multiple data sources. As new timesteps are input into the system each datastream can be assigned a symbol. When all datastreams have been assigned a symbol the expert system can then compare the symbols with the rule base to determine if any fault has occurred. If a positive correlation occurs this timestep is marked with the corresponding fault type. Over time it is possible to build up a history of any faults that have occurred historically in the asset, an example of this over a small time period is shown in Table 5.

Table 5: **Example processing of 4 datastreams for 5 timesteps.**

Datastream	T1	T2	T3	T4	T5
A	∿	—	\	—	∿
B	—	/	—	/	\
C	—	∿	—	—	/
D	∿	—	\	\	—
Fault	N/A	7	6	N/A	N/A

4. CASE STUDY: AGR BOILER FEED PUMPS

Following the proposed methodology, a case study was performed for data gathered from boiler feed pumps of an advanced gas-cooled (AGR) reactor in the UK. This case study was selected as these assets are critical to the continued operation and electrical generation of an AGR power station, therefore, it is imperative that the pumps are monitored for any abnormal behaviour that may contribute to accelerated plant degradation or to tripping the plant which would result in reduced or zero power generation.

The diagnostics rules for the asset were supplied by the domain experts at the beginning of the project. This determined each data stream necessary to diagnosis a given predefined list of faults, These 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, speed, vibration, and flow. The additional data required to formalise this knowledge into the rules for a rule-based expert system were acquired through knowledge elicitation meetings following the proposed methodology.

Having captured and formalised the domain experts knowledge it was possible to develop a prototype demonstrator for quickly and accurately identifying faults in the boiler feed pump data in real-time. To implement the knowledge base, all the knowledge that was acquired from the knowledge elicitation meetings is stored in a Microsoft Excel spreadsheet. It was stored in this format so that any engineers using the system would easily be able to view all the captured knowledge and therefore provides greater acceptance of the system and also that the captured knowledge is correct. If the analyst wishes to amend a specific rule or add a new fault type, this can be done by editing the spreadsheet directly. Any updates that are made to the rule-based are automatically detected by the system, and displayed to the analyst in the "Changes to Rule Base" panel in formatted text, see Fig. 7. Currently, any amendments made to this file will only be saved for the same session, however, the functionality to load the rule base from historical sessions can be added in the future. This functionality will also require for validation of any new, or amended rules using historical data to ensure that the quality of the knowledge base is maintained. When the analyst is satisfied with the knowledge stored in the rule base the analyst can begin to perform the analysis. The average analysis for the current rule base (37 faults and 10 datastreams) takes less than 0.5 seconds to complete one timestep of the analysis. When potential faults are detected, they are displayed to the analyst in the "Potential Faults" panel, see Figure 7. The date of any fault detected is displayed and the analyst can select this to open up a new window (Fig. 8). This presents the analyst with a drop down menu that contains all faults detected and the associated data streams used in the analysis. The analyst can then view each of these datastreams to display the data covering the time in question on the axis to the right of this window and confirm that the correct fault has been identified.

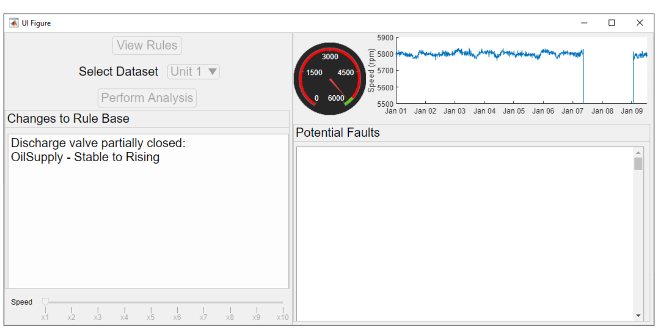

Figure 7: Main GUI for automated boiler feed pump diagnostics

This methodology has allowed for the rapid development of a rule-based expert system for the purpose of fault detection in boiler feed pumps. Due to the novel approach adopted for the knowledge elicitation process, it was possible to minimise the amount of time required from the domain experts but still accurately elicit all the knowledge necessary to develop the system.

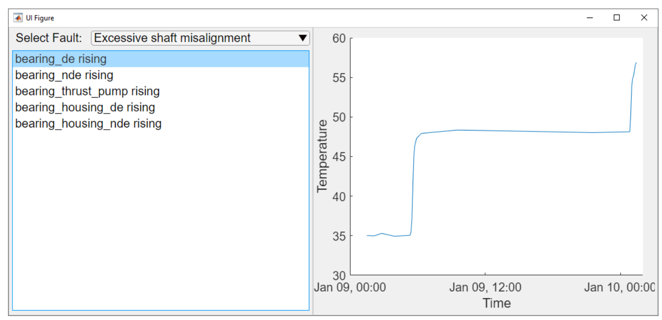

Figure 8: Fault justification window for boiler feed pump diagnostics

5. CONCLUSIONS & FUTURE WORK

This paper has proposed a new approach to knowledge elicitation for the development of a knowledge-based fault detection system, specifically a rule-based expert system. The benefits of knowledge-based systems over data-driven approaches are the increased explicability, however, the increased cost in the development time has been highlighted as a disadvantage. The methodology discussed in this paper attempts to reduce the burden placed on the domain experts by streamlining the knowledge elicitation process, the most time consuming part of developing an expert system. Through the use of symbolic representation of knowledge and the parametrisation of these symbols, it was possible to set out a framework to follow for these streamlined knowledge elicitation sessions.

Using this framework it was possible to develop a rule-based expert system for boiler feed pumps from an AGR power station in the UK. Having further developed the expert system beyond the knowledge elicitation process it has been possible to implement all 37 faults that occur on the boiler feed pumps for the corresponding 10 datastreams. The resulting system can detect faults in the data in real-time due to the segmentation of timesteps into symbols and the efficient inference engine deployed in CLIPS (C Language Integrated Production System) [11].

Future work will involve the development of a human in the loop automated system to improve the elicited knowledge during the system operation. By initially using the methodology discussed in this paper to set the initial parameters for the knowledge and the formalisation of the rule base it should be possible to develop an active learning system [7] to query the analyst to determine any false positives. These labelled false positives will then be used to amend the current parameters to improve the overall system performance. Due to the change in the system knowledge a further piece of work will also explore how to verify and validate this new knowledge using historical data without the need for the domain experts to manually label all historical occurrences.

ACKNOWLEDGEMENTS

This work was funded by the Engineering and Physical Sciences Research Council under grant EP/R004889/1.

REFERENCES

[1] Ayodeji, A., kuo Liu, Y., and Xia, H. (2018). Knowledge base operator support system for nuclear power plant fault diagnosis. *Progress in Nuclear Energy*, 105:42 – 50.

[2] Cooke, N. J. (1994). Varieties of knowledge elicitation techniques. *Int. J. Hum.-Comput. Stud.*, 41(6):801–849.

[3] Costello, J. J. A., West, G. M., McArthur, S. D. J., and Campbell, G. (2012). Self-tuning routine alarm analysis of vibration signals in steam turbine generators. *IEEE Transactions on Reliability*, 61(3):731–740.

[4] Cullen, J. and Bryman, A. (1988). The knowledge acquisition bottleneck: time for reassessment? *Expert Systems*, 5(3):216–225.

[5] Grosan, C. and Abraham, A. (2011). *Rule-Based Expert Systems*, pages 149–185. Springer Berlin Heidelberg, Berlin, Heidelberg.

[6] Guan, X. and He, J. (2019). Life time extension of turbine rotating components under risk constraints: A state-of-the-art review and case study. *International Journal of Fatigue*, 129:104799.

[7] Holzinger, A. (2016). Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain informatics.*, 3(2):119–131.

[8] Nii, H. P., Feigenbaum, E. A., and Anton, J. J. (1982). Signal-to-symbol transformation: Hasp/siap case study. *AI Magazine*, 3(2):23.

[9] O'Hagan, A. (2019). Expert knowledge elicitation: Subjective but scientific. *The American Statistician*, 73(sup1):69–81.

[10] Tang, S., Yuan, S., and Zhu, Y. (2020). Deep learning-based intelligent fault diagnosis methods toward rotating machinery. *IEEE Access*, 8:9335–9346.

[11] Wygant, R. M. (1989). Clips — a powerful development and delivery expert system tool. *Computers & Industrial Engineering*, 17(1):546 – 549.

[12] Xiao, C., Jin, Y., Liu, J., Zeng, B., and Huang, S. (2018). Optimal expert knowledge elicitation for bayesian network structure identification. *IEEE Transactions on Automation Science and Engineering*, 15(3):1163–1177.

[13] Young, A., West, G., Brown, B., Stephen, B., and McArthur, S. (2019). Improved explicability for pump diagnostics in nuclear power plants. 2019 ANS Winter Meeting and Nuclear Technology Expo ; Conference date: 17-11-2019 Through 21-11-2019.