This paper appeared in ANZIIS'93

Embedding Neural Networks in On-line Monitoring Applications

Mathew J. Boek[1], Jacob L. Cybulski[2] and Andre S. Szczepanik[3]

[1,3] Vimac Pty. Ltd.
10 Lynden Gve
Mt Waverley, Vic 3149, Australia
Ph: (03) 888 2207, Fax: (03) 807 4368

[2] Dept of Comp. Sci. and Eng.
La Trobe University
Bundoora, Vic 3083, Australia
Ph: (03) 479 1270, Email: jacob@latcs1.lat.oz.au

Abstract

This paper describes our approach to the application of intelligent information processing in a large engineering system for machine condition monitoring. The paper starts with an overview of the vibration signature analysis methodology currently in use by maintenance engineers to assess machine condition. We then describe an alternative approach utilising an experimental neural network system capable of automating certain aspects of analytical and diagnostic tasks performed in the monitoring process. The discussion is subsequently extended to incorporate on-line and real-time monitoring, and the use of knowledge and expert systems to enhance the currently taken approach.

1. Introduction

Modern machine condition monitoring (Braun 1986) is based on the analysis of machine vibration and is usually performed by skilled engineers. It involves the collection and analysis of vibration signals to identify patterns of signal features which indicate deterioration in machine condition. Traditionally this procedure is conducted with the use of either a tape recorder or a portable data collector to record data to be later analysed off-line; we have developed software supporting this traditional approach and have gained experience in its application to specific industrial problems. More sophisticated systems allow on-line, real-time data collection and analysis; currently we are developing such a system.

The method has many advantages over other more intrusive methods of assessing plant condition as it can be performed during the machine's normal operation without disrupting the production cycle. The technique involves the following basic steps:

- **Machine Specification.** Before any vibration or non-vibration data could be collected and analysed, an engineer must define the fundamental characteristics of the machine and its operating environment. This usually includes machine design, its natural and forcing frequencies, position of measuring points, types and periodicity of measurements, alarm levels and signal masks, trending algorithms and parameters, inter-signal correlations, maintenance routines and their impact on machine condition, etc.

- **Signature Recording.** As the first step in the monitoring cycle, an engineer records the machine's vibration signature, i.e. a collection of measurements characterising the machine's good condition, which can be used as a reference in all future data collections.

- **Data Acquisition.** This step involves progressive recording of vibration signals as measured by transducers, storage of signals on a magnetic tape, signal amplification with a charge amplifier, digitising, averaging and some data reduction through FFT by a spectrum analyser or a portable data collector (which also eliminates the need for an intermediate tape recorder).

- **Feature Analysis.** Previously specified signal features are identified and their values are calculated from the collected vibration data. The feature values of greatest interest usually include: the spectrum increase above the alarm, mask or the signature, the RMS value at specific spectrum points or their ranges, peak values over the entire

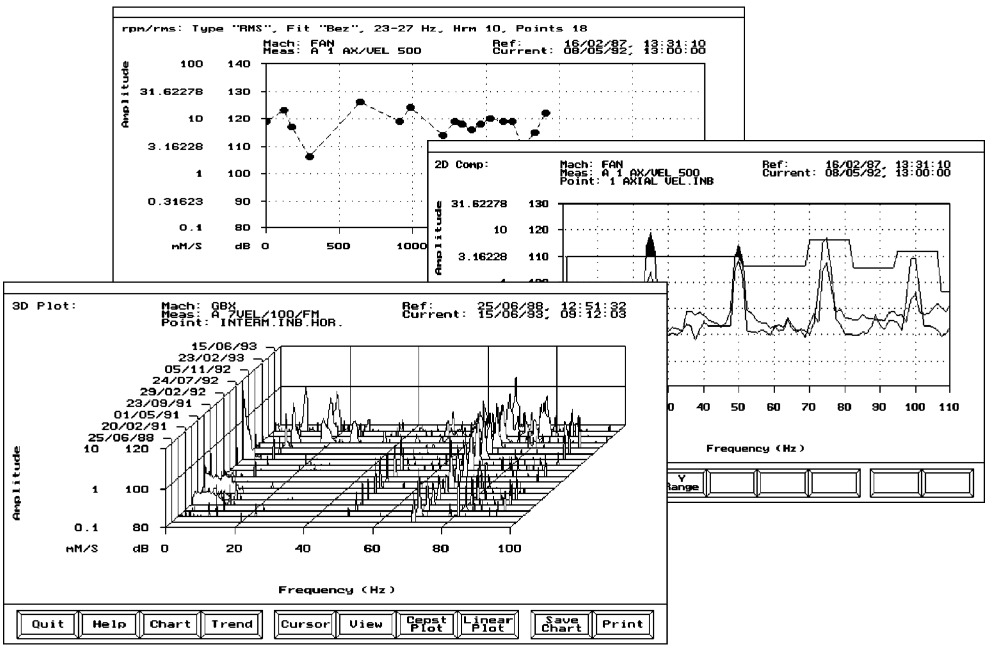

Figure 1 - Sample vibration data and its features

spectrum, phase difference between measuring points, trend values of specific parameters, etc. This phase is usually aided by a software tool calculating the relevant feature values and storing them in a database, allowing their visualisation, analysis and some additional manipulation (Cf. figure 1).

- **Condition Analysis.** Collected signal features need to be interpreted in terms of machine condition. This task usually requires services of experienced engineers with knowledge of not only engineering principles but also of machine characteristics and previous history. Condition is usually assessed in terms of component wear, risk assessment, prediction of time to failure, or impact of current machine condition on product quality.

- **Recommendations.** Based on machine condition assessment, a set of recommendations is produced to assist in the scheduling of machine maintenance.

In practice the association between complex patterns of signal features and fault conditions developing in machines can only be performed by highly skilled mechanical engineers. Their experience cannot be easily explained in terms of fundamental engineering principles or be transliterated into a set of written procedures which could become the basis for writing a computer program to be used in monitoring tasks. This reliance on human expertise in the machine monitoring cycle renders the construction of an on-line system for machine condition diagnosis virtually impossible. Nevertheless, modern expert system techniques provide tools, methods and techniques for capturing human expertise (McGraw and Harbinson-Briggs 1989), which could be used for building fully automatic on-line monitoring systems (Milne 1990) or off-line fault detection, identification and rectification systems (Bannister, R.H. and Moore, M.P. 1986). Neural networks (Rumelhart and McClelland 1986) offer an alternative approach to expert system knowledge acquisition techniques; they offer the additional advantages of an ability to learn implicit facts from examples and the simplicity of their implementation.

2. Neural Networks in Machine Condition Monitoring

The condition analysis phase of machine monitoring can be viewed as a problem not dissimilar to that of pattern classification, where a huge collection of highly correlated (and in some cases redundant) features of vibration signals

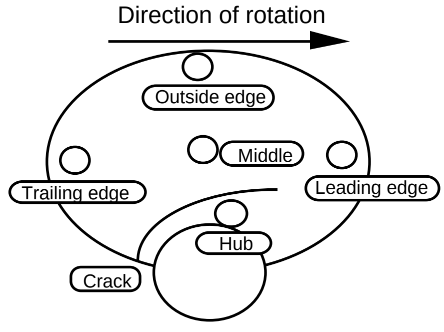

Figure 2 - Weights and a crack in fault simulation

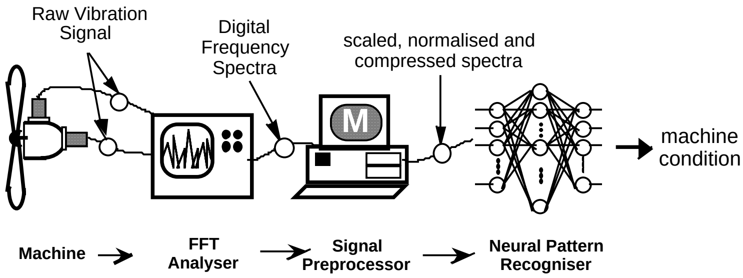

Figure 3 - Neural network vibration signal classification

could be associated with a number of (frequently fuzzy) machine conditions.

An experimental automatic classification system for detection and diagnosis of certain defects in rotating machinery was proposed and implemented with the use of neural networks (Boek 1992). The system was tested using data generated by simulating faults (load imbalance, shaft imbalance, and a cracked blade) commonly found in heavy industrial fans by attaching small weights to the fan blades or cutting a blade on a 3-speed, 3-blade oscillating desktop fan (figure 2).

Two accelerometers were attached to the fan, one in the radial and one in the axial plane of the shaft (cf. figure 3). The fan was run at a constant speed, and the signals from the accelerometers were processed by a charge amplifier, which fed the modified signals to a signal analyser. The analyser digitised the signal, transformed it from a time signal into a frequency spectrum, averaged a number of such spectra, and then passed the averaged spectrum to condition monitoring software (the M system, developed by Vimac Pty Ltd) for feature extraction, storage and analysis. The set of extracted features was then presented as input to the neural network classifier (based on the PDP package - McClelland and Rumelhart 1988) for training and assessment of machine condition (cf. figure 3).

In our experiments the spectrum consisted of 400 values, which ranged between 0 and 130 dB. As every state of the fan was represented by two spectra, one from the vertical plane of vibration and one from the axial plane, a total of 800 values were available to describe each snapshot of fan condition. It is known, however, from reference to engineering principles that for rotating machines (such as a fan) these values consist mostly of noise. It is also known that the vibrational information important for detecting rotor defects is carried in the amplitudes of the rotational frequency and its associated harmonics. Thus, some of our experiments used only the peak values of the harmonics of the rotational frequency as a representation of the fan condition. This reduced the dimensionality of the input space from 800 to just 18. All spectral peaks extracted were scaled from the range 0..130 to a value between 0 and 1.

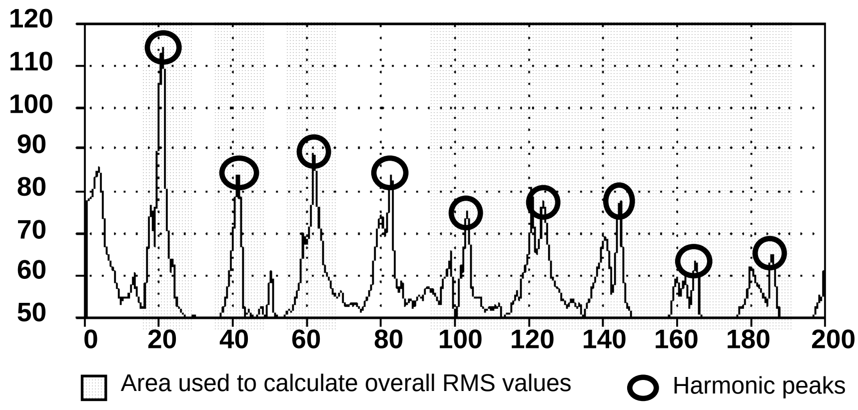

Figure 4 - Spectral feature extraction

In an effort to further reduce the dimensionality of the input space, in our experiments we used the first 3 harmonics, with the higher harmonics summarised by a single *overall* value (the average RMS value of the high frequency end of the spectrum); see figure 4. Although dimensionality reduction was achieved, it was possible that some potentially important information was lost. As it is virtually impossible to determine in advance the precise feature values that characterise a specific machine condition the features selected for our experiments were determined by trying different combinations of harmonics, and observing which appeared to vary the most for different fan conditions.

Training was conducted using the back-propagation algorithm (Rumelhart, Hinton and Williams 1986). The neural network used for this project had one hidden layer containing between 5 and 20 processing units. The learning rate was fixed at 0.02 and the momentum control parameter at 0.9.

The performance of the network was tested on five classification tasks: detection of imbalance, classifying imbalance into subtypes (deposit on 1/2/3 blades, slight/medium/heavy or concentrated/distributed deposit) and distinguishing between imbalance and a cracked impeller blade. The average test set result across all classes for the first four tasks are presented on figure 5. As expected, the best results were obtained in the detection of imbalance, as this simple task could also be done by a simpler system than a neural network. In general, the results of the experiments to classify type of imbalance were quite poor, although in most cases the network detected that imbalance was present, but could not distinguish one type of imbalance from another. In order to put the results for these experiments in perspective, it should be noted that some of the

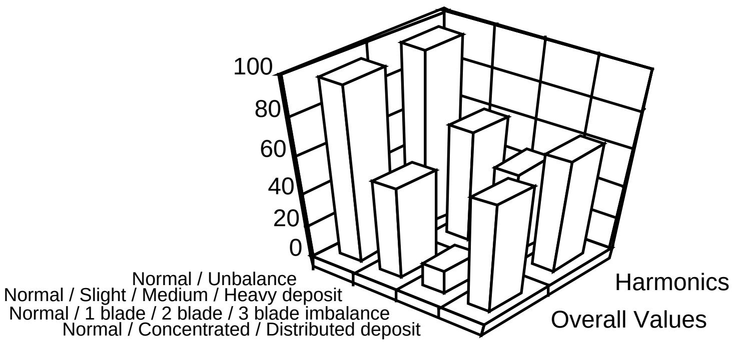

Figure 5 - Average test set performance

classifications being attempted were of questionable practical importance; it is far more important to know that there is imbalance of some sort present in a machine than to know for example that the deposit is on one or two blades. This poor performance may also be explained by the limited number of examples available for training and testing the network. Good results, however, were obtained when the network had to distinguish between physically different faults, i.e. imbalance vs a crack in the blade. The good results for these experiments are the most important, as they show that a multi-layer perceptron neural network trained with back propagation is capable of developing rules to classify the complex signals provided by vibration spectra (with suitable preprocessing and feature extraction) to certain level of accuracy. The fact that the network coped successfully with data generated from a real fan, together with all the noise that is unavoidable in such an environment, also confirms to a certain degree the much heralded noise tolerance of neural pattern classification systems.

These experimental results should translate well to fault detection in and diagnosis of industrial fans, since they are more rigid in structure, and thus tend to produce less noise in the vibration signal. The transducers used in collecting the data for these experiments were attached to the plastic housing of the motor, and the impeller itself was made of flexible plastic. This flexibility introduces more noise into the signal, and thus in some sense the noise could be seen as “worst case” noise in these experiments.

A study was performed to compare the performance of an experienced mechanical engineer on the same data sets that were used for training and testing a network (figure 6 shows the results on the training set). This was done for the task in which the network had to classify the normal, imbalance, cracked

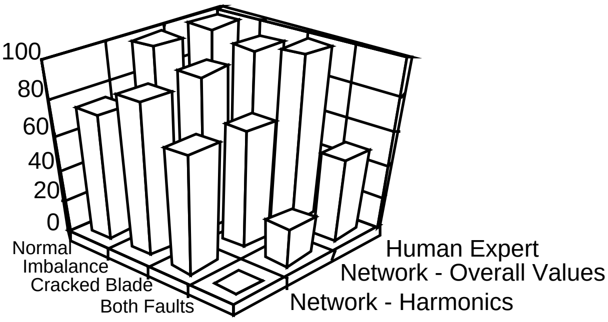

Figure 6 - Network vs human

impeller, and both-faults-together states of the fan. The human expert clearly outperformed the network for the cracked blade and both-faults-together classes. However, a network using overall value spectral representation obtained almost equivalent performance to that of the expert for the normal and imbalance classes. It should be noted that this result was obtained by the expert having access to the entire spectra of a faulty and a good machine condition as opposed to the restricted set of features used by the neural network. Further work in extracting better features could lead to improved performance by the network.

As the automation of machine condition analysis significantly reduces the dependence of monitoring tasks on vibration analysts, the proposed method is being investigated by the authors for the development of a system for continuous, on-line monitoring of machine vibrations. The software, currently under development, will be capable of rapid classification of vibration signals into a set of high-level messages and alarms associated with the current machine condition, process parameters and product quality. Due to the fact that vibration is often the very first sign of malfunction, the system will give the process operators an opportunity to detect, accurately assess and quickly correct deficient manufacturing processes well before conventional monitors could register a fault developing in the machine, process or product.

3. Other Applications of Intelligent Processing in Machine Condition Monitoring

Although machine condition analysis is a critical phase of machine condition monitoring, there are many other aspects of the monitoring cycle that cannot benefit from the use of neural networks, either due to type of processing to be performed or due to the nature of data to be manipulated.

Certainly, the elementary tasks of data collection, storage and presentation are eminently suited for conventional data processing and algorithms. Other tasks, however, such as machine specification, task scheduling, risk assessment, fault prediction, and recommending a course of action, may greatly benefit from the use of intelligent process models and more elaborate information representation.

For instance, we have developed a machine type database, supporting type classification and machine instantiation with inheritance. Such machine organisation greatly reduces the initial effort of machine specification, which normally includes a description of all possible measurements to be performed on the machine measurement points, down to the details of analyser settings, data filtering options and trend calculation. Thus, the use of machine taxonomies allows the popular machine types and specific brands to be quickly incorporated in a particular customer plant database.

The real-time, multi-tasking operation of an on-line machine monitoring system requires intelligent scheduling of measurement, analysis, and diagnostic tasks. In our system the scheduler is controlled by an agenda, which outlines the tasks of which higher level monitoring activities are composed, information and control flow between these tasks, concurrency specification, timing information/constraints, etc. Currently, the intelligence is supplied by the agenda designer, but in future it will be possible to automatically generate (or provide assistance in the design of) the agenda from a specification of the machine to be measured, its operational parameters, and for some machines the product or process parameters.

Other monitoring tasks, such as risk assessment, fault prediction and action recommendation, are grounded in the knowledge and experience of expert condition monitoring engineers, rather than based on a clever algorithm, past case classification, or elaborate data structures. These tasks are being considered for the implementation in a rule-based machine condition monitoring expert system which models the knowledge acquired from the maintenance and monitoring engineers.

4. Conclusions

Vimac developed and is currently distributing an off-line machine condition monitoring system (The M System). At the moment, Vimac is developing a large on-line monitoring system, to be used by several Australian heavy-industry companies. This system will be used as an environment for the application of the neural network and expert system based intelligent diagnosis system. The experimental results from the prototype neural-network-based fault classification system show the method can be feasibly used in the vibration monitoring. Vimac has compiled a data base consisting of a very large number of machines and their measurement points. This data could be used to train a neural network. All machines have been analysed and diagnosed, and described in some computer readable form. This has taken many man-years of effort but was conducted as a matter of Vimac's normal business operation. The major emphasis in the development of the system's intelligent component at the moment is focused on the interpretation of data and feature extraction.

Based on Vimac's experience, the process of embedding an intelligent information processing in a product requires a lot of research and development effort, but this is only a small part of solving the general business problem. Most of the costs are involved with the establishment of the application system - hardware, testing, software development, data collection, feature selection and manual analysis of the data - rather than with the development of the intelligent software.

5. References

Bannister, R.H. and Moore, M.P. (1986): General Rotational Machinery Expert System. In *Research and Development in Expert Systems III: Proceedings of Expert Systems '86*, Bramer, M.A. (ed), British Computer Society Specialist Group on Expert Systems, Brighton, England, pp. 140-151.

Boek, M.J. (1992): An Application of Neural Networks to Machine Condition Monitoring, Master of Applied Science Thesis, Faculty of Applied Science, Royal Melbourne Institute of Technology, Melbourne, Australia.

Braun, S. (Ed) (1986): Mechanical Signature Analysis: Theory and Applications, Academic Press, London.

McClelland, J.L. and Rumelhart, D.E. (Eds) (1988): Explorations in Parallel Distributed Processing: A Handbook of Models, Programs and Exercises, MIT Press, Cambridge, Mass., Vol. 3.

McGraw, K.L. and Harbinson-Briggs, K. (1989): Knowledge Acquisition: Principles and Guidelines, Prentice-Hall Int. Ed., London.

Milne, R. (1990): Amethyst: Vibration Based Condition Monitoring. In *The Handbook of Expert System Applications in Manufacturing*. McGraw Hill, Keyes, J. and Maus, R (eds).

Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986): Learning Internal Representations by Error Propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition Vol. 1 Foundations*. MIT Press, Rumelhart, D.E. and McClelland, J.L., pp. 318-362.

Rumelhart, D.E. and McClelland, J.L. (1986): PDP Models and General Issues in Cognitive Science. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition Vol. 1 Foundations*. MIT Press, Rumelhart, D.E. and McClelland, J.L., pp. 110-146.