

# 20th Australian International Aerospace Congress

ISBN number: 978-1-925627-66-4

## Successful Predictive Maintenance

Peter D M Brady, Emmanuel Blanchard

MathWorks Australia

Level 6, Tower 2, 475 Victoria Avenue, Chatswood, NSW 2067

### Abstract

Predictive Maintenance is a topic that is fast-growing in popularity but is somewhat of a buzzword with many different implementations. A common and simple starting point involves the use of a “black box” sensor on a plant or machine part to capture readings. While this does provide some results, it misses the opportunities that a full-scale implementation of Predictive Maintenance delivers. Instead, the most successful Predictive Maintenance implementations that we see are fully integrated with data recording systems and automatically handle model lifecycle management. After all, as your plant operates, and loads change, the models will drift to a point that they need to be retrained or even removed and replaced.

This talk showcases best practices and lessons learned from successful projects and high-level walk through of the full workflow of a typical Predictive Maintenance journey. Specifically, we begin with the raw data and early gains made with condition monitoring before concluding with full model life cycle management. Case studies illustrate these steps.

**Keywords:** black box model, lifecycle management, model drift, predictive maintenance

### Introduction

Our strongest lesson learnt from supporting multiple organisations to deploy active systems is that PdM is a *journey*. It takes time to develop mature systems that manage the life cycle plant assets and PdM models. Typically, we see the main components of the PdM journey as:

1. Initial data gathering.
2. Anomaly detection.
3. Condition based monitoring.
4. Static predictive maintenance models.
5. Total life cycle maturity.

### Initial Data Gathering

The most important realisation during the initial data-gathering phase is that some of your assumptions will prove incorrect. In our general PdM task work, we found that the leading indicator of project success is the quantity and quality of the data representing the plant items under consideration. Conceptually, the guidance is simple: gather more than you need. For example, with rotational machinery, you generally need to sample much higher than the basic Nyquist frequency to ensure capturing the high-speed system dynamics generating harmonics. The main complaint with this approach is the sheer volume of data generated across an asset fleet. However, this is just phase one. Established techniques can handle high-frequency data capture at the edge in a mature system.

## Anomaly Detection

Anomaly detection is usually the first number-crunching step of the predictive maintenance journey. Here we start finding challenges with data cleaning, including handling gaps and determining if the signal outliers are real. We would also see the first steps in feature engineering: looking at fundamental time domain statistics such as mean, mode, median, min/max, and frequency space parameters such as peak frequency and power, interpeak spacing, and peak to noise ratios.

Interestingly, this step can be applied from day one, even with very little data. This low-hanging fruit can deliver rapid returns based on the fundamental assumption that most units will be within nominal bounds for a regular fleet. The key is to have enough representative plant units allowing the rapid deployment of *control charts* on a per-

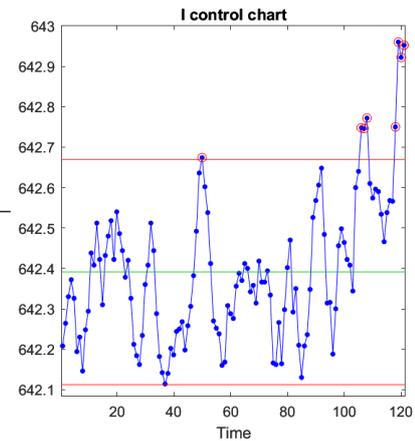
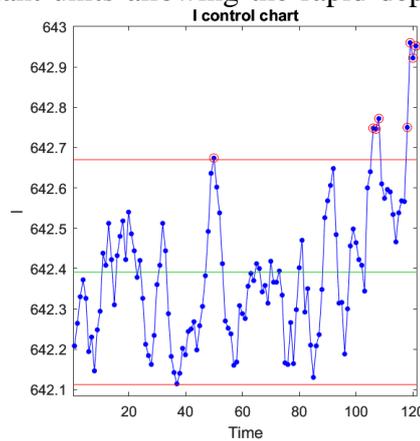


Figure 1 – sample control chart



signal basis, as shown in

Figure 1.

## Condition-Based Monitoring

Condition-Based Monitoring is usually the next logical step from anomaly detection because it relies on a higher volume of clean data on which we have performed feature engineering. We can then deploy unsupervised learning techniques to generate stop light style analyses, as shown in Figure 2. This example utilised Principal Component Analysis for data reduction, but we also regularly see unsupervised clustering algorithms in use.

## Static Predictive Maintenance Models

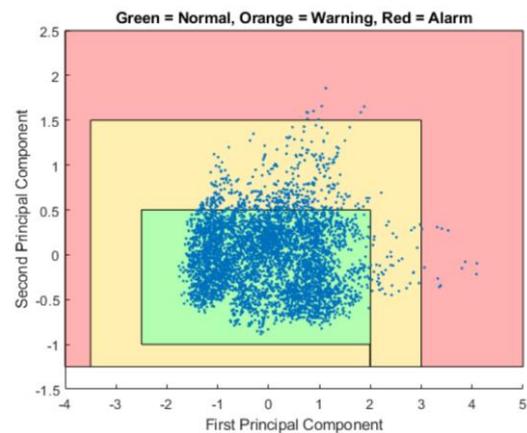
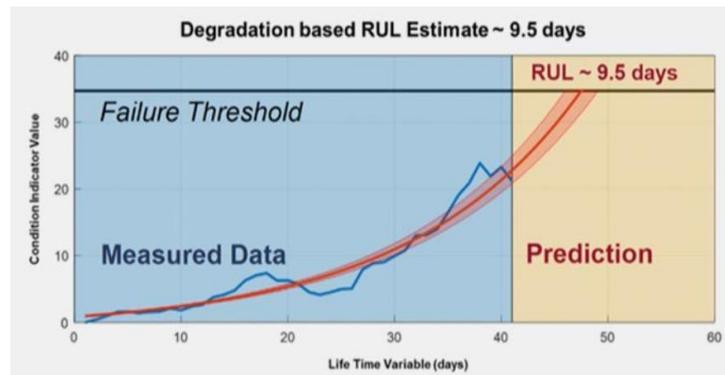


Figure 2 – example of traffic light plots

The next step in the journey – where we find a conceptual block with most conceptual models – is the development of the first full Predictive Maintenance models. Many models that can be used, each tailored to different data availability. For example, you may or may not have run-to-failure data or known failure threshold data [1]. Your data will generally guide your model choices. Broadly, we classify

these into three groups based on data availability: (1) Run-to-failure history; (2) know failure thresholds, and; (3) lifetime data with or without covariates. A typical example of a mid-life application of a remaining useful life model is shown in Figure 3. *But this is just the beginning of your full predictive maintenance solution.*



*Figure 3 – example remaining useful life model with measured data and prediction.*

### Total Life Cycle Maturity

The final step often missed in a predictive maintenance program is the total life cycle that considers model drift and additional insights as more data is gathered, as shown in Figure 4. Model drift is common in machine learning and is easily accounted for, even automatically, through a retrain and deploy process, akin to a quality assurance program, where models are continually trained on incoming data and, if performing better, are deployed.

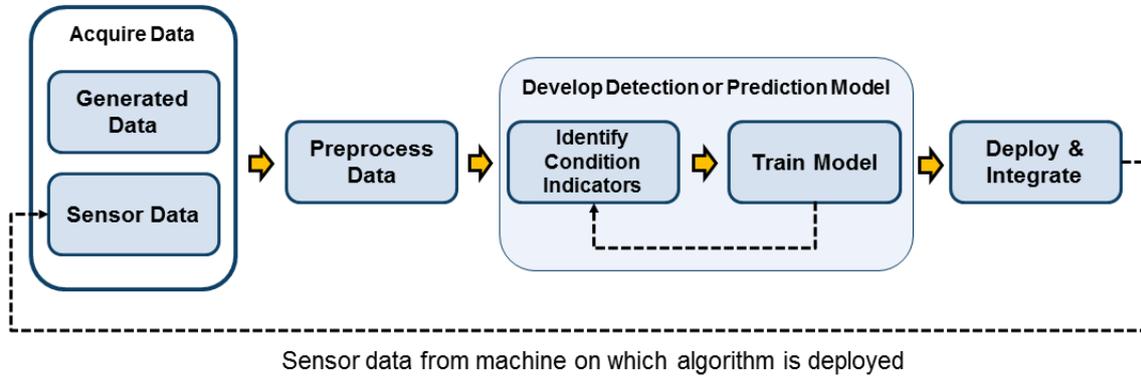


Figure 4 – Continuous improvement process for predictive maintenance algorithms

A final advantage gained at the maturity stage is deployment pathways that allow pushing the models from a data centre to the edge. A stereotypical example in this space is for high-speed applications where the algorithm is deployed onto an edge computing device that works on the full real-time data feed. It is usually configured to signal anomalies in real-time to the operator and to periodically return sample sets, at reduced time intervals, for data archiving and retraining.

### Case Study 1: Wind turbine without run-to-failure data

The case study is representative of a typical wind turbine and is taken from a 2MW wind turbine high-speed shaft driven by a 20-tooth pinion gear [2]. While the data does not run to failure, an observation of the accelerometer readings, Figure 5, indicates that the magnitude of the vibrations is increasing over time, a result confirmed through an examination of the spectral kurtosis, Figure 6.

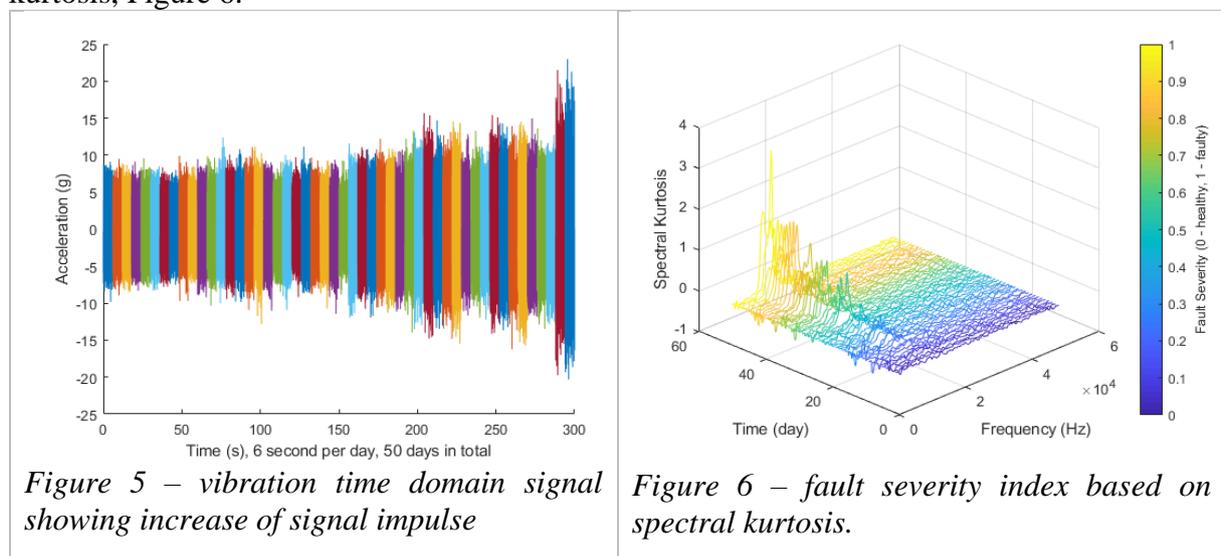


Figure 5 – vibration time domain signal showing increase of signal impulse

Figure 6 – fault severity index based on spectral kurtosis.

However, while this is indicative of a degradation signal to quantify, we next undertake feature extraction and smoothing. As shown in Figure 7, smoothing is necessary as many remaining useful life models rely on monotonic signals with constant changes in direction. Therefore, a short five-period window is generally applicable, as shown in Figure 8.

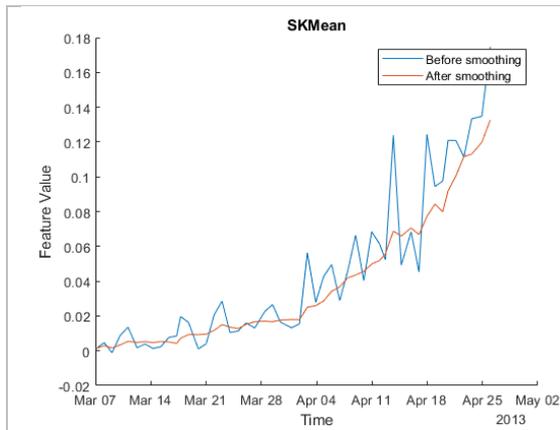


Figure 7 – Comparison of raw and smoothed data

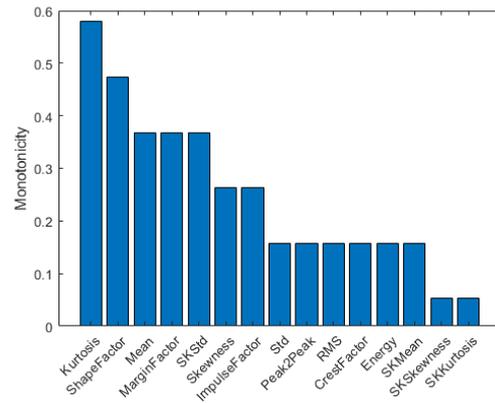


Figure 8 – monotonicity estimates of the engineered features.

The final two steps before fitting a remaining useful life model are to (1) reduce the dimensionality of the engineered features and; (2) synthesise a health indicator from the reduced data. In this case, Principal Component Analysis was applied to reduce the engineered data set with the first two components shown in Figure 9. Finally, after inspection, the first principal component was chosen as a lead health indicator, shown in Figure 10.

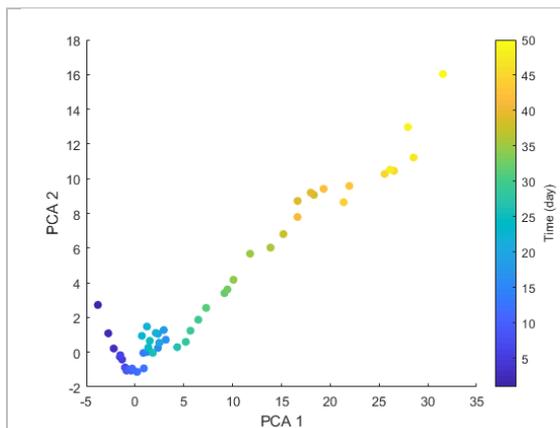


Figure 9 – scatter plot of the first two principal components, coloured by time of measurement.

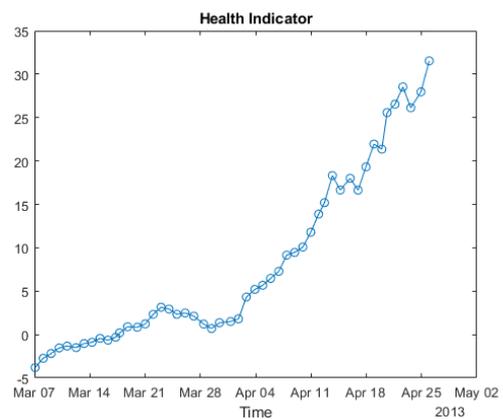


Figure 10 – computed health indicator synthesised from the principal component analysis.

From this an exponential degradation model was fitted and used for prediction, as shown in Figure 11, which is indicating an impending failure on the day 49 data.

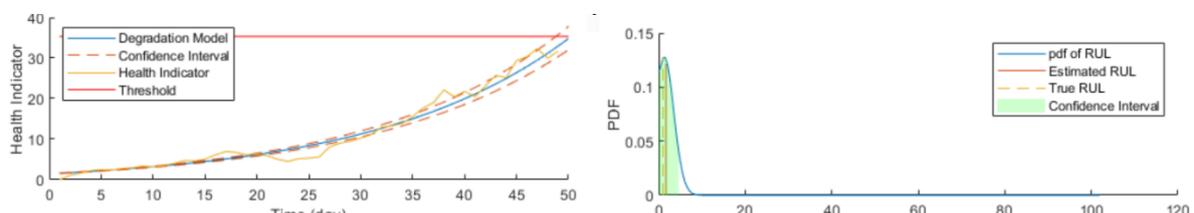


Figure 11 – Remaining useful life model with a prediction from day 49 indicating imminent failure.

## Future Directions

The workflow discussed in Case Study 1 is what we see as a typical and valid workflow for predictive maintenance. However, this is a somewhat “easy” case as there was a consistent signal generated from a fault on the inner race of the main bearing. We see more sensor fusion and deep-learning techniques deployed for cases with more subtle data trends. In Figure 12, for example, we can see the relationship between an instantaneous spectral kurtosis calculation and a short time window approach. The next key step is that these short time window plots are treated as images fed into convolutional neural networks for classification. The classification outcome can be used at a first pass for traffic light-style condition-based monitoring or in and of itself generate signals that can be used to develop health indicators and, in turn, remaining useful life models.

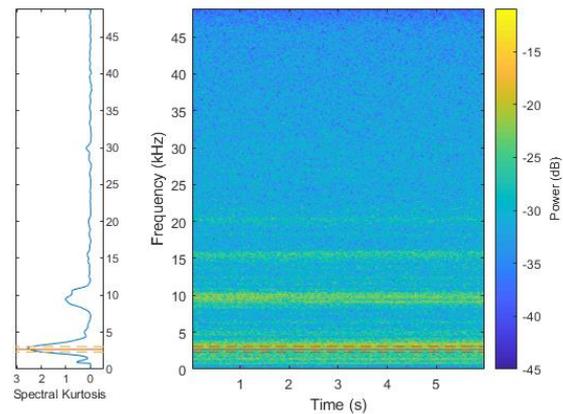


Figure 12 – relationship between instantaneous spectral kurtosis and short time window methods.

## Case Study 2: Deep Learning for Underground Cable Predictive Maintenance

Here we discuss the use of deep-learning methods deployed by IMCORP to predict and locate partial discharge failures in underground electrical cables [3]. For context, over 90% of the failures of underground electrical cables are due to bush-type failures, as shown in Figure 13.

This analysis is particularly suited to deep learning classification as the system for measuring conditions is based on interpreting the peak-to-peak signal returns, as shown in Figure 14. Specifically, a two-layer Long-Short Term Memory network was trained on a data set comprising a four million-labelled time series resulting in a network prediction success rate of over 95% after tuning the hyperparameters.

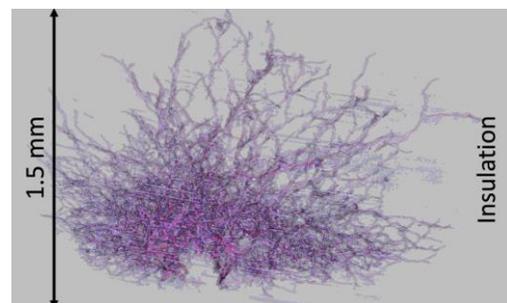


Figure 13 – micrograph of “bush-

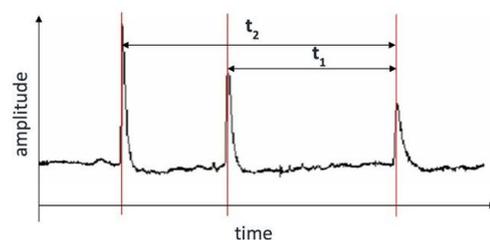


Figure 14 – peak-to-peak timing intervals

## Conclusion

In conclusion, we have presented our learnings from customer projects into the current state of the art in predictive maintenance and future directions and trends we observe. For current methods, we generally observe a trend from project inception with anomaly detection and condition based-monitoring to a mature deployed solution automatically accounting for model drift and quality assurance steps. We see the deployment of more deep learning-based methods with future trends to tease out more subtle results taking additional human effort with traditional signal-based efforts. Deep learning methods can be a double-edged sword because training and development requirements often increase significantly without assurance of improvement. Further, deep networks are generally

deployed to edge devices for which the process is straightforward but with a slight increase in edge compute requirements.

### Acknowledgments

The authors would like to acknowledge IMCORP for allowing publication of the details of their predictive maintenance work and support from MathWorks USA Predictive Maintenance team based in Natick, Massachusetts.

### References

- [1] MathWorks (2022) *What Is Predictive Maintenance?* Available at: <https://au.mathworks.com/discovery/predictive-maintenance-matlab.html> (accessed 2022-10-30)
- [2] Bechhoefer, Eric, Brandon Van Hecke, and David He. *Processing for improved spectral analysis*. Annual Conference of the Prognostics and Health Management Society, New Orleans, LA, Oct. 2013.
- [3] Ziegler, Steffen, IMCORP (2022) *Using Deep Learning for Predictive Maintenance of Underground Power Cable Systems* Available at: <https://au.mathworks.com/company/newsletters/articles/using-deep-learning-for-predictive-maintenance-of-underground-power-cable-systems.html> (accessed 2022-10-30)