Practical Predictive Maintenance Workflows

Peter Brady and Emmanuel Blanchard

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

Abstract

This session will focus on a deep dive of the data cleansing, condition monitoring, and model development phases. To illustrate, we will be using two technical data sets: (1) turbine run to failure data and (2) wind turbine main bearing prognosis. We examine data cleansing methodologies and the different numerical techniques to predict remaining useful life using survival, degradation, or similarity models, depending on your system data.

While a full predictive maintenance solution to estimate remaining useful life requires data, which takes time to gather and process, early gains can be achieved through condition monitoring. We will spend some time dissecting this approach and how to use condition monitoring algorithms to develop traffic light dashboards as a precursor to predictive maintenance maturity.

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

Introduction

This paper walks through a deep dive of a standard predictive maintenance workflow, which nominally follows the following steps:

1. data import and exploration
2. feature extraction and postprocessing
3. feature importance ranking and fusion
4. model fitting and prediction
5. performance analysis.

This relates the central "dashed" region of the block diagram shown in Figure 1.

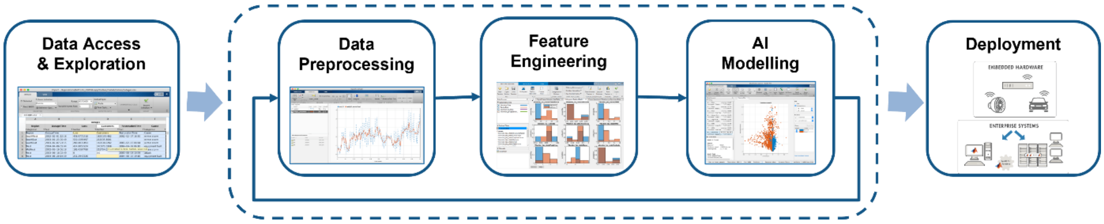

Figure 1 – Block diagram of a typical data driven predictive maintenance workflow.

It is important to note that this paper will not discuss the final, critical steps of deployment and model drift maintenance of a predictive maintenance. Instead, it will focus on the initial development of a model. Conceptually, drift detection is merely an extension of this training workflow, which carries on throughout the plant's useful life, somewhat like a continuous improvement quality assurance process.

The data set that this deep dive is based on was collected from a 2 MW wind turbine high-speed shaft driven by a 20-tooth pinion gear [1]. An interesting facet of this data set is that it shows a very good use case for high-fidelity analysis on the turbine rather than streaming and central processing, specifically for data reduction. To capture the system dynamics, particularly the high-frequency harmonics of a bearing fault, the data was captured at a sample rate of approximately 100kHz. However, a continuous-time signal captured at 100 kHz for a single channel of double-precision data would result in a data stream of 800 kb/s, or ~40GB per day, per turbine. This level of data capture is unnecessary, and the edge controller was programmed to sample six seconds per day, which is a much more manageable 5MB per turbine per day.

Data Import and Exploration

The first step in any predictive maintenance is importing the data for preliminary exploration. Ideally, the exploration would be graphically driven because, intuitively, graphs are easier to interpret than columns of figures. This step is easier for lower dimensional data than higher as humans can more readily interact with three dimensions than four or more. For the remainder of this paper, we will refer to the data set as an *ensemble*, which is where each *member* of the ensemble represents one six second measurement for the remainder of this paper. This ensemble structure allows us to associate metadata at the relative levels of the ensemble, specifically:

1. at the day level, where we can associate lifetime data such as number of days operating, average weather conditions, etc.
2. at the sample level, where we can associate technical data such as sampling frequencies and tachometer specifications.

It can be helpful to visualise the raw signals, as shown in Figure 2, when beginning exploration. In this case, we observe what appears to be a trend in the signal impulsiveness, as the measured amplitude seems to increase through the 50 days of measurement.

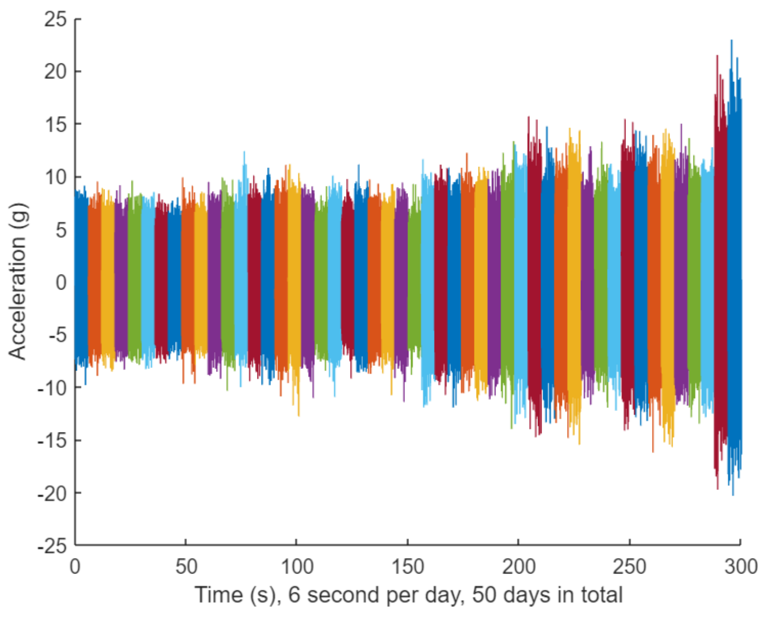

Figure 2 – Raw acceleration plot per ensemble

Next, as this data is from a rotational plant, we know there will be changes in the power spectrum and other spectral quantities as the bearings degrade. In this case, we focussed on spectral kurtosis with a window of 128 samples (~1.3 ms) because spectral kurtosis has been used with success in wind turbine prognosis [2], as shown in Figure 3.

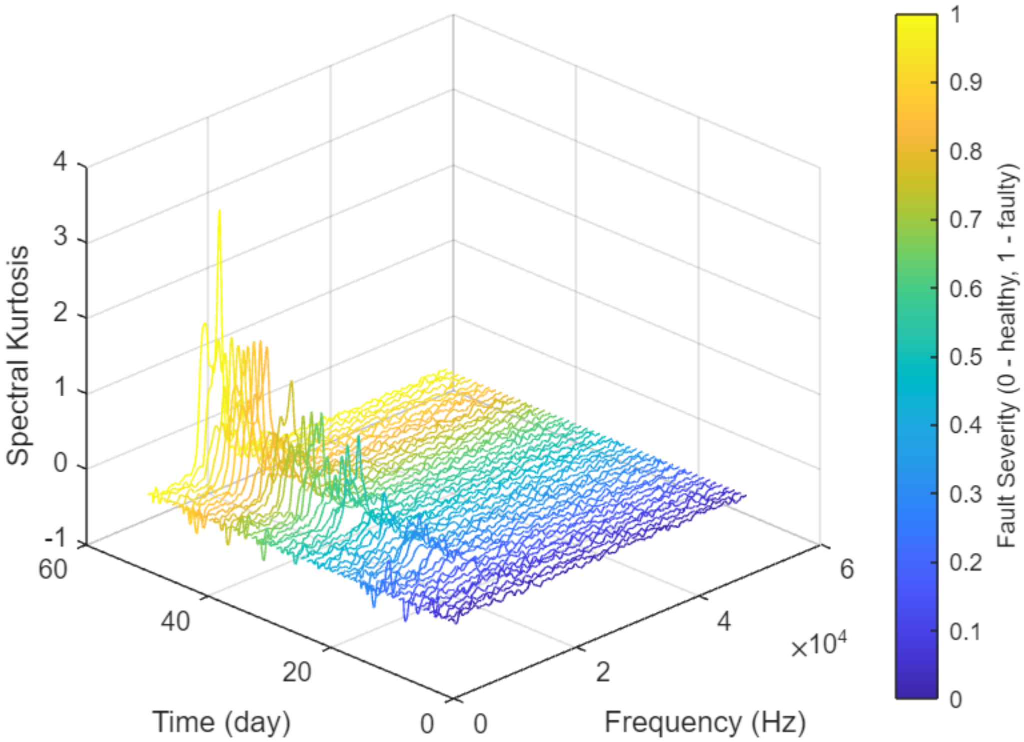

Figure 3 – evolution of spectral kurtosis as a function of operating days

In this figure we have indicated fault severity on a scale of zero to one, indicative only of the normalised date between the start of the measurements and failure and so intended as a visual identification of temporal progress. However, we can see from the peaks in the spectral kurtosis at around 10 kHz through a gradual increase in strength, indicative of a degradation signal. This strongly indicates that spectral statistics would be useful as predictor signals.

Feature Extraction, Importance Ranking and Fusion

With the indicators discussed above, the next step is to again reduce the time domain ensemble components down to representative statistical measures. This is a form of dimensional reduction that shunts the compute time into a pre-processing step, simplifying model creation. The measures chosen were:

1. Time domain: mean, standard deviation, skewness, kurtosis, peak to peak interval, signal RMS, crest Factor, shape factor, impulse factor, margin factor, and energy.
2. Spectral kurtosis domain: mean, standard deviation, skewness, and kurtosis.

Because most remaining useful life models assume a monotonic trend, the raw indicator signals must be smoothed, as shown in Figure 4, and then ranked on monotonicity, as shown in Figure 5. Finally, we can compute a numerical value for monotonicity,

$$\text{Monotonicity}(x_i) = \frac{1}{m}\sum_{j=1}^{m} \frac{\left|\text{number of positive diff}\left(x_i^j\right) - \text{number of negative diff}\left(x_i^j\right)\right|}{n-1} \tag{1}$$

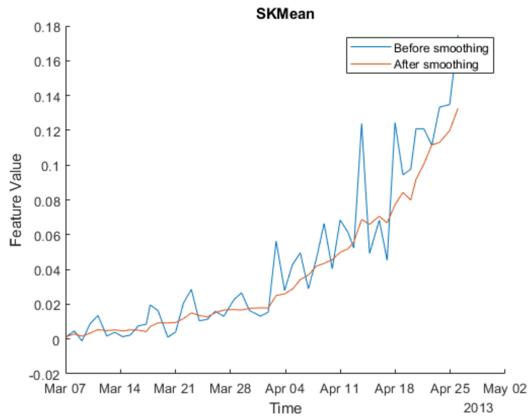

Figure 4 – Comparison of raw and smoothed data

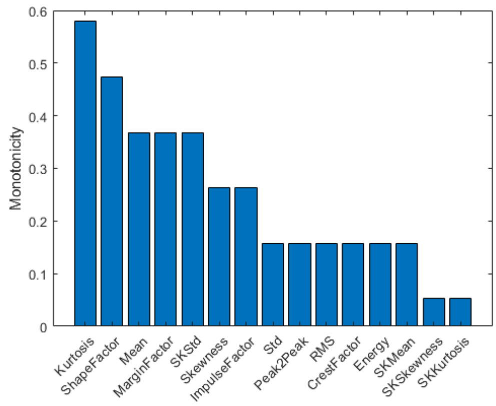

Figure 5 – monotonicity estimates of the engineered features.

This monotonicity value can then be used to filter out signals with low monotonicity. We set the threshold at 0.3, which reduced the data set from 15 signals to five. Further, we can fuse these final five signals down using principal component analysis, two principal components, as shown in Figure 6. This indicates that the first two principal components have a clear trend as their magnitudes increase through the working life, with a generally monotonic trend in the final 25 days to failure.

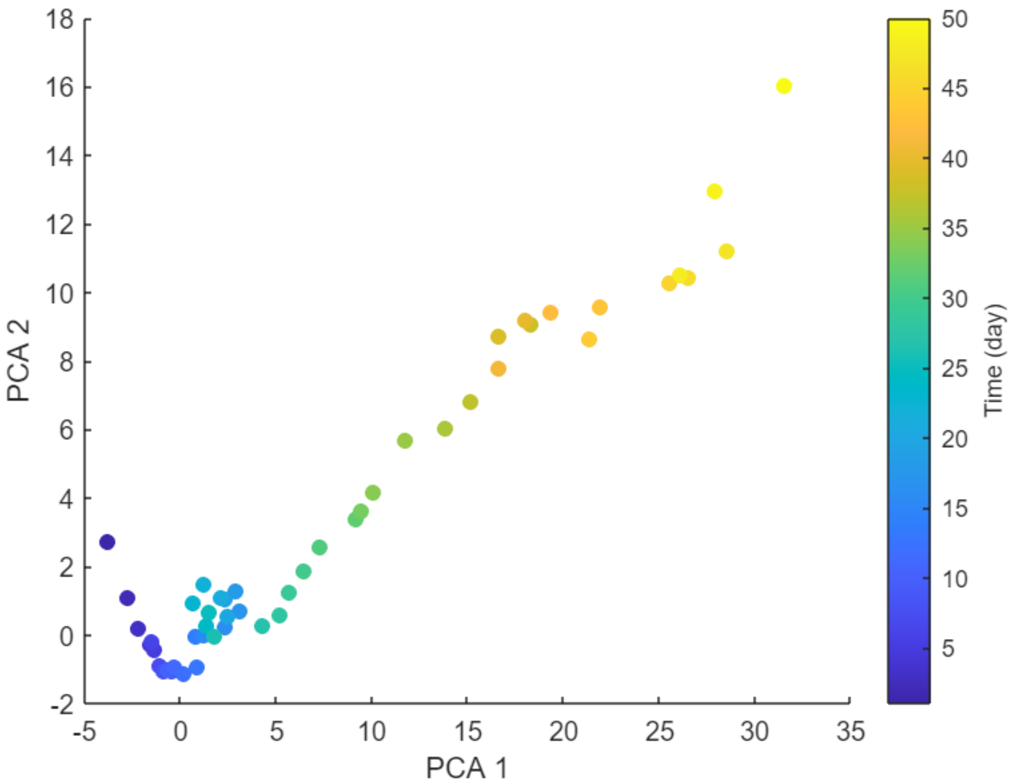

Figure 6 – scatter of the first and second principal components coloured by day of measurement

Model Fitting

The final step here is to fit an exponential degradation model,

$$h(t) = \phi + \theta e^{\beta t + \varepsilon - \frac{\sigma^2}{2}} \tag{2}$$

where *h(t)* is the health indicator as a function of time, ϕ is an intercept term and generally considered to be constant, θ and β are semi-random parameters that determine the slope of the model, while ε is a Gaussian white noise yielding to $N(0, \sigma^2)$. The final term, $-\frac{\sigma^2}{2}$ is to satisfy the expectation of *h(t)*

$$E[h(t)|\theta, \beta] = \phi + \theta e^{\beta t}. \tag{3}$$

Procedurally, we shift the health indicator such that the first value is zero and then use a standard fitting technique to fit the model. New daily data is fed iteratively to evaluate the model, with the model "remembering" its state from day to day. For example, the model state and life predictions for day 32 are shown in Figure 7.

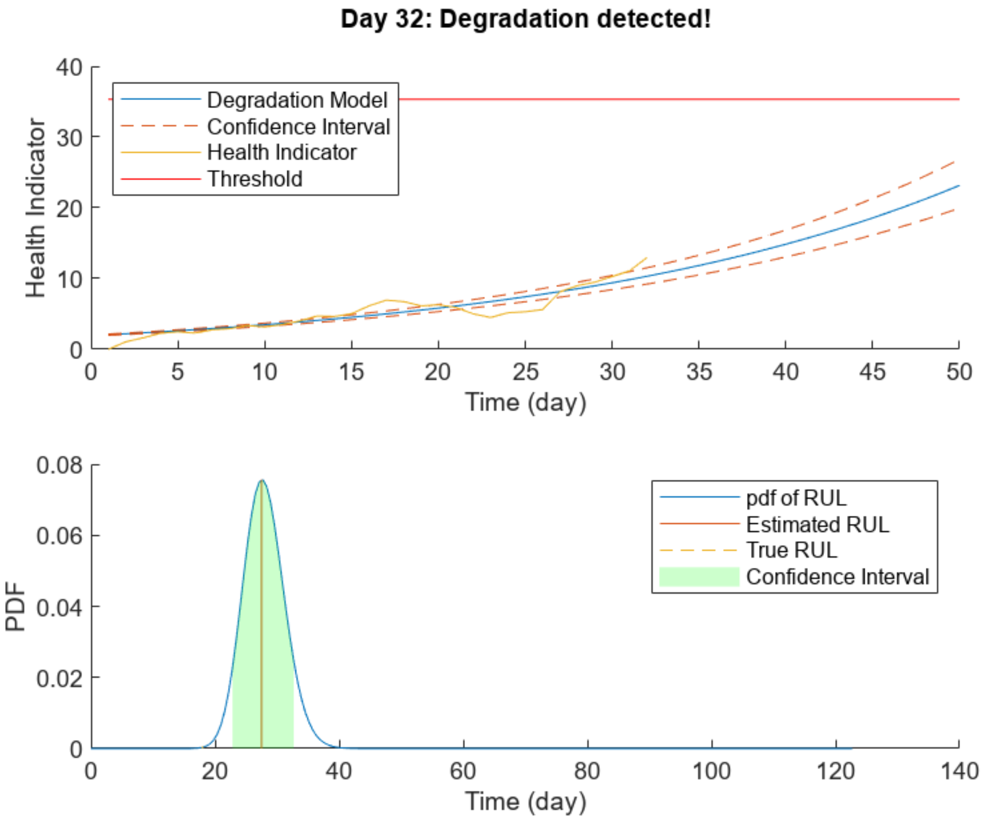

Figure 7 – example evaluation of remaining useful life model at day 32

Model Performance Evaluation

A common performance method check for a remaining useful life model is an α-λ plot, shown in Figure 8, in which α is set to 20% where the probability that the estimated remaining useful life is between the α bound of the true remaining useful life. In this case, the match begins poorly, as would be expected, where there is relatively little known of the fault or for the model. We assumed a linear degradation of RUL, which, based on an inspection of the spectral kurtosis, may not be correct. This can only be improved by dismantling the gearbox and measuring the bearing failure profile. However, as more data points are added to the prediction, the model tends to the true value, so despite the previous uncertainty, we have confidence in the final lifetime predictions – as we tend to fail, we have confidence.

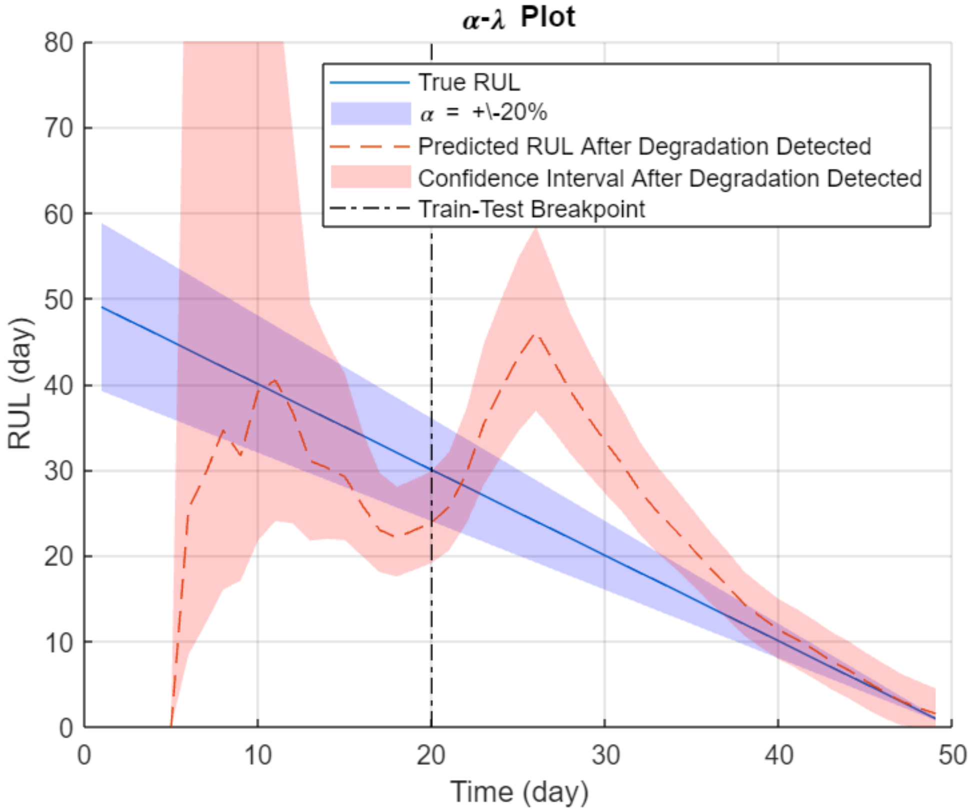

Figure 8 - α-λ plot of model performance

Conclusion

This paper has presented a deep dive into a sample predictive maintenance workflow based on a set of wind turbine data. We have walked through the initial ingestion of the data, exploration, and feature engineering through to fitting a model based on reduced and fused data sets before evaluating the model performance via an α-λ plot. As discussed in the introduction, the final step not discussed here is the deployment of the model into operations and handling of model drift.

References

[1] 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.

[2] Saidi, Lotfi, et al. *Wind turbine high-speed shaft bearings health prognosis through a spectral Kurtosis-derived indices and SVR.* Applied Acoustics 120 (2017): 1-8.

[3] Saxena, Abhinav, et al. *Metrics for offline evaluation of prognostic performance.* International Journal of Prognostics and Health Management 1.1 (2010): 4-23.