CONDITION MONITORING INSIGHT USING BAYESIAN INFERENCE AND ROTOR DYNAMICS MODELLING FOR ROTATING MACHINERY

Greg Nelson
Frazer-Nash Consultancy, a KBR company
United Kingdom

Iain Palmer
Frazer-Nash Consultancy, a KBR company
United Kingdom

Originally published in the Proceedings of ASME Turbo Expo 2022 Turbomachinery Technical Conference and Exposition

ABSTRACT

Rotor dynamics modelling can be used to predict vibration levels for given inputs, such as unbalance levels and location, which may be of interest for condition monitoring or diagnosis. However, given measured vibration, using rotor dynamics models to find the corresponding root cause inputs is not straightforward.

In the method presented in this paper, Gaussian Process models are developed as surrogates for the rotor dynamics finite element models, and are used with Bayesian Inference to determine the probability distributions of model inputs for a given vibration response. This method allows parameters describing the machine condition, such as unbalance location and magnitude, and bearing clearances, to be determined as well as the confidence in these predictions.

The method is demonstrated by simulating the vibration response of a compressor rotor, adding noise to it, and then using the technique to accurately infer useful information such as the unbalance magnitude and location, and the clearance in each bearing.

This technique can be applied as a risk-based approach to condition monitoring of rotating machinery. Further development of this approach as part of a digital twin which uses in-service measurements would provide operators with insight into the likelihood of different root causes of vibration, and the corresponding machine condition.

1 NOMENCLATURE

1X	Vibration response filtered to a multiple of 1x the running speed
x_i	The ith parameter point at which an amplitude response is recorded
$d(x_i, x_j)$	The Euclidean distance between two parameter points x_i and x_j
l	Kernel length scale for the Gaussian Process
σ	Kernel scale factor for the Gaussian Process
k	Kernel function used in the Gaussian Process
c	Vector describing the machine condition – the unbalance magnitude and location, and journal bearing clearance parameters to be estimated
M	Vector of measured amplitudes for varying speed
n	The number of measurements taken, which is twice the number of discrete speed points at which measurements are taken and GP predictions are made in the run down, as there are two proximity probes
r	The vector of residuals – the error between the measurement amplitude predicted by the GP, and the true measured amplitude
Σ_M	The $n \times n$ covariance matrix of assumed measurement noise
GP	Gaussian Process
MCMC	Markov Chain Monte Carlo

2 INTRODUCTION

2.1 Physics-based methods for rotating machinery

Rotor dynamics models are commonly used in the rotating machinery sector to predict vibration levels of rotating machinery [1], ensure they operate away from critical speeds, and confirm they do not suffer from other related issues such as instability. These models are often built using simple finite element beam representations of the rotor geometry, with spring-dampers representing the bearings. They are typically run with given out-of-balance sets across the speed range of the machine, to predict the vibration response as a function of speed. They are generally built and used either at the design stage of the machinery, to ensure that the machine operates free from vibration problems, or as a tool to help diagnose issues where this has not been adequately carried out at the design stage. In this instance, the models are often verified by manually cross-checking the vibration response with the measured data.

However, their use for interpreting measured vibration responses is limited, because they take inputs which correspond to the conditions of interest (such as unbalance magnitude or location, bearing clearance etc) and use these to predict outputs which correspond to the measured vibration. Therefore, to fully make use of them in combination with the measured vibration data, we are required to solve the inverse problem – that is, which combinations of model inputs best explains the measured data? This is a problem which also manifests itself in many other areas of rotating machinery, for example when interpreting blade tip timing results.

Therefore, despite their utility in predicting vibration response for a given machine condition, the use of such models in a condition monitoring context, for example as part of a 'digital twin', is currently limited. Physics-based methods suitable for on-line condition monitoring of rotating machinery, at least for interpreting vibration data, are therefore not currently widely used. This paper addresses the inverse problem applied to vibration diagnostics.

2.2 Current condition monitoring methods for rotating machinery

Vibration measurements, whether taken using proximity probes or accelerometers, are commonly used to give an indication of the condition of rotating machinery [2]. However, automatically determining the root cause of increased vibration is a more difficult proposition, and interpreting this in terms of risk is harder still. Two main approaches are used to interpret vibration data – one which uses limits and trending [3], and the other Artificial-Intelligence (AI) based [2].

2.2.1 Limit-based methods

Trending or limit-based approaches to condition monitoring assign limits to vibration levels [3], and look at trends through time. This technique is commonly used in industry, with alarm limits set on the vibration recorded in machines. However, they do not always use or provide information about the root cause – such as whether the increase in vibration is caused by unbalance (where this unbalance is located and what may be the cause), or by bearing issues such as increased clearance due to wear. They also do not provide information about the risks that are brought about by the change in machine condition, such as wear or instability.

2.2.2 Artificial-Intelligence methods

Another approach in increasing use is AI-based approaches, which use machine learning techniques such as neural networks, as in references [2], [4] and [5] to spot and diagnose potential faults based on the vibration data. This is the current state of the art for automatically interpreting condition monitoring data for rotating machinery.

Some examples of these approaches include Principal Component Analysis, Support Vector Machines and Artificial Neural Networks, as in references [4] and [6]. These methods are useful techniques for classifying faults into different categories, thereby diagnosing the potential issues behind the data received. However, these machine learning techniques suffer from two main limitations:

1. They often rely on extensive training data. This data may not always be available – for example on rare/unique machines, new designs, and/or machines with high consequence of failure, it may not be possible to see or generate all of the training data required to make accurate predictions. Indeed, it may be impossible or unsafe to operate in some regimes which the condition monitoring system is required to detect and diagnose.

2. They have no built-in knowledge of the physics in the system. This has two further implications:

   a. It is difficult to understand how reliable the predictions made by the system are, particularly if the machine is operating in a regime not seen in the training data, as may be the case during a fault
   b. Significant training time and data is required to 'learn' physics that is in fact well understood, leading to a lost opportunity to understand more fully what is being seen.

2.2.3 Other methods for detecting out of balance

A number of alternative techniques have been developed in an effort to detect the magnitude and location of unbalance in rotating machinery. Reference [7] describes a method to estimate unbalance and misalignment of a machine from a single run down, by constructing matrices that represent the rotor-bearing system including its misalignment and solving the resulting equations to minimise the least-squares error across a range of frequencies. This provides a point estimate of the parameters of interest – namely the unbalance location, magnitude and rotor misalignment.

Reference [8] provides an alternative method which works at steady state, by estimating the location and magnitude of unbalance and rotor misalignment based on a residual force generation technique. Reference [9] gives a recent method to estimate unbalance location and magnitude which eliminates the use of modal expansion and is not based on calculating residual forces. It uses a Finite Element simulation model to generate a set of reference cases, and compares the predicted response amplitudes (displacements) at the measurement locations with the measured values to estimate first the unbalance location and then the magnitude.

These methods provide a variety of means of deterministically estimating the unbalance magnitude and location, typically using a least-squares minimisation approach. However, being point estimates, they do not offer the advantages of a probabilistic approach in directly estimating the confidence in their predictions.

2.2.4 Stochastic methods

Stochastic methods are not often applied in the field of rotor dynamics, but there are some recent developments. In 2015, reference [10] used Bayesian inference to identify unbalance magnitude and location, and reference [11] also used Bayesian inference to estimate journal bearing parameter uncertainty and proposed its use as part of a design process. Markov Chain Monte Carlo (MCMC) was used to perform the Bayesian Inference. MCMC returns a posterior probability distribution for each parameter of interest, rather than a single value; this is an advantage, as this carries information about the degree of uncertainty in each parameter being estimated. This is useful in a condition monitoring context, where the operator must make fundamentally risk-based decisions about how and whether to operate the machinery – and therefore an estimate of the uncertainty on a given parameter of interest and/or corresponding risk will be useful.

More recently, reference [12] has used Polynomial Chaos Expansion (PCE) methods to perform the Bayesian inference step to identify journal bearing wear and unbalance parameters. A similar method is presented in this paper, with a key difference being the use of Gaussian Processes in a MCMC to perform the Bayesian inference, rather than the PCE methods as in reference [12] or directly on the Finite Element model as in references [10] and [11]. This gives the opportunity to provide additional information about the machine condition, by running the complete Markov chain of identified parameters back through a Gaussian process to obtain further information about the machine condition, such as vibration amplitudes at locations other than at the measurement locations.

3 A PHYSICS-BASED METHOD FOR CONDITION MONITORING OF ROTATING MACHINERY

3.1 Method overview

A method is presented to address the problem of interpreting measured vibration data from rotating machinery, by combining known techniques from rotor dynamic analysis and statistics, and which is made computationally feasible with the use of Gaussian process surrogate models. The method makes use of the physics inherent in a rotor dynamics model to diagnose the root cause of the observed vibration, and the corresponding risks. The method combines the physical understanding built into standard rotor dynamics finite element models with Bayesian inference, a statistical technique to relate the observed outputs to the most likely combination of inputs that caused it. The method is shown in the diagram in Figure 1, and was implemented using the Python programming language.

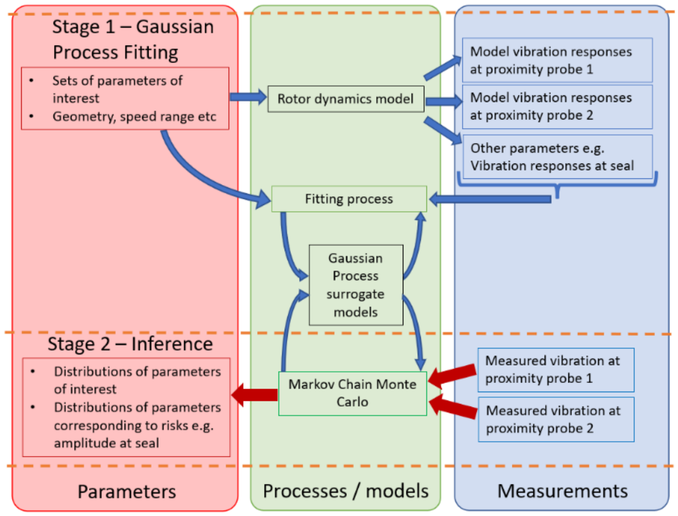

FIGURE 1: DIAGRAM OF METHOD

As shown in Figure 1, the method proceeds in two stages – Stage 1, in which the Gaussian processes are fitted to data generated by the FE models, and Stage 2, where the parameters of interest are inferred. The right-hand side of the figure corresponds to parameters which are outputs from a rotor dynamics model, and which generally correspond to parameters which may be measured, such as the 1X vibration amplitude at proximity probe locations. The left-hand side corresponds to parameters which may not always be easily measurable, but are of interest for condition monitoring, and often are inputs to a rotor dynamics model. Examples include unbalance magnitude and location.

In Stage 1, the parameters to be inferred are defined, along with the other parameters required to run the rotor dynamics model, such as the geometry, speed range etc. The model is then built and run for a range of discrete points across combinations of all the parameters of interest. This is done in an automated manner, and the displacement amplitudes at the proximity probe locations are recorded for each parameter combination as a function of speed. In addition, if other outputs from the models are of interest, such as the corresponding vibration amplitudes at a seal or other region of close clearance, these are also stored. Gaussian process surrogate models are then fitted to provide a mapping between the input parameters and corresponding recorded output parameters from the rotor dynamics models.

The Gaussian processes are then used in the Bayesian inference step for Stage 2. Gaussian processes are used as they are much less computationally expensive to evaluate than the FE model. Bayesian inference solves the inverse problem, allowing the input parameters which best correspond to the measured vibration responses at the proximity probe locations to be estimated.

Bayesian inference is carried out using a MCMC method to infer the posterior distributions of the model inputs of interest, in a similar manner to reference [10], but acting on the Gaussian Processes rather than the FE model.

The following sections outline the key elements of this method.

3.2 Rotor dynamics model

The purpose of the rotor dynamics model is to provide a mapping between possible machine conditions and their corresponding observed vibration response as a function of speed. This model provides the physical understanding inherent in the method, giving it accuracy and robustness.

The model itself consists of a finite element model of the rotor-bearing system. Typically, this would consist of a Timoshenko beam element model, with mass and inertia elements, gyroscopic effects included, and spring-damper elements representing bearings and seals. The bearing elements should include direct and cross-coupled terms as required, with the coefficients being speed-dependent and calculated as appropriate for the bearing or seal in question. This could be using empirical correlations, bespoke software or simple fluid dynamics models. In the case reported here, the coefficients were calculated using bespoke software, and the analysis was run using an ANSYS solver.

The model is solved as for the forced response across a speed range for multiple configurations of unbalance magnitude and location, bearing clearance and other system parameters as required. The vibration response at the proximity probe locations (and/or other measurement points) is then recorded. In addition, other model outputs of interest are also recorded, such as the vibration amplitude at locations where rubs or wear may be an issue, or a stability value (from a modal analysis).

The rotor dynamics model is run in this way to simulate a run up or run down for different combinations of parameters of interest, in a design-of-experiments approach, for a tractable number of runs. For the case reported in Section 4, since the vibration response scales linearly with the applied unbalance magnitude, the corresponding points for varying unbalance magnitude were not all generated using the rotor dynamics model, but by scaling the response linearly.

3.3 Gaussian Process surrogate model

3.3.1 Background to Gaussian Process Surrogates

The purpose of the Gaussian process surrogate models is to enable rapid predictions of vibration response for a given set of input parameters, orders of magnitude faster than is possible by evaluating the rotor dynamics model directly. This is because the Bayesian inference method requires a large number of simulation runs, and likely would not be possible to use with the full rotor dynamics model for realistic geometry seen in industry, and in time-scales relevant to condition monitoring.

Gaussian process surrogate models [13] are a form of surrogate model which assume that the outputs at neighbouring points are similar and vary with a covariance. By aggregating these points and supplying points with known output and output uncertainty (which can tend to zero), a mean line can be drawn passing through each point supplied in the fitting process (or near in the case of non-zero uncertainty in the fitting data). In this way, both the mean line and the variance can be computed at all points in the input parameter space, with greater certainty (lower variance) nearer to the fitting points.

Gaussian process surrogate models are a good choice of surrogate model for a variety of reasons. They are very flexible for a wide range of response types, they generalise readily to multiple input dimensions (they remain fairly computationally efficient up to 10 or 20 dimensions, which would likely be an upper limit on the number of parameters able to be estimated from a few vibration sensors), and their statistical nature is appropriate in the probabilistic approach employed in the Bayesian inference. This is because when evaluated at a given point, they can return both a mean and a variance, and this degree of uncertainty can readily be combined with the assumed measurement uncertainty in the Bayesian inference step.

3.3.2 Implementation of Gaussian process models

The Gaussian processes used here each map all of the input variables of the rotor dynamics model, including rotor speed, to one model output – for example, the magnitude of vibration measured at a proximity probe. They are trained using the data produced by the rotor dynamics model runs, and Figure 2 shows an example output from the Gaussian processes fitted to the rotor dynamics model training data, when one parameter is allowed to vary and others are kept fixed.

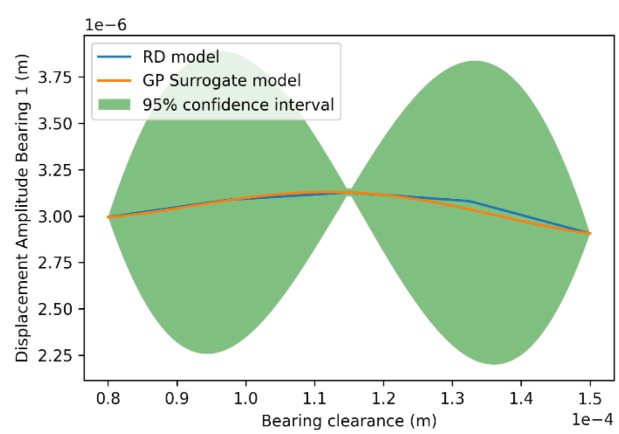

FIGURE 2: EXAMPLE GAUSSIAN PROCESS SURROGATE MODEL PREDICTIONS FOR VARYING BEARING CLEARANCE

The scikit-learn [14] implementation of a Gaussian process in Python was used for this work, with a Radial Basis Function kernel (covariance) function. This kernel is shown in Equation (1), where $d(x_i, x_j)$ is the Euclidean distance between two points in question, l is the length scale of the kernel, and σ is a scale factor. The length scale and scale factors are optimized as part of the fitting process of the Gaussian process.

$$k(x_i, x_j) = \sigma^2 \, exp\left(-\frac{d(x_i,x_j)^2}{2l^2}\right) \tag{1}$$

Using this kernel and the training data points run, the mean (expected) value at any point (including points not run to date) can be evaluated as outlined in Reference [13] and implemented in Reference [14].

In our implementation of the overall method, the input parameters were normalized to be between zero and one for reasonable magnitudes of the parameters, due to the vastly different scales of input parameters – for example, RPM varies up to 20,000 rpm, whereas the bearing clearance can be down to 0.00008 m. These different scales can make fitting and optimization difficult, with the potential for numerical errors becoming significant for some parameters. Normalizing the parameters removes this source of potential error.

3.4 Bayesian Inference method

3.4.1 Background to MCMC for Bayesian Inference

As discussed above, in this condition monitoring application for rotating machinery we are attempting to solve an inverse problem, that is, which combination of model inputs best describes the observed output, to better understand the machine condition. There are several ways this could be achieved, including posing the problem as an optimization problem and attempting to minimize the difference between the observed measurements and the measurements ‘predicted’ by the surrogate models for a given set of input parameters. This is discussed in reference [12].

However, there are advantages to posing the problem in probabilistic terms – i.e., given the observed measurements, what is the probability distribution of each of the model inputs which correspond to the machine condition? This gives the user an explicit understanding of the uncertainty around each of the parameter predictions, allowing better risk management decisions to be made, and allows for cases where different scenarios could explain the observed behavior.

$$P(c \mid M) = \frac{P(M \mid c)\, P(c)}{P(M)} \tag{2}$$

Bayesian inference is a probabilistic approach to solving the inverse problem, and Equation (2) shows Bayes formula applied to condition monitoring, where c is the machine condition and M are the measurements. It shows how to calculate the probability distributions of parameters representing the machine condition, given the observed measurement. The machine condition represents the parameters to be estimated, such as unbalance magnitude and location etc., and the measurements are the vibration amplitudes across a range of speeds, for example taken during a run-down.

The formula shows that the machine condition is a function of the probability of observing the given measurements given the machine condition parameters, as well as the prior distributions on the condition parameters and on the measurements.

However, the equation is difficult to evaluate directly, as the normalizing factor $P(M)$ is difficult to quantify. The Metropolis-Hastings algorithm, a MCMC method [15], allows the posterior distribution $P(c \mid M)$ to be approximated without knowledge of this normalizing factor. The method works by taking a random walk through parameter space by proposing random steps outwards from the current point, and calculating $P(M \mid c)\, P(c)$, which is the likelihood function multiplied by the prior on the machine condition. The step is then accepted or rejected randomly with a probability proportional to this calculated value if the step moves to a less likely point, or accepted if it moves to a more likely point. The distribution of steps taken will then tend to the posterior distribution to be calculated.

3.4.2 MCMC Implementation details

Due to the small probabilities encountered in the MCMC method, the probabilities were handled in log form, in order to minimize the effects of numerical error. Therefore, the top line of Equation (2) becomes the log likelihood added to the log of the prior on machine condition. Since the MCMC method works without knowledge of the normalizing factor $P(M)$, it is only required to use a function proportional to the gaussian likelihood function. A Gaussian likelihood function was used [16], as shown in Equation (3),

$$\log P(M \mid c) \propto -\frac{1}{2} r^T {\Sigma_M}^{-1} r \tag{3}$$

$$r = GP(c) - M \tag{4}$$

where r is the vector of residuals – the error between the predicted measurement values (in our case the amplitudes predicted by the models at each speed point) for a given condition vector, and the true measured values – and is of length n, where n is the number of measurements. Since we measure the amplitude at two proximity probes, n is equal to two times the number of discrete speed points taken in the run down. Σ_M is the $n \times n$ covariance matrix. In this paper, Σ_M was taken to be a constant assumed measurement noise multiplied by the Identity matrix. This was chosen as there is no expected variation of noise with speed/frequency, either from the proximity probes, or from other forcing (aerodynamic or otherwise) in the machine. If however, it was known for a given machine that certain frequencies contained more noise, this could be incorporated in the Σ_M matrix.

Uniform prior distributions were used throughout this work, although user experience could be incorporated into the prior distribution – for example, if it is known that a particular component, such as a disc, is more likely to become unbalanced, then this could be captured in the prior distribution.

In our implementation, an adaptive MCMC method [15] was used, updating the proposal covariance matrix based on the existing knowledge of the posterior distribution from the steps taken so far. Python was used to implement this method. The precise details of the MCMC steps to accept or reject moves through the parameter space are well covered in references [15] and [16] and are not discussed here.

The output from the MCMC is a chain of steps through the parameter space, the distribution of which tends to the posterior distribution of $P(c \mid M)$, which is the distribution of the machine condition, given the observed measurements. In addition, now that we have a set of parameters distributed according to their posterior distribution, we can run each case back through additional GPs, which have been fitted to give further information such as the vibration amplitude at other locations of interest. The output from these GP(s) will therefore be the corresponding distributions of amplitude (or other parameters of interest) and be used to calculate the risk of wear occurring at locations of close clearance, for example. This additional step is enabled by the MCMC chain and the use of GPs, and represents a key difference to the work presented in reference [12].

4 Application of Method to High Speed Compressor

4.1 Context and Methodology

This section describes how the method outlined in Section 3 was implemented for a simulated six stage high speed compressor, and the corresponding results of the study. The compressor geometry was chosen to emulate real-world geometry and speeds, with the rotor operating above its first critical speed.

The goal of the case study is to estimate the bearing clearance at each bearing, and the unbalance magnitude and location, which may be located at one of the compressor wheels, and potentially due to a blockage in one of the channels. These are chosen as the parameters of interest to be inferred from the measured vibration response, as well as the risk of wear at the seal location as discussed above. We also intend to estimate the amplitude of vibration at the seal location, a location of close clearance, in order to understand the risk of wear occurring for the given machine condition.

In the absence of the means to induce a known out of balance on the physical machine, simulated measurement data was produced by adding noise and error to results from an FE model.

There are a number of potential sources of error between simulated and measured vibration responses. Firstly, there is the potential for the model to be incorrectly calibrated, which is something that the method presented in this paper is able to address, by setting the unknown parameter as a variable to be estimated. Second, there is noise in the measurement process and other uncertain or unknown forcing occurring in the system. This is represented by noise in the measurement data, and was simulated by adding randomly distributed noise to results produced by the FE model. Thirdly, there is systematic error, where the model is not a fully accurate description of the system – for example, the model may not have the correct rotor stiffness, or may not account for the flexibility in the bearing supports. This third case was also simulated, by adding noise to results generated from a separate FE model with small changes to its stiffness and oil viscosity, so that the model used to produce the measured data and the model used for inference are different. This is discussed further in Section 4.4.

The method described in Section 3 was then applied, to infer the unbalance magnitude and location, the bearing clearances and the seal vibration amplitude. Section 0 describes the underlying geometry and FE model used in this example, Section 4.3 shows the Gaussian Process surrogate models fitted, Section 4.4 shows the simulated measurement data produced from which to carry out the inference, and Section 4.5 shows the results of the study.

4.2 Model set up

Figure 3 shows the rotor model used in this paper. It consists of a high speed six stage compressor, mounted on short journal bearings. The rotor is modelled with Timoshenko beam elements, and mass/inertia elements are used to represent the compressor wheels, as these do not significantly contribute to the rotor stiffness. At step changes in the rotor diameter, the stiffness radius of the rotor is reduced, with material outboard of this radius contributing mass and inertia but not stiffness, as would be the case in a 3D rotor model. Proximity probes are located near to both bearings.

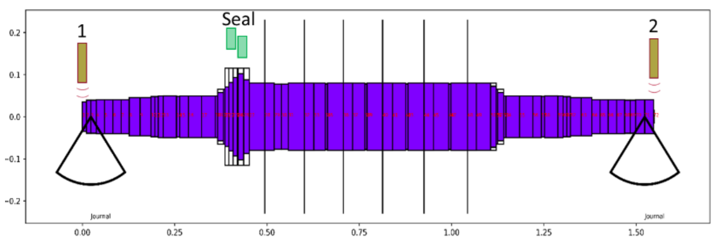

FIGURE 3: ROTOR MODEL IN ROTOR DYNAMICS TOOL, WITH PROXIMITY PROBE AND SEAL LOCATIONS SHOWN

The model was built and solved using an ANSYS solver. The journal bearing characteristics as a function of speed were calculated using the method presented in reference [1].

In order to fit the GPs, the model was run for ranges of four parameters of interest: unbalance magnitude, unbalance location, and bearing clearances at each bearing (which affect the stiffness and damping terms of the bearings). The ranges were generated using a uniform sampling methodology, with the exception of the speed term, where more samples were concentrated at speeds near the peaks in response. The models were run as forced response models through the running range of the machine – up to 20,000 rpm. Results were recorded for all combinations of six unbalance locations, three bearing clearance values at bearing one and three at bearing two and two unbalance amounts, each for a range of 100 speed values. This resulted in 108 parameter combinations run, each for 100 speeds. The resulting vibration amplitude was recorded at the proximity probe locations as functions of speed. The vibration amplitude was also recorded at a location corresponding to a seal with known clearance. This leads to a total of 10,800 points per GP surrogate model.

Figure 4 shows an example plot from the rotor dynamics model, for a 0.002 kg-m unbalance at located 0.602 m, with 0.115 mm bearing clearance at both sides. The figure shows a clear mode just below 5,000 rpm, with a split critical due to bearing anisotropy, and further modes at 10,000 rpm and 15,000 rpm. It also shows significantly higher vibration at some speeds for the seal than is seen at the measurement locations, which may not be apparent from the proximity probe measurements but would be useful for an operator to understand.

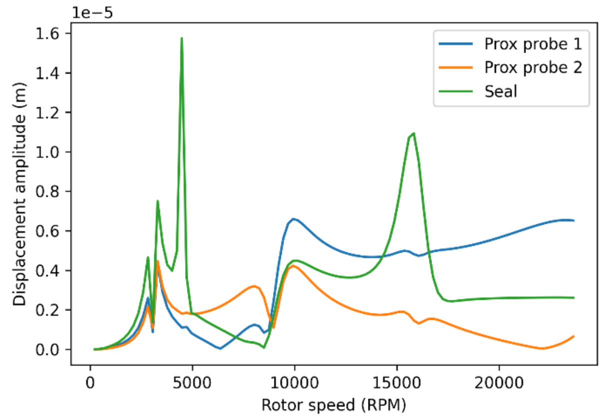

FIGURE 4: VIBRATION FOR 0.002 KG-M UNBALANCE AT 0.602M, 0.115MM BEARING CLEARANCE AT BOTH SIDES

4.3 Gaussian process surrogate models

Gaussian process surrogate models were fitted using the data generated by the rotor dynamics model. Three surrogate models were generated – two for the vibration amplitude responses at the two proximity probe locations, and a further surrogate model for the response at the seal location. Each model takes inputs corresponding to the rotor speed, unbalance magnitude, unbalance location and the clearance at each journal bearing (as two separate parameters). The rotor speed is assumed to be known, as each vibration measurement will be taken at known rotor speeds. The other parameters represent the machine condition, which we wish to infer based on the observed vibration response.

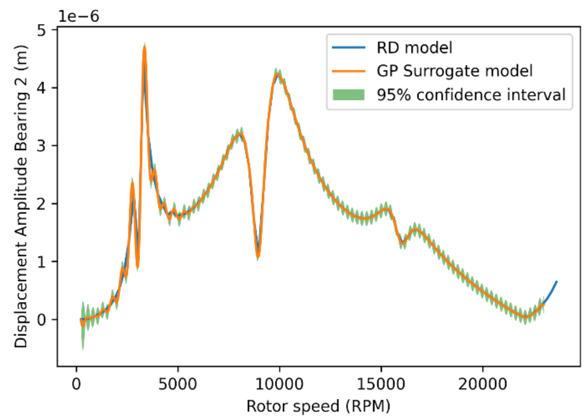

FIGURE 5: GAUSSIAN PROCESS MODEL COMPARED WITH ROTOR DYNAMICS MODEL FOR 0.002 KG-M UNBALANCE AT 0.602M LOCATION, WITH CLEARANCES FOR 0.115 MM AT EACH BEARING

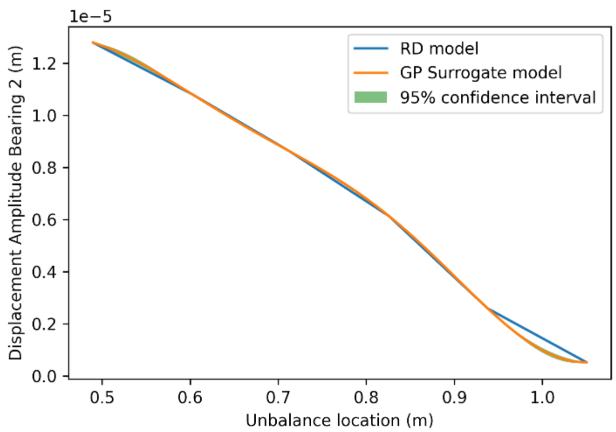

FIGURE 6: SURROGATE MODEL FIT FOR 0.002 KG-M UNBALANCE, 5198 RPM AND 0.115 MM CLEARANCES

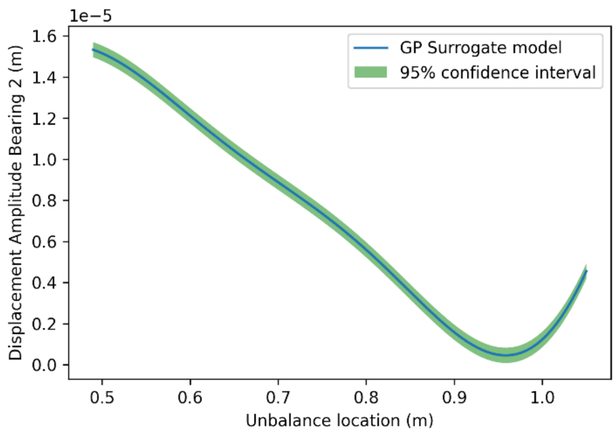

FIGURE 7: SURROGATE MODEL FIT FOR 0.002 KG-M UNBALANCE, 6050 RPM AND 0.115 MM CLEARANCES

Figure 5 to Figure 7 show plots from the surrogate model responses at proximity probe two, to illustrate the quality of fit of the surrogate model. The models show a good fit to the rotor dynamics model data. Figure 7 shows the predicted amplitude as a function of unbalance location for a speed which was not run in the rotor dynamics model, which therefore shows an increased level of uncertainty in its prediction.

4.4 Simulated Measurement

In order to simulate vibration responses from the rotor taken in the field, noise was added to a set of forced responses generated by the rotor dynamics model. This captures the uncertainty created by the noise in the measurement process or other noisy forcing in the machine. The simulated machine condition is shown in Table 1, and is not a parameter set that was used in the generation of the points to fit the GPs, requiring the GPs to correctly interpolate at this condition.

Two measurement cases were run corresponding to high and low noise cases. The high measurement noise case was made by adding normally distributed error to the rotor dynamics model output, with a standard deviation of $5.0\text{x}10^{-7}$ m, and the standard deviation for the lower noise case was 1.0x10-7 m. Figure 8 and Figure 9 show the 'measured' noisy vibration response which correspond to the simulated machine condition which is to be inferred, compared with the underlying data produced by the rotor dynamics model before the addition of noise.

Another low noise case was run, where the stiffness of the rotor in the model used to generate the simulated measured data was increased by 5%, and the oil viscosity increased by 10%. This is to represent a case where there is systematic error in the model – i.e. the model does not completely match the real machine due to some small error or unknown in the modelling. Figure 10 shows this case.

Table 1 SIMULATED MACHINE CONDITION

Parameter	Value
Unbalance location	0.602 m (Compressor wheel 2)
Unbalance magnitude	0.000333 kg-m
Bearing 1 clearance	0.0975 mm
Bearing 2 clearance	0.1325 mm

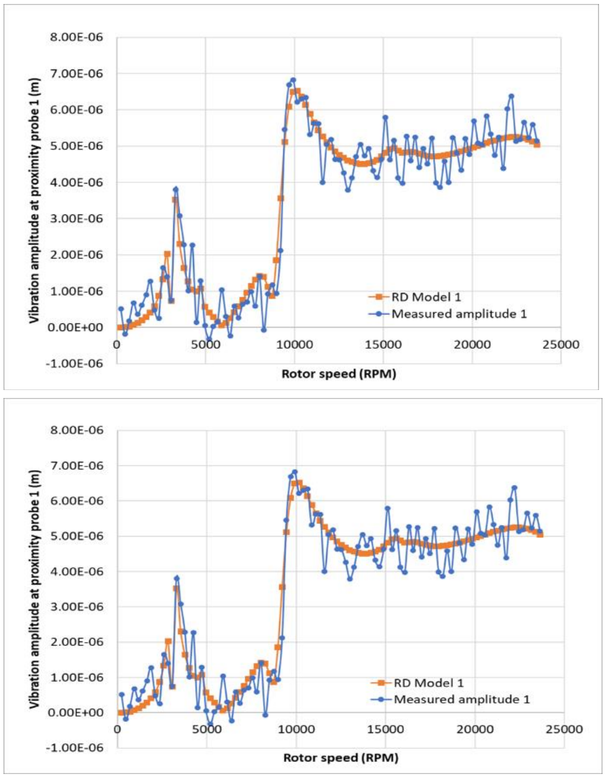

FIGURE 8: HIGH NOISE SIMULATED MEASUREMENT DATA COMPARED WITH UNDERLYING DATA FROM MODEL

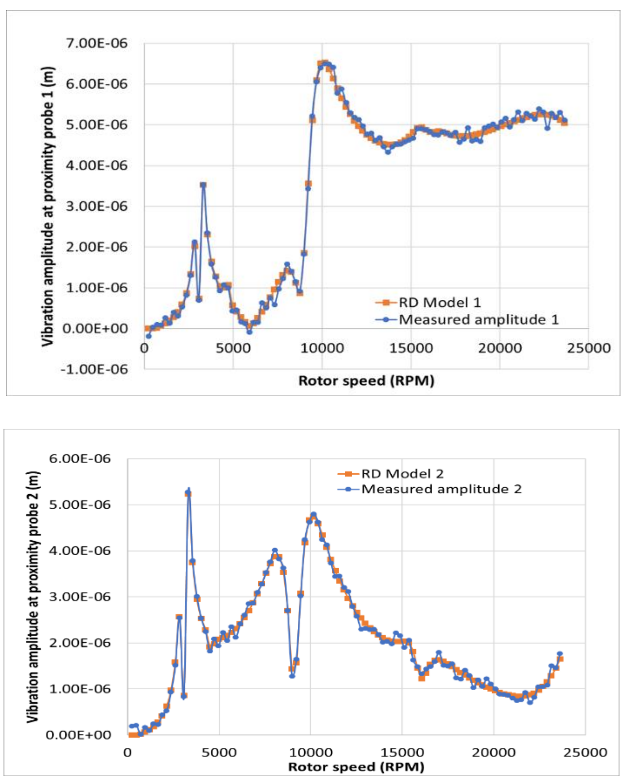

FIGURE 9: LOW NOISE SIMULATED MEASUREMENT DATA COMPARED WITH UNDERLYING DATA FROM MODEL

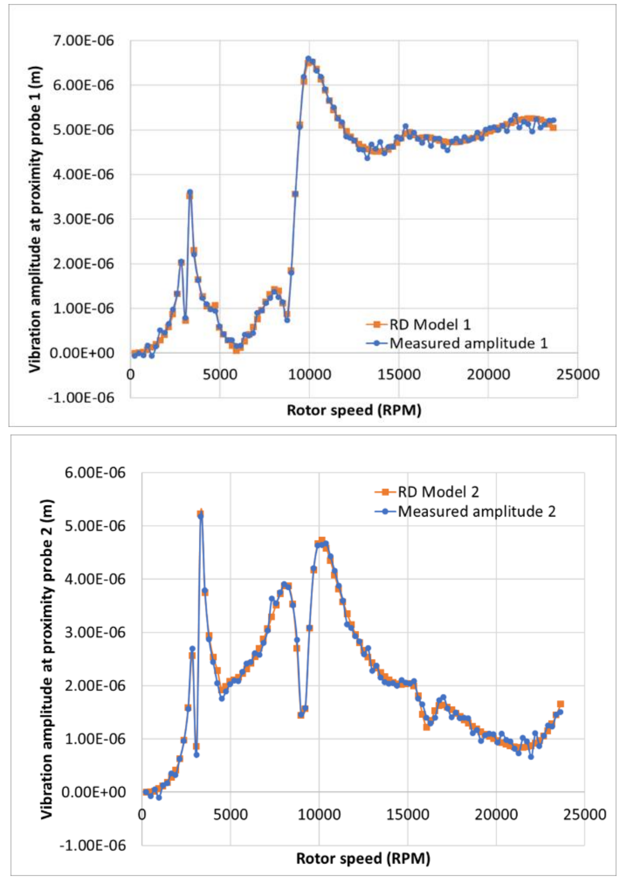

FIGURE 10: SIMULATED MEASUREMENT DATA COMPARED WITH UNDERLYING DATA FROM MODEL FOR CASE WITH SYSTEMATIC ERROR

4.5 Bayesian Inference

4.5.1 High noise

For all cases, the MCMC chain was run for a length of 50,000 samples, with the first 10% removed for the burn-in period. Figure 11 to Figure 15 show the results of the MCMC chain for the high noise case. Figure 11 to Figure 14 correspond well with model inputs, and are shown on with the x axis corresponding to the full range of potential values considered, with the exception of Figure 11, for which only the first half of the range is shown, for clarity. The true value is shown by a green vertical line.

The figures show that even in the presence of relatively high noise, good estimates of the parameters of interest can be obtained. The distributions of all uncertain parameters bound the true values, and true values of the bearing clearances and unbalance location are very near the centre of the distributions produced. This indicates that the method should work well in practice to infer the true values from measured in-service data.

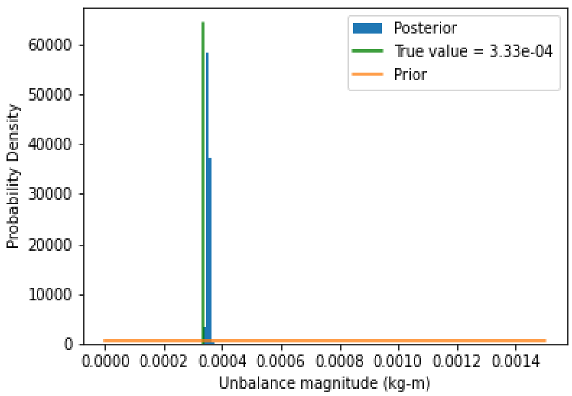

FIGURE 11: POSTERIOR DISTRIBUTON OF OUT OF BALANCE MAGNITUDE COMPARED WITH TRUE VALUE

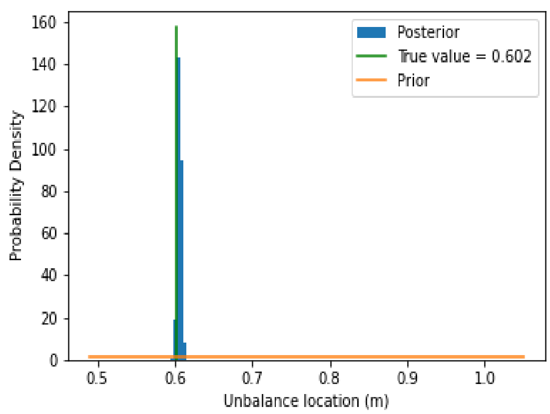

FIGURE 12: POSTERIOR DISTRIBUTON OF OUT OF BALANCE LOCATION COMPARED WITH TRUE VALUE

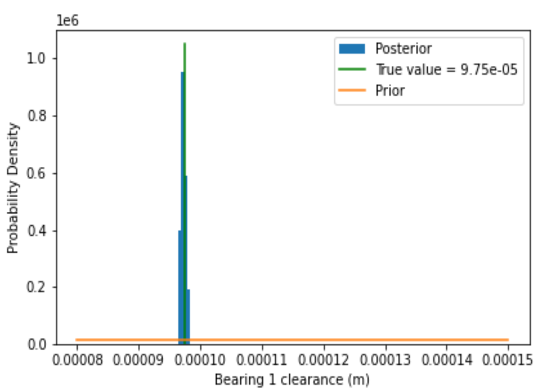

FIGURE 13: POSTERIOR DISTRIBUTION OF JOURNAL BEARING 1 CLEARANCE COMPARED WITH TRUE VALUE

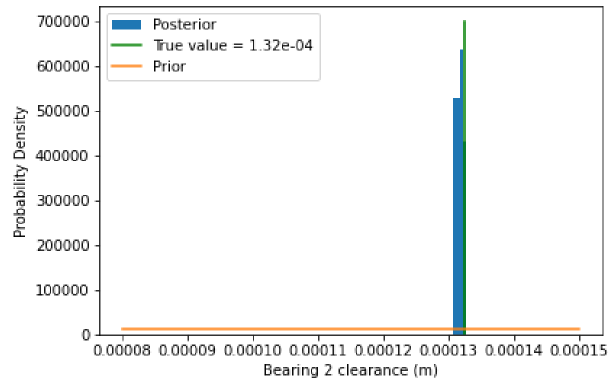

FIGURE 14: POSTERIOR DISTRIBUTION OF JOURNAL BEARING 2 CLEARANCE COMPARED WITH TRUE VALUE

Figure 15 shows the posterior distribution of the vibration amplitude at the seal location, which also bounds the true value. Just as the posterior distributions of input parameters are found by taking the distributions of MCMC sampled points, Figure 15 is produced by running these distributions of input parameters through the GP model for the seal displacement amplitude and plotting the corresponding distribution of seal displacement.

For this plot, although the distribution appears wide compared with the other figures, it is only due to the scale on the x axis, as there is no prior on this value, and all points are within 10% of the true value. The width of the distribution is governed ultimately by the widths of the distributions of input parameters and their corresponding effect on the amplitude at the seal clearance location.

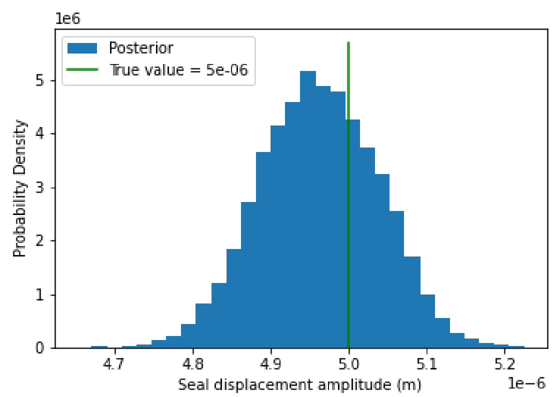

FIGURE 15: POSTERIOR DISTRIBUTION OF MAXIMUM VIBRATION AMPLITUDE AT SEAL LOCATION

4.5.2 Low noise

The results from the low noise case are similar to that reported in Section 4.5.1, with distributions that are well centred around the true values of the parameters of interest. The distributions themselves are slightly tighter compared to the high noise case, meaning that there is greater certainty in the predicted value.

This is to be expected, as the effects of noise are smaller. Figure 16 shows an example plot of the distribution of out of balance location predicted for the low noise case.

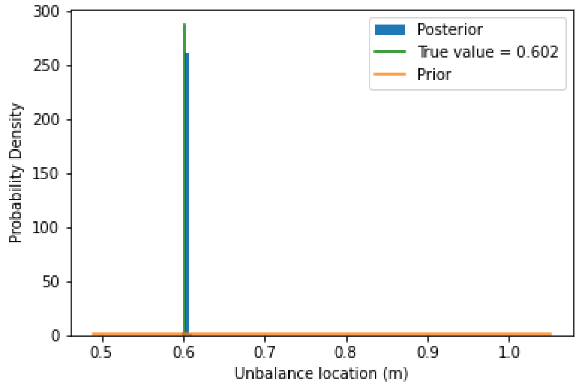

FIGURE 16: POSTERIOR DISTRIBUTON OF OUT OF BALANCE LOCATION COMPARED WITH TRUE VALUE

4.5.3 Systematic model error

The results from the systematic model error case were very similar to the low noise results, and Figure 17 shows an example results plot, showing the distribution of out of balance location predicted, which compares well with the true value. This shows that the method is robust even in the presence of small systematic errors in the models used – for example, where the model does not fully match the behaviour of the real machine.

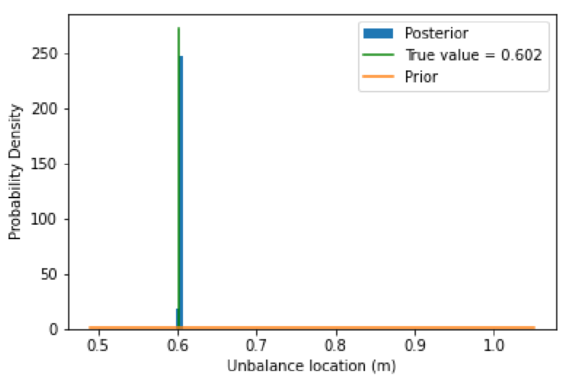

FIGURE 17: POSTERIOR DISTRIBUTION OF OUT OF BALANCE LOCATION COMPARED WITH TRUE VALUE

5 DISCUSSION AND EXPLOITATION

This paper outlines a method to infer useful information about the condition of rotating machinery from measured vibration data. This includes unbalance magnitude and location, and bearing clearances, and the method also gives other useful information about risks, such as the probabilities of rubs or wear at given locations. The method was shown to work well in the simulated example shown, accurately inferring the parameters of interest from noisy simulated vibration data, as well as in a case where the underlying models had some small systematic error, as would likely be the case when modelling real machinery.

This method provides benefits relative to conventional limit and trending based approaches to interpreting vibration data, as well as non-physics-based methods which rely purely on artificial intelligence or data analytics approaches. The benefits are given by incorporating knowledge of the underlying physics with the rotor dynamics models, allowing more insight to be gained from commonly measured data, and allowing more reliable interpretation of that data. It also allows information about the machine condition to be inferred in the absence of sensor data that can measure the condition more directly, perhaps because the parameter is difficult to measure, or sensors have not been built in to the equipment. This is achieved by using the physics to relate what has been measured to the influence of the non-measured parameter.

The method relies on rotor dynamics models, which can be hard to calibrate well with experimental data, due in part to uncertainty in some of the key inputs, such as bearing characteristics. However, where there is uncertainty in the values of these model input parameters, this approach can actually be used to mitigate the problem and help calibrate the model, by taking these inputs as parameters to be estimated rather than setting them as fixed values. This allows the method to be employed for calibrating rotor dynamics models in circumstances other than condition monitoring, such as for vibration problem solving and diagnosis.

The method is also probabilistic in nature, which provides operators with insight into the likelihood of different root causes of vibration. This provides a benefit to the user, who wants to know not just the machine condition but the uncertainty in that estimate, in order to make risk-based decisions.

The method could be practically implemented by carrying out Stage 1 in Figure 1 prior to use, and setting up the Gaussian Processes. Then, whenever the machine is run up or run down, the measured vibration response at the proximity probes would be filtered to the 1X response, and the MCMC step in Stage 2 would be run, which is computationally tractable to run each time the machine is switched on or off due to the use of GP surrogate models rather than the full FE models. This would then give the operator probability distributions on clearances in the machine, unbalance magnitudes or locations, and these results could be tracked each time the process is carried out, allowing risks to be understood and managed, and appropriate maintenance planned in.

The approach can also be readily applied to other classes of problems in rotating machinery (such as blade vibration and high-cycle fatigue), or even in entirely separate areas. This can be achieved by changing the rotor dynamics models employed here for another physics-based model, which maps inputs of interest to measured outputs. This overall approach is therefore not limited to vibration in rotating machinery, and would also provide an effective diagnostic and condition monitoring tool in other contexts. We foresee the overall approach being applicable to owners, operators, maintainers or Original Equipment Manufacturers, as a key component of a predictive maintenance or digital twin system.

6 CONCLUSIONS

We have presented a new approach to interpreting measured vibration data, allowing it to give insight into other aspects of the condition of the system that are more difficult to measure. This is achieved by using physics-based rotor dynamics models, Gaussian Process surrogate models, and Bayesian inference, to infer the likely distributions of input parameters which result in the measured vibration. The method was shown to work well, and could readily be adapted to other classes of problem in rotating machinery or in other fields, as an effective condition monitoring or digital twin system.

References

[1] Friswell, Garvey, Penny and Lees: Dyamics of Rotating Machines, Cambridge University Press, 2010.

[2] M. Hazwan: M. Ghazali and W. Rahiman: "Vibration Analysis for Machine Monitoring and Diagnosis: A Systematic Review," *Shock and Vibration, vol. 2021,* p. 25, 2021.

[3] A. Lifshits: H. R. Simmons and A. J. Smalley: "More Comprehensive Vibration Limits for Rotating Machinery.," *ASME. J. Eng. Gas Turbines Power,* p. 583–590, 1986.

[4] G. Zurita, V. Sánchez and D. Cabrera: "A REVIEW OF VIBRATION MACHINE DIAGNOSTICS BY USING ARTIFICIAL INTELLIGENCE METHODS," *INVESTIGACION & DESARROLLO 16(1),* 2016.

[5] N. Vyas and D. Satishkumar: "Artificial neural network design for fault identification in a rotor-bearing system," *Mechanism and Machine Theory,* vol. 36(2), pp. 157-175, 2001.

[6] R.B.Walker, R. Vayanat, S. Perinpanayagam and I. Jennions: "Unbalance localization through machine nonlinearities using an artificial neural network approach," *Mechanism and Machine Theory,* vol. 75, pp. 54-66, 2014.

[7] J. Sinha, A. Lees and M. Friswell: "Estimating unbalance and misalignment of a flexible rotating machine from a single run-down," *Journal of Sound and Vibration,* vol. 272, p. 967–989, 2004.

[8] A. K. Jalan and A. Mohanty: "Model based fault diagnosis of a rotor–bearing system for misalignment and unbalance under steady-state condition," *Journal of Sound and Vibration,* vol. 327, pp. 604-622, 2009.

[9] T. Choudhury, R. Viitala, E. Kurvinen, R. Viitala and J. Sopanen: "Unbalance Estimation for a Large Flexible Rotor Using Force and Displacement Minimization," *Machines,* vol. 8, 2020.

[10] N. C. Tyminski and H. F. d. Castro: "Application of Bayesian Inference to Unbalance Identification in Rotors," in *Proceedings of the 9th IFToMM International Conference on Rotor Dynamics. Mechanisms and Machine Science*, 2015.

[11] N. C. Tyminski and F. W. S. Tuckmantel: "Bayesian inference applied to journal bearing parameter identification," *Journal of the Brazilian Society of Mechanical Sciences and Engineering,* 2017.

[12] G. Y. Garoli, D. S. Alves, T. H. Machado, K. L. Cavalca and H. F. d. Castro: "Fault parameter identification in rotating system: Comparison between deterministic and stochastic approaches," *Structural Health Monitoring,* vol. 20(6), 2021.

[13] N. Durrande and R. L. Riche: "Introduction to Gaussian Process Surrogate Models," 2017. [Online]. Available: https://hal.archives-ouvertes.fr/cel-01618068/document.

[14] Pedregosa et al: "Scikit-learn: Machine Learning in Python," *JMLR 12,* pp. 2825-2830, 2011.

[15] G. O. Roberts and J. S. Rosenthal: "Examples of Adaptive MCMC," http://probability.ca/jeff/ftpdir/adaptex.pdf, 2006.

[16] Sandia National Laboratories: "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.14 User’s Manual," https://dakota.sandia.gov/sites/default/files/docs/6.14/Users-6.14.0.pdf, 2021.