GT2019-91339

PREDICTING THE OPERABILITY OF
DAMAGED COMPRESSORS USING MACHINE LEARNING

J.V. Taylor[1] , **B. Conduit**[2] , **A. Dickens**[1] , **C. Hall**[2] , **M. Hillel**[2] , **R.J. Miller**[1]
[1] Whittle Laboratory, University of Cambridge, United Kingdom
[2] Rolls-Royce plc, Derby, United Kingdom
Email: jvt24@cam.ac.uk

ABSTRACT

The application of machine learning to aerospace problems faces a particular challenge. For successful learning a large amount of good quality training data is required, typically tens of thousands of cases. However, due to the time and cost of experimental aerospace testing, this data is scarce. This paper shows that successful learning is possible with two novel techniques: The first technique is rapid testing. Over the last five years the Whittle Laboratory has developed a capability where rebuild and test times of a compressor stage now take 15 minutes instead of weeks. The second technique is to base machine learning on physical parameters, derived from engineering wisdom developed in industry over many decades.

The method is applied to the important industry problem of predicting the effect of blade damage on compressor operability. The current approach has high uncertainty, it is based on human judgement and correlation of a handful of experimental test cases. It is shown using 100 training cases and 25 test cases that the new method is able to predict the operability of damaged compressor stages with an accuracy of 2% in a 95% confidence interval; far better than is possible by even the most experienced compressor designers. Use of the method is also shown to generate new physical understanding, previously unknown by any of the experts involved in this work. Using this method in the future offers an exciting opportunity to generate understanding of previously intractable problems in aerospace.

INTRODUCTION

This paper has two aims: First, it addresses the specific industry problem of predicting the effect of blade damage on compressor operability. Second, it presents a new methodology that enables machine learning to be applied to aerospace problems. It achieves this by coupling machine learning with rapid experimental test and physical parameters distilled from engineering wisdom.

Specific industry problem

Figure 1 shows damaged compressor blading that has been "boroblended", these blades have been cut back to remove cracks that were found during an aeroengine inspection. An accurate prediction of the impact of blade damage on operability needs to be made quickly in order to sentence the hardware and minimise disruption. Inaccurate predictions will result in hardware being condemned unnecessarily, this incurs significant cost penalties.

This expense is compounded by the increasing use of blisks. Now, instead of being able to replace individual damaged blades, the entire blisk must be scrapped or subjected to a complex repair. In order to overcome this problem in the future the industry needs to dramatically improve the accuracy of sentencing damaged hardware.

Current strategies for sentencing damaged compressors use simple correlations, evolved over many years, based on historic test
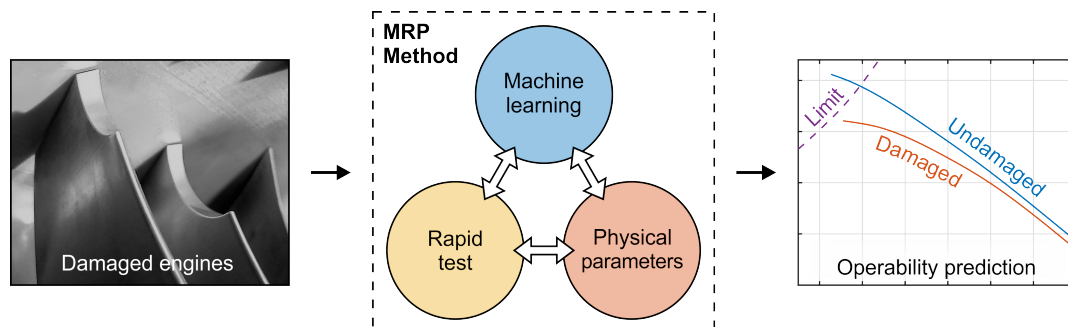

FIGURE 1: Predicting the impact of blade damage on operability using the MRP methodology

experience. These have very limited accuracy as they have little capability to take account of the detailed distribution of damage around the compressor. Relying on skilled human aerodynamicists to implement these strategies also has considerable cost, as their time is taken up with in-service issues instead of focusing on new design.

CFD techniques are also time consuming, labour intensive and not accurate enough. The damaged geometry must be measured, gridded, solved and analysed, this would currently take days of an expert's time. The prediction of the stall point also relies upon accurate modelling of large separations caused by the blunt blades; this is beyond current industry CFD methods.

It is clear that a fast and accurate method is required to tackle this problem and it must be based upon experimental data. The problem is gathering and processing enough training data so that the underlying mechanisms governing the operability of damaged compressors can be successfully learnt.

Machine learning problem

The new methodology is called the MRP method, it is shown in Figure 1. It is comprised of three components: **Machine learning** with neural networks, **Rapid test** to provide the training experimental data and **Physical parameters** drawn from engineering wisdom.

Machine learning will be used to predict the relationship between blade damage and compressor operability. The problem with machine learning is that it does not have prior understanding of the physics, therefore the quantity of data required for training is very large. Common practice suggests at least 10 data points are required for every dimension of the problem [1], in the case of leading edge blade damage the number of dimensions is the number of blades in the ring. This means order 1000 data points would likely be required, this far exceeds the number of rig tests commonly available to industry. To overcome this challenge two novel approaches are used:

Rapid test capability developed in the Whittle Laboratory has reduced the time to test a complete compressor stage by three orders of magnitude. A single engineer can now rebuild the compressor stage and test a new damage configuration in 15 minutes. This has allowed the characteristics and stall points of 125 different configurations of damaged blades to be tested in the course of this work.

Physical parameters were drawn from industry designers and academic researchers who have built experience over several decades. They are used to describe the most important aspects of a damage distribution for the neural network. This reduces the number of dimensions of the problem and therefore the required quantity of training data.

The combination of rapid testing and physical parameters ensures that enough training data can be taken to cover the number of dimensions of the problem. This has two powerful benefits: First, the mechanisms are successfully learnt by the machine and damaged blade operability can be predicted accurately. Second, the trained machine is able to generate new physical understanding of the problem.

This paper is presented in three sections. First, description of the individual components of the MRP method. Second, development of the whole method and how its components were used to inform each other in this process. Third, demonstration of the accuracy and generation of new physical understanding.
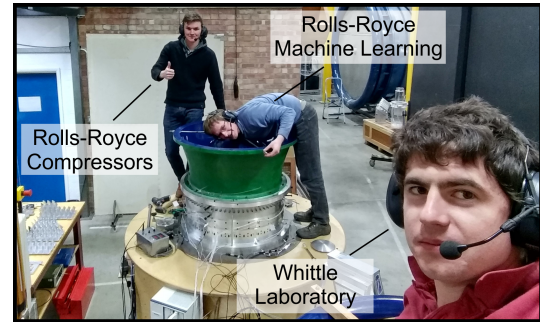

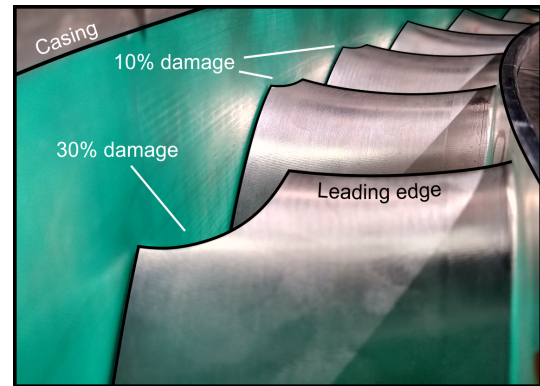**FIGURE 2**: Selfie of the industry and academic team

FIGURE 3: Damaged rotor blades installed in compressor rig

METHOD DESCRIPTION

In this section the three components of the MRP methodology are described: Rapid test, physical parameterisation and machine learning.

Rapid test

The rapid test capability was initiated by a value stream analysis of current compressor and turbine test techniques. Each process in the design, build and test cycle of a rotating rig was identified and analysed. The most time consuming processes were targeted and the compressor rig was redesigned to accelerate them or remove them from the testing cycle completely. It is now possible for a single engineer to test a configuration of damaged rotor blades every 15 minutes.

To make full use of the benefits of rapid testing a new philosophy of working between academia and industry is required. Instead of a top down bureaucratic style of project management, small autonomous teams are formed from both academic and industrial engineers. These teams allow the project to move more quickly and fully exploit the rapid test capability.

The test campaign was run by a single team of three people. Figure 2 shows a Rolls-Royce compressor aerodynamicist, machine learning specialist and Whittle Laboratory researcher building the next configuration to test. This combination ensured flexibility in both the experimental campaign and the iterative development of the MRP method.

The compressor rig is a single stage machine that is representative of a modern high pressure aero engine compressor. It has a rotor inlet relative Mach number of 0.4 and a Reynolds number of 5×10^5, further details are given in [2] and [3].

Four groups of rotor blades were CNC machined to model
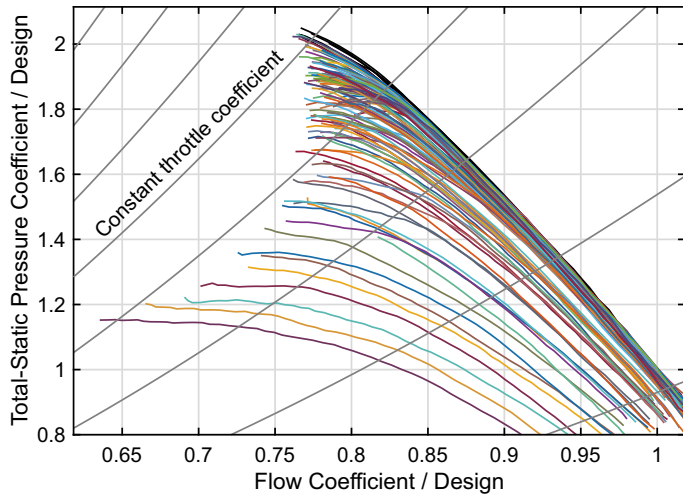
FIGURE 4: Measured characteristics of 125 damaged configurations

boroblended geometries with leading edge cutbacks of 5, 10, 30 and 50% of the chord. An example of a group of damaged blades installed in the rig is shown in Figure 3. The tip leading edges were cut in the ratio 1:1 in both the chordwise and spanwise direction.

The specific rapid testing improvements made to accelerate the damaged blade test campaign were:

- Design of the inlet and bearing assembly to allow access to the blades with minimal deconstruction; as the rotor disc is never removed alignment does not need to be repeated between builds.
- Blades are machined from solid with root, platform and blade in one piece so they can be assembled without stagger or tip clearance setting.
- Blade root clamps are segmented so blades can be replaced one or two at a time.
- Each rotor is individually counterweighted to account for the material lost in the damage, the whole rotor assembly does not need balancing between builds.

The total-static pressure rise (ψ) and the inlet flow coefficient (ϕ) was measured for every case, the characteristics of the 125 different configurations tested are shown in Figure 4. The static pressure measurement is located at the casing downstream of the stator. The inlet total pressure is measured with midspan Pitot probes at rotor inlet. These have been calibrated to calculate the total flow coefficient and integrated pressure rise with five hole probe area traverses. The compressor is throttled into stall three times and the stalling throttle coefficient is taken from the last stable recorded points.

Physical parameterisation

The parameterisation quantifies the damage configuration, if this can be based on physical understanding then the required quantity of training data can be reduced from order 1000 to 100. It is able to do this because it conditions the machine learning to the physics of the problem. It reduces the number of dimensions from the number of blades in the annulus down to the number of physical sensitivities identified.

To identify the physics of the problem a questionnaire was circulated to current and retired compressor aerodynamicists in the Whittle Laboratory and Rolls-Royce. In the questionnaire was a description
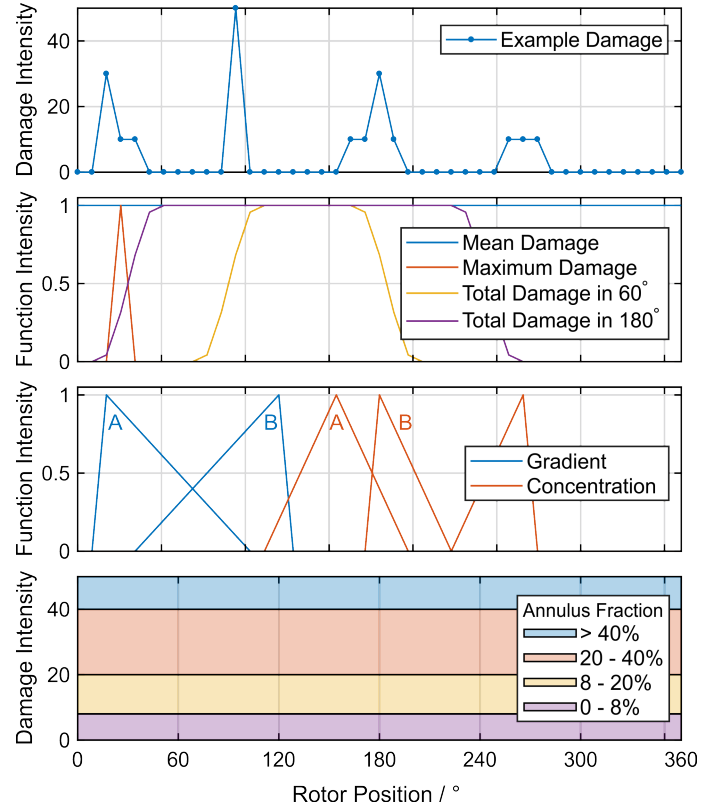
FIGURE 5: Top: Example damage distribution with four clusters
Second: Basis functions to measure circumferential distribution
Third: Basis functions combined to measure damage details
Bottom: Conditions to measure distribution of damage intensity

of the problem, a sample plot of damage around the annulus and example physical parameters (mean and maximum damage in the annulus). Both academic and industry engineers responded to provide exact functions that can be used to parameterise the damage, e.g. a DC60 measure, or more general ideas of sensitivities, e.g. it matters whether heavily damaged blades are on the edges of a cluster or in the centre. The questionnaire was phrased to encourage both of these types of responses to maximise the information that could be gained.

The information from the questionnaire was used to inform the choice of ten variables, they are shown in Figure 5 and summarised:

1. Mean magnitude of damage in the whole annulus
2. Maximum magnitude of damage in the whole annulus
3. Total damage within the worst 60° sector
4. Total damage within the worst 180° sector
5. Greatest damage gradient in a 75° sector
6. Greatest damage concentration in a 75° sector
7. Fraction of the row with greater than 40% damage
8. Fraction greater than 20% and less than 40% damage
9. Fraction greater than 8% and less than 20% damage
10. Fraction greater than 0% and less than 8% damage

Parameters 1-4 are calculated by convolving rotor damage as a function of position with the basis functions plotted in the second subplot of Figure 5. Take parameter 2 as an example, by convolving this with the damage function in the first subplot the magnitude of damage on a single blade is returned. The basis function is then
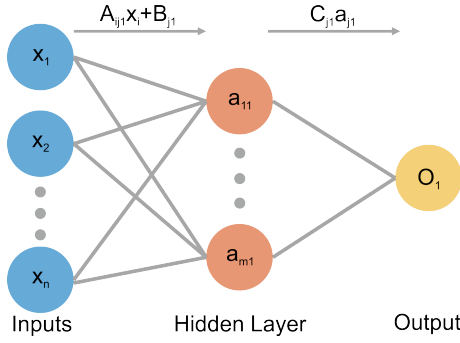
FIGURE 6: Schematic of neural network, where A_{ij}, B_j, C_j represent the machine learnt weights

indexed by one rotor position and the convolution is repeated. The maximum value is taken as the parameter for input into the neural network. This is equivalent to the maximum damage in the row. This process is repeated for parameters 1-4 to give a measure of the circumferential distribution of damage in the row.

Parameters 5-6 are calculated in two steps. First the damage function is convolved with basis function A, shown in the third subplot, and the maximum value is recorded after indexing to all positions. Then this is repeated for function B, the parameter input into the neural network is the difference between these two values. If the value is positive that means the damage function looks more like A, if it is negative it looks more like B.

Parameters 6-10 are conditional statements based upon the intensity of damage, they are shown in the fourth subplot of Figure 5. It was found that the stall point was a non-linear function of damage magnitude which the first six functions do not clearly discriminate. For example, the third parameter would treat the case of one blade with 50% damage the same as the case of five blades with 10% damage. Parameters seven and nine would highlight this particular difference clearly.

Machine learning

Figure 6 shows a schematic representation of the neural network model. This part of the paper describes three aspects of the network: Processing of the input data, the functions used in the neurons and the cross-validation strategy used to determine the complexity of the hidden layer.

Data inputs   Neural network modelling can be used to construct quantitative relationships in vast-arrays of quantitative data [4]. Neural networks are used in this paper to establish the unknown correlation between the defined physical parameters and the operability of the damaged compressor. Two networks are used: one to model the pressure rise characteristic and one to model the stall point.

The data for these two networks is presented in two different formats. For a damaged compressor the characteristic is calculated as a reduction in pressure rise coefficient relative to the datum undamaged case. 50 points across the stable part of the compressor's characteristic were used to train the pressure rise network. The stalling throttle coefficient (k) was used to train the stall point network, this is a measure of the exit area at the last stable operating point and is shown in Equation 1. Lines of constant k are shown in Figure 4.

$$k = \frac{\psi}{\phi^2} \qquad (1)$$

Neuron functions   A neural network is assembled using a combination of transfer functions, linear and non-linear curves contained within one or more "hidden layers" that can be used to model the data. These are combined to form a flexible mathematical function, which can well describe both linear and non-linear relationships. Therefore, neural networks can describe a large number of different mathematical relationships without laborious selection of basis functions. This makes neural networks ideal for the problem in this paper, at the start of this work the functions describing the effect of damage on operability were unknown and were likely to be a complex combination of linear and non-linear relationships.

To get the correct number of curves or neurons within the neural network function, cross validation [5] is used. By using a simple neural network we can exploit connected weights to understand the importance of each variable [6]. The neural network used can be described as follows:

$$O = \tanh \sum_{j=1}^{m} C_j \tanh \left(B_j + \sum_{i=1}^{n} A_{ij} x_{ji} \right) \qquad (2)$$

Where C_j, B_j, A_{ij} are learned weights from the fitting process, i is the number of input variables and j is the number of hidden nodes. To learn the weights from the data, we minimise the following equation which gives equal bias to each datapoint as a function of increasing error:

$$E = \sum_{k=1}^{l} \left| w_k \left(\tanh \sum_{j=1}^{m} C_j \tanh \left(B_j + \sum_{i=1}^{n} A_{ij} x_{ji} \right) - I_k \right) \right| \qquad (3)$$

Where k represents the number of datapoints, w_k represents the randomly chosen importance of the datapoint chosen by a cross validation strategy and I_k represents the relevant experimental result.

The complexity of the network is determined by the number of hidden nodes, which is given as i in Equations 2. The output is then predicted by the sum of the functions within the hidden layer. An adapted simulated annealing approach [7] was used with up to 10^7 iterations to ensure that the global minimum were found [8–10]. This robust fitting mechanism ensures a global minimum, in large neural networks this is not an issue due to their flexibility, but in small neural networks the possibility of obtaining a poor local minima or a saddle point is more probable [11].

Hidden layer complexity   In order to ensure the correct balance between generalisation and fit to the training data, nine models with a number of neurons ranging from 3 to 19 were created using a cross validation strategy. This ensures that over-fitting or retaining redundant features in the neural network [12] is avoided.

For the stalling throttle coefficient prediction, a neural network with 125 datapoints and 10 physical parameters was created. For the pressure rise neural network there were 6250 datapoints and the same 10 physical parameters plus the flow coefficient itself (ϕ), 11 variables total.

Data is randomly removed from the training set and a committee of models is generated with a range of different complexities. The standard deviation between these models is used to calculate the

4
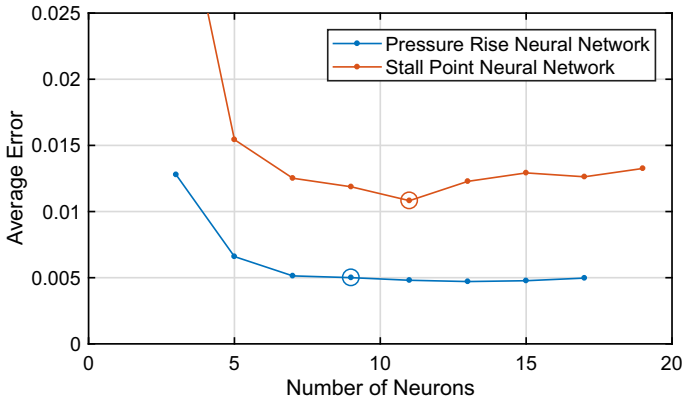
FIGURE 7: Average error from the committee of neural network models, showing the selected number of neurons
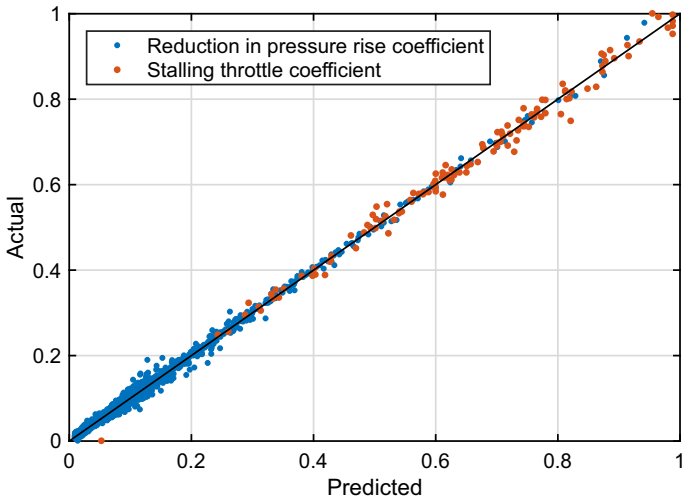

FIGURE 8: Fits to 100% of the data, stall point network $R^2 = 0.990$ and pressure rise network $R^2 = 0.997$

uncertainty. The results of this analysis are shown in Figure 7, the network with the minimum overall error is considered to be the best. For the neural network which models the stall point, the optimal number of neurons was found to be 11. Nine neurons were selected for the pressure rise neural network; in this case the error becomes constant at 0.005, to avoid over-fitting it is best to choose a simpler network and fewer neurons.

Figure 8 shows the fits achieved with both neural networks trained on 100% of the data. The data has been scaled between the actual observed maximum and minimum values in both cases. The coefficient of determination R^2 for the stall coefficient network was $R^2 = 0.990$ and for the pressure rise network it was $R^2 = 0.997$. The unpredictable variance is lower in the case of the pressure rise network compared to the stall coefficient network, 0.3% of the variance is unpredictable compared to 1.0%.

METHOD DEVELOPMENT

An iterative method was used to develop the MRP methodology. Results from the machine learning, physical parameterisation and rapid testing were fed back into each other during the development so that each component could be improved. Three parts of this process
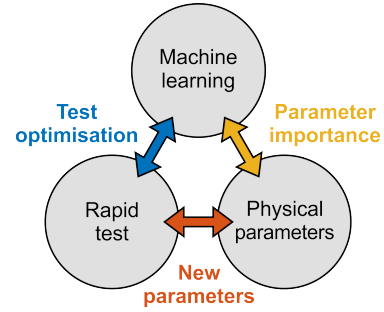
FIGURE 9: Iterative development of the MRP methodology

are shown in Figure 9 and are described in this section.

Test optimisation

In total 125 different damage configurations were tested, this data was taken in stages with analysis of the previous set before undertaking further tests. The initial configurations for test were chosen with the following criteria:

1. The neural network should interpolate between known configurations. Configurations with many heavily damaged blades ensure that the design space covers future in service issues.
2. The configurations should be representative of real life damage. Groupings and magnitudes of damage observed in engines on wing were replicated in the test compressor.
3. The accuracy should be highest close to sentencing limits where a damaged configuration is marginally acceptable. Greater numbers of configurations were focussed near this limit.
4. Systematically varying cases should be tested. These allow human aerodynamicists to investigate some trends without use of a neural network.

Once 50 cases had been tested a simple neural network was trained to help inform which configurations should be tested next. Two methods were used to select data for the remainder of the set:

First, the local uncertainty to a first approximation correlates with the local data density. Random configurations of damage were generated and the trained network was used to predict their operability. Cases with high uncertainty are likely to be in regions with a large distance between neighbours and so these configurations were tested. This approach effectively allowed the neural network to select the next cases to be tested.

Second, the network was trained on a random subset of the data gathered so far. Predictions of the operability of the remaining cases were made and compared to actual test data. Some of these predictions were poor because the case in question was not well represented by the training data. Further similar damaged configurations were tested and added to the set.

New parameters

While the authors were in the middle of the experimental campaign they plotted every new configuration tested against the last few cases, the datum and other extreme points. This allowed them to build an understanding of the effect of damage on the operability of the compressor and anticipate what the next result would be. This
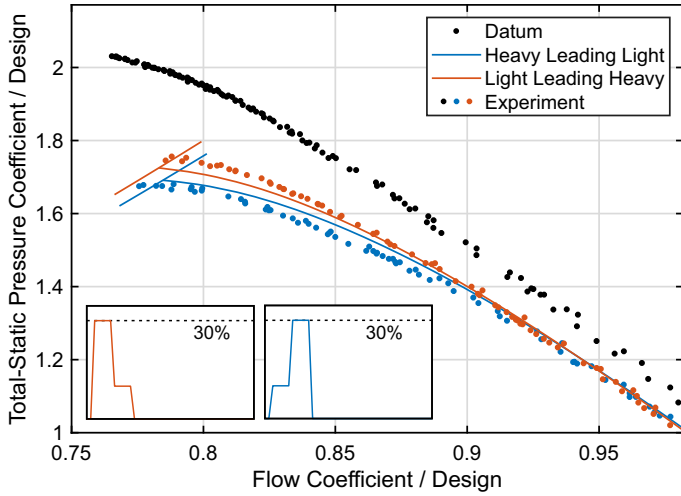
FIGURE 10: Prediction of two cases with a reversed order of damaged blades

FIGURE 11: Relative importance of parameters to the neural network

procedure allowed shortcomings in the current parameterisation to be noticed, explored and corrected.

One case where this became apparent was when the order of damaged blades in a cluster was reversed. Figure 10 shows two cases; a cluster is made up of five blades with 30% damage and five with 10%. If the more heavily damaged blades were placed at the leading edge of the cluster (in the direction of rotation) then the pressure coefficient and stalling throttle coefficient were reduced compared to the case where the more heavily damaged blades were placed at the trailing edge of the cluster. The parameterisation that was in use at the time could not discriminate between these two cases and so this sensitivity was impossible to predict.

A new parameter was introduced to measure the ordering of uneven damage across a sector of the annulus, this is the gradient parameter, shown in blue in the third subplot of Figure 5. If heavily damaged blades are leading light damage in a 75° sector then the basis function B is amplified while A is attenuated, parameter 5 is then negative. If light damaged blades lead heavy then the parameter is positive.

When this new parameter was incorporated into the neural network it was able to predict this sensitivity. These two test cases were removed from the training data and the operability was predicted, the results are shown as the solid lines in Figure 10.
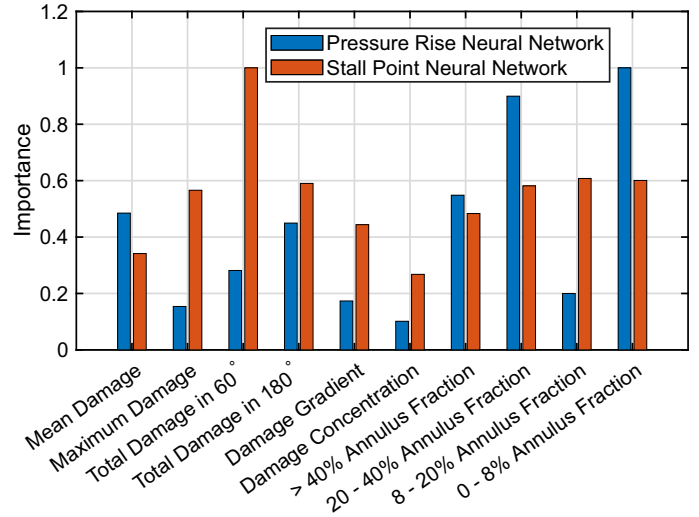
Parameter importance

Although machine learning can cope with many dimensions, the accuracy suffers if too many parameters are used or if they are not independent from each other. Machine learning was used to improve the parameterisation by using "connected weights" to calculate the importance of each variable [6]. The importance I_i, of a variable i is given by:

$$I_i = \sum_{j=1}^{m} A_{ij} C_j \qquad (4)$$

Where A_{ij}, C_j represent the non-bias machine learnt weights.

This method was used to optimise the sector sizes of the basis functions shown in Figure 5. Networks were trained with many sector sizes and the importance of each variable and the correlation

between predicted and actual results was examined. The parameters with low importance and poor correlation were discarded to give the final optimised parameters.

The importance is calculated for the final parameters for both the pressure rise and stall point neural networks and is plotted in Figure 11. This shows what kinds of damage are most important in affecting operability.

Figure 11 shows that the pressure rise depends upon longer length scale parameters. Mean damage and the damage intensity parameters are most important. The 180° damage measure is more important than 60° and the damage gradient and concentration have low importance.

The stall point depends upon shorter length scale parameters. Total damage in the worst 60° sector is by far the most important parameter. Maximum damage, gradient and concentration are also more important than in the pressure rise network. The conditional parameters also show that lighter damaged blades are most important for stall point prediction; 5%, 10% and 30% are more important than 50%.

METHOD DEMONSTRATION

In this section of the paper the MRP methodology is used for two purposes: First, the accuracy of the method's predictive capability is tested. 20% of the cases from the training data set are randomly removed, this means there are 100 cases to train and 25 to test the accuracy against. Second, the method is used to increase our physical understanding. By examining the predictions of outliers and selected cases it is possible to probe the underlying physical mechanisms that govern the operability of damaged compressors.

Method accuracy

Damage incidents that occur on engines are likely to be unique and so the trained neural network must be used to interpolate the experimental data. To simulate new cases and test the prediction capability of the method 20% of the cases are randomly removed and set aside for validation. The network was then retrained on the remaining 80% of the dataset. The pressure rise coefficients and stalling throttle coefficients were predicted for the test data set and the results were

Copyright © 2019 by ASME
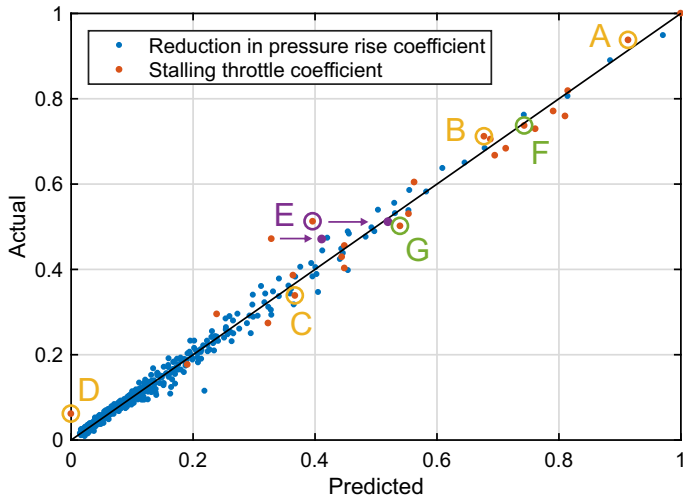
FIGURE 12: Prediction capability of the method for 20% of unseen test cases. Characteristics of cases A-D plotted in Figure 13, case E in Figure 14 and cases F-G in Figure 16

compared to those measured. The agreement is plotted in Figure 12.

As expected, the correlation is not as strong as that of 100% of the training data presented in Figure 8. The coefficient of determination R^2 for the stalling throttle coefficient network was $R^2 = 0.960$ and for the reduction in pressure rise network it was $R^2 = 0.989$. Therefore, for the stall coefficient network, 4.0% of the variance in the predicted verses experimental values is unaccounted and for the pressure rise network, 1.1% of the variance is unaccounted. The accuracy of the stalling throttle coefficient prediction is 2% in a 95% confidence interval.

To determine the benefit of using the physical parameterisation over other methods a separate network was trained to predict the stall point. This separate network was trained on the raw values of damage magnitude at each blade location. This resulted in many more input variables than for the physical parameterisation; 42, one for each blade, instead of 10, based on physical importance. When this network is used to predict the operability of the same 20% of removed test cases, it was only able to predict the stall point with $R^2 = 0.82$. It can be concluded that the physical parameterisation dramatically improves the accuracy of the method.

To visualise what the accuracy looks like to an aerodynamicist four cases are selected and the predicted characteristics are plotted. The four cases chosen are marked in Figure 12 as A, B, C and D, their characteristics are plotted in Figure 13. These points are chosen as they cover the wide range of cases investigated. It can be seen that the accuracy of this tool is extremely good, especially when compared to other predictive tools conventionally used in aerospace, such as 3D CFD.

Generating understanding - Case 1

Outlier cases can be used to guide the generation of new physical understanding. Consider case E plotted in Figure 12, it is the most extreme outlier in the current model. The measured and predicted characteristics for this case are shown in Figure 14. It can be seen that the predicted stalling throttle coefficient, shown by the red line, is lower than that measured, shown by the final blue marker.
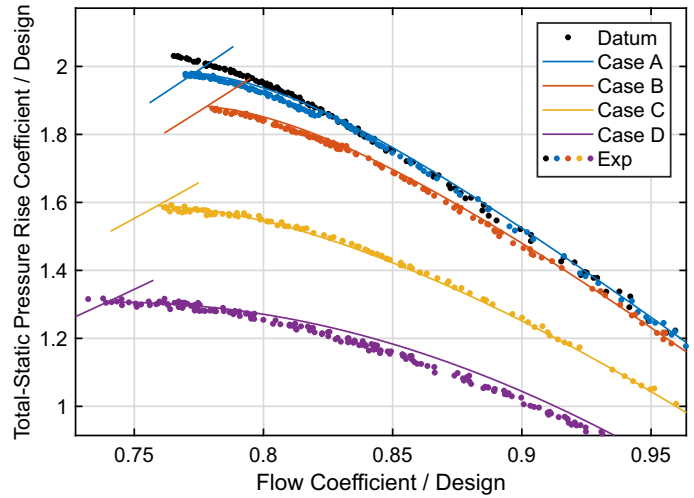
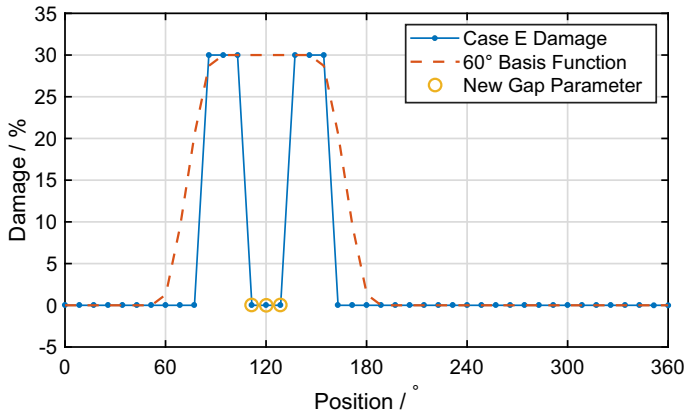**FIGURE 13**: Prediction of the stall coefficient and pressure rise of four random unseen cases A-D
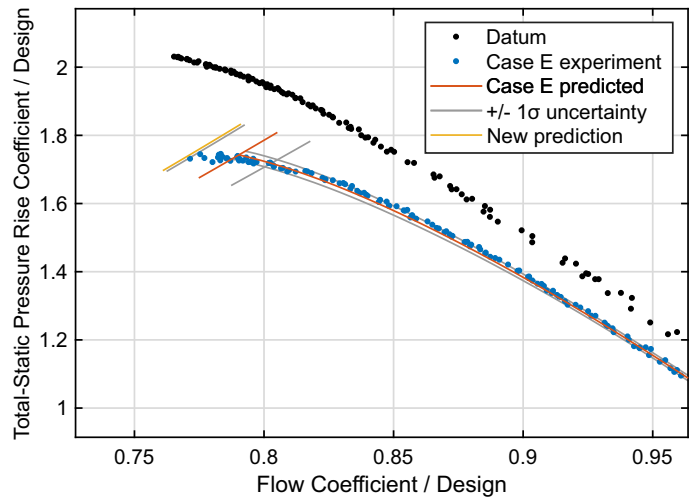
FIGURE 14: Operability prediction of outlier case E

In absolute terms the accuracy is still good and the neural network does recognise a higher uncertainty in this region, this is shown by the grey lines which cover the actual stall point. In practice an aerodynamicist would also use this uncertainty data to inform their decision successfully. However, in this paper it is useful to probe the cause of this discrepancy so that we can learn about the physical mechanisms which set the operability.

Inaccuracies can occur because of three reasons: First, there is a lack of similar cases in the training data set. This was avoided by checking for inaccurate predictions during the experimental campaign so gaps in the data set could be filled. Second, the neural network has the incorrect complexity and either over-fits or over-smooths the data. This was avoided by performing the cross-validation strategy to achieve the optimal network complexity. Third, the physical parameterisation is incomplete and it is impossible for the machine to learn the physical reason why case E has the operability that is does.

In this case the parameterisation is incomplete. Shown in Figure 15 is the distribution of damage in case E, it is made up of two clusters of three 30% damaged blades separated by three undamaged blades. The gap will be important in increasing the operability

FIGURE 15: Development of new parameter to measure gap between clusters

FIGURE 16: Prediction of two cases with different intensities of damage

FIGURE 17: Ensemble averaged, filtered casing pressure at rotor leading edge

compared to a case where all six damaged blades were neighbouring each other, but the presence of this gap cannot be detected by the current parameterisation. The basis function for the 60° parameter is plotted in Figure 15, note that this would have the same value whether there was a gap or not. The other nine parameters also fail to flag the gap's presence.

A new parameter must be developed so that the neural network can learn about gaps between clusters and incorporate this into its predictions. Over the 60° window with the worst damage the number of undamaged blades are counted, this number of undamaged blades as a fraction of the annulus is then used as the new parameter. The new parameter along with the original ten are calculated for the same 80% of the training data, the stall point network is then retrained and the operability of case E is predicted again.

The results are plotted in Figure 14, the stall point is shown by the yellow line which now matches the measured stall point. This improvement is also shown in Figure 12 where the new predictions are marked in purple. Note that the neighbouring outlier to case E also has a gap between clusters and its prediction is also improved.

The new physical parameter was shown to a number of experts who developed the original 10 physical parameters. They could immediately reason why the gap parameter must be important in determining the operability. Close to the stall point the gap inhibits the propagation of three-dimensional separations from one cluster to the next. It is clear that the MRP methodology has augmented their ability to develop new physical understanding.

Generating understanding - Case 2

It is also possible to generate physical understanding by inspecting cases which perform contrary to expectation. Consider cases F and G, shown in Figure 12. Both of these cases have the same mean level of damage but case F is made up from two 50% damaged blades and case G is made up from ten 10% damaged blades. Figure 16 shows that the 10% damaged blades have a reduced stalling throttle coefficient compared to the 50% damaged blades.

The operability depends non-linearly on the magnitude of the damage. If the cluster is made up of many lightly damaged blades then the compressor was found to lose flow range but maintained a similar pressure rise to the undamaged case at high flow rates. If the cluster is made up of heavily damaged blades the compressor
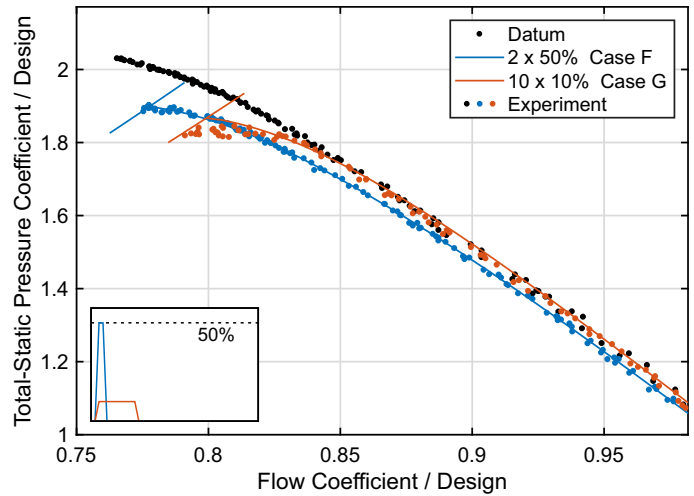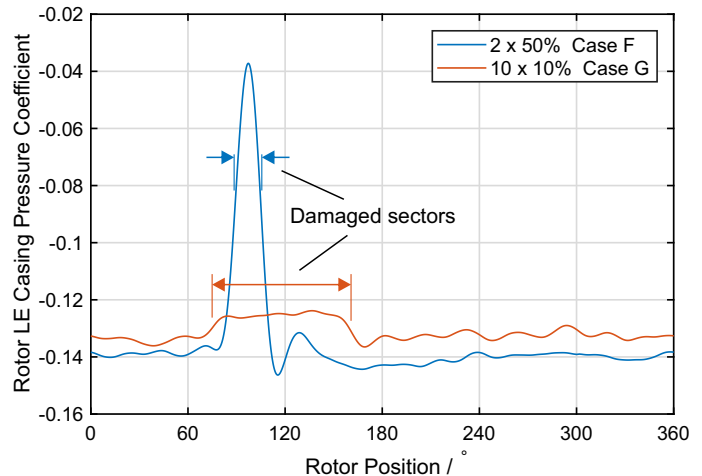
did not lose as much flow range but did have a reduced pressure rise.

For the machine to learn this mechanism it must be able to correlate changes in the damage configuration from multiple parameters. Between these two cases changes in five of the parameters are seen: the maximum damage, the damage in a 60° and 180° sector and two of the conditional measures of damage magnitude. From Figure 16 it can be seen that the network is able to make accurate predictions of the operability in both cases, the change in the pressure rise characteristics and the stalling throttle coefficients are both well predicted.

To investigate the cause of this non-linear behaviour detailed measurements were performed on the two cases. Unsteady static pressure measurements were made on the casing upstream of the rotor leading edge. The results are filtered to remove the blade passing signal and then ensemble averaged at a normalised flow coefficient of 0.84, the resulting static pressure coefficient is plotted in Figure 17.

Figure 17 shows a lower static pressure at the tip of all of the undamaged blades in case F, this must be caused by an increased flow rate through the undamaged sector. The blockage caused by the 10% damaged blades is not as great and thus little flow redistribution occurs in case G. For the compressor to stall, leading edge separations

8

must occur and propagate on the undamaged blades [13] and [14]. The increased flow in the undamaged blades in case F therefore delays the onset of separation to a lower flow coefficient. In case G the lack of flow redistribution increases incidence and allows separations to propagate more readily from the damaged sector onto the undamaged blades.

Building a complete model of the flow mechanisms using the MRP methodology made this understanding possible. Previous methods correlating 1 or 2 parameters across a handful of test cases have not been able to detect the subtlety of the flow mechanisms which are occurring in these two cases.

CONCLUSIONS

The power of the MRP methodology (Machine learning - Rapid test - Physical parameterisation) in tackling complex aerospace problems has been demonstrated. In the future this method will become common in tackling problems in design, manufacture and through life performance.

This paper has shown three benefits of applying the MRP methodology to concess the aerodynamic impact of blade damage: First, the impact can be predicted more accurately than by human designers; this method can predict the operability with an accuracy of 2% in a 95% confidence interval. Secondly, the impact can be judged with more certainty; the statistical methods in the machine learning algorithm determine the uncertainty of every prediction. Finally, the use of the tool ensures the time of skilled aerodynamicists can be saved and their efforts can be directed to problems of design.

The authors believe that in order to successfully apply machine learning in aerospace three changes in approach are required: First, a new generation of experimental test facilities must be developed. New geometries must now be tested in minutes or hours in order to generate enough training data for machine learning to succeed. Secondly, engineering wisdom must be formalised and recorded, this reduces the necessary quantity of training data. The success of the aerospace industry to date has been based upon building and retaining wisdom, it is no surprise that in order for machine learning to succeed this wisdom must be leveraged. Third, a new approach to research and development should be taken. Experts in machine learning, and experts in aerodynamics, should be co-located in autonomous teams rather than managed from the top-down. This gives the teams the freedom and flexibility necessary to successfully develop the new methodologies.

Finally, the most exciting outcome of this paper is the way in which the MRP methodology augmented the human in the generation of new physical understanding. This represents an opportunity for developing new understanding in problems throughout the aerospace sector.

REFERENCES

[1] Haldar, M., 2015. How much training data do you need? https://medium.com/@malay.haldar/how-much-training-data-do-you-need-da8ec091e956.

[2] Taylor, J. V., 2015. "Three-dimensional mechanisms in compressor flows". PhD thesis, Department of Engineering, University of Cambridge.

[3] Taylor, J. V., and Miller, R. J., 2016. "Competing three-dimensional mechanisms in compressor flows". *Journal of Turbomachinery, 139*(2), Oct., pp. 021009–021009–10.

[4] Verpoort, P., MacDonald, P., and Conduit, G., 2018. "Materials data validation and imputation with an artificial neural network". *Computational Materials Science, 147*(1), pp. 176–185.

[5] Stone, M., 1974. "Cross-validatory choice and assessment of statistical predictions". *Journal of the Royal Statistical Society. Series B (Methodological, 36*(2), pp. 111–147.

[6] Ona, J., and C., G., 1974. "Extracting the contribution of independent variables in network models: a new approach to handle instability". *Journal of the Royal Statistical Society. Series B (Methodological, 36*(2), pp. 111–147.

[7] Conduit, B., and Conduit, G., 2013. Method and system for designing a material, EP2778990, US20140236548, GB201302743D0, 02.

[8] Ingber, L., and Rosen, B., 1992. "Genetic algorithms and very fast simulated reannealing: A comparison". *Mathematical and Computer Modelling, 16*(11), pp. 87–100.

[9] Mahfoud, S., and Goldberg, D., 1995. "Parallel recombinative simulated annealing: Agenetic algorithm". *Parallel Computing, 21*(1), pp. 1–28.

[10] Gelman, A., Roberts, G., and Gilks, W., 1996. "Efficient metropolis jumping rules". *Bayesian Statistics, 5*(1), pp. 599–607.

[11] Choromanska, A., Henaff, M., Mathieu, M., and Arous, G., 2015. "The loss surfaces of multilayer networks". *Processings of the 18th International Conference on Artificial Intelligence and Statistics, 38*(1), pp. 176–185.

[12] Denil, M., Shakibi, B., Dinh, L., Ranzato, M., and Freitas, N. D., 2013. "Predicting parameters in deep learning". *arxiv.org/abs/1306.0543v2*.

[13] Pullan, G., Young, A. M., Day, I., Greitzer, E., and Spakovszky, Z., 2014. "Origins and structure of spike-type rotating stall". *ASME J. Turbomach.*

[14] Hewkin-Smith, M., Pullan, G., Grimshaw, S. D., Greitzer, E., and Spakovszky, Z., 2018. "The role of tip leakage flow in spike-type rotating stall inception". *Journal of Turbomachinery*, Dec.