

GT2021-60283

REDUCING UNCERTAINTY IN THE ONSET OF COMBUSTION INSTABILITIES USING DYNAMIC PRESSURE INFORMATION AND BAYESIAN NEURAL NETWORKS

Michael McCartney*
GE Aviation
Freisinger Landstr. 50
85748 Garching, Germany
Email: michael1.mccartney@ge.com

Ushnish Sengupta
Matthew Juniper
Department of Engineering
University of Cambridge
Trumpington Street
Cambridge CB2 1PZ, UK

ABSTRACT

Modern, low emission combustion systems with improved fuel-air mixing are more prone to combustion instabilities and therefore use advanced control methods to balance minimum NOx emissions and the presence of thermoacoustic combustion instabilities. The exact operating conditions at which the system becomes encounters an instability is uncertain because of sources of stochasticity, such as turbulent combustion, and the influence of hidden variables, such as un-measured wall temperatures or differences in machine geometry within manufacturing tolerances. Practical systems tend to be more elaborate than laboratory systems and tend to have less instrumentation, meaning that they suffer more from uncertainty induced by hidden variables. In many commercial systems, the only direct measurement of the combustor comes from a dynamic pressure sensor. In this study we train a Bayesian Neural Network (BNN) to predict the probability of onset of thermoacoustic instability at various times in the future, using only dynamic pressure measurements and the current operating condition. We show that, on a practical system, the error in the onset time predicted by the BNNs is 45% lower than the error when using the operating condition alone and more informative than the warning provided by commonly used precursor detection methods. This is demonstrated on two systems: (i) a premixed hydrogen/methane annular combustor,

where the hidden variables are wall temperatures that depend on the rate of change of operating condition, and (ii) full scale prototype combustion system, where the hidden variables arise from differences between the systems.

NOMENCLATURE

BNN Bayesian Neural Network
 $DFA(P')$ Detrended Fluctuation Analysis transform
 $FFT_w(P')$ Fast Fourier Transform using Welch's method
 $insta(P')$ Binary indication of instability transform
FPCS Full-scale Prototype Combustion System
LL Log Likelihood
MAP maximum a posteriori
ML Machine Learning
OOD out-of-distribution
OP Operating Parameters
P2P Peak to Peak
ReLU Rectified Linear Unit
 \dot{m} Normalised air mass flow rate
 Δt_{pred} Difference between prediction and current time
 t_{TTI} Time to instability
 P' Pressure Fluctuations
 T_w Wall Temperature
 Φ Equivalence Ratio

*Address all correspondence to this author.

INTRODUCTION

With increasingly stringent regulation on emissions from aero and power gas turbines, manufacturers are often turning to lean premixed combustion systems in order to reduce peak temperatures and hence the production of NOx. With this shift to lean premixed combustion, comes an increased propensity for combustion systems to exhibit thermoacoustic instabilities. These instabilities are difficult to model and complex in nature which means that, despite efforts to eliminate them through good design, they often still exist under certain conditions late in the design process, where a complete redesign incurs significant costs. In addition to good design, passive and active control strategies are therefore often employed in order to suppress instabilities e.g. resonators are integrated into the combustor design and tuned to the frequency of instabilities that occur near on-design conditions as a form of passive control and fuel staging has been found to be an effective form of active control in both power [1] and aero [2] applications. In fuel staging, a rich-burn pilot injector is used alongside the lean-burn main injector to allow an extra degree of freedom to vary the conditions within the combustor whilst maintaining a constant fuel flow rate. While this is often used for maintaining stability at low power conditions, the GE TAPS combustor combined fuel staging with advanced control methods to optimise the trade offs between NOx productions, combustion efficiency, operability and combustion dynamics across the entire operating envelope [2].

Although fuel staging allows the complete avoidance of combustion instabilities, it necessitates margins around instability zones in the operating space wherein more optimum operating points, in terms of NOx emissions or combustion efficiency, may lie. Furthermore, the exact operating condition where the system will encounter instability can vary between systems in addition to the uncertainty caused by stochastic forcing of the system from turbulence in the flow or turbulent combustion. These differences can be attributed to hidden (i.e. unmeasured or unmeasurable) variables such as combustor wall temperatures, where the hostile environment makes measurements on a production system prohibitively expensive, or differences in geometry between systems that are within manufacturing tolerances but nonetheless subtly affect the dynamics of the system. Additionally, degradation of the system over time (e.g. component wear or coking of fuel nozzles) can also cause these instability onset conditions to shift with continued use [3]. In this study, we show that dynamic pressure measurements of the combustor can be used to reduce the uncertainty in the operating condition where that specific system will exhibit instabilities, which could allow the system to operate nearer to a local optimum, without triggering instability, than when relying on operating conditions alone.

There has been significant work by the thermoacoustic com-

munity around the detection of precursors to combustion instabilities using dynamic pressure measurements. All of the methods proposed thus far use some kind of statistical measure to detect a transition away from the low amplitude stochastic behaviour that characterises normal combustion noise. The most successful methods either look for a departure from chaotic behaviour in the state space or look for specific symptoms of precursors to instability. The methods that look for departure from chaotic behaviour borrow for the plethora of techniques for analysing dynamical systems in mathematics literature [4]. The first attempts created a representation of the state-space using time-embedding methods and then looked at the predictability of the embedded state over time using methods such as the Translation error [5], Lyapunov exponent [6], Symbolic Time Series Analysis (STSA) [7] and Complex Networks [8], which were all able to show a measurable change in behaviour of the system before there was an increase in the pressure fluctuation magnitude. In the case of precursors, it was noted that intermittency could often be observed in combustion systems as brief sections of periodic behaviour amongst the combustion noise. Intermittency is a phenomena known to occur when a system is within the fold of a sub-critical Hopf bifurcation where the system jumps between the initial linearly stable branch to the limit cycle branch due to external forcing [9] and the Hurst exponent was shown to be an effective measure of this by Nair et al. [3] [10].

Machine Learning has often been used in efforts to detect precursors including attempts to: estimate the level of chaos by examining the error in predictions of pressure fluctuations made by neural networks trained on the signal itself [11]; use Hidden Markov Models to classify the state using the output of STSA [12] or directly from pressure measurements [13]; or use nonlinear methods (e.g. SVM, Random Forest, Neural Network) to combine multi-dimensional outputs from precursor detection into a single prediction [14] [15]. One commonality of the literature in this area is that the methods warn that an instability is approaching, rather than give an indication of when it will occur, as would be required for it to be used to decide how to control the system. They also look only at the pressure fluctuations and ignore the operating parameters, which are good at indicating when an instability will occur as well as providing information on how fast the system is moving through the operating space. In this study we use Bayesian Neural Networks to predict the probability of the system encountering instability in the future operation of the system along with a confidence in the prediction that reflects whether the prediction point is close to the training domain and thus reliable. We also show that combining information from the operating parameters with the dynamic pressure information significantly improves the prediction of when an instability will occur beyond using operating parameters alone.




FIGURE 1. Left: Image of the model gas turbine under operation. Right: Schematic of side and top view.

TEST DATA

Two sources of data were used in this study, each demonstrating that our method can improve the predictions in the onset of instability with two different types of hidden variable.

The first data set for this study was taken from experiments conducted using the atmospheric, 50kW-100kW annular combustion chamber from NTNU described in [16] and shown in Figure 1. In this rig, a premixed fuel-air mixture is fed into a cylindrical plenum that conditions and divides the flow into 12 axisymmetric burners. The burners are comprised of 150 mm long tubes with bluff bodies of 18.95 mm diameter equally spaced around the circumference of the annular combustor. The inner and outer walls of the chamber are 120 mm and 300 mm long respectively, with diameters of 127 mm and 212 mm. The chamber walls are water cooled, enabling long run times and equilibrium wall temperatures. The fuel is composed of 87.6% hydrogen and 12.4% methane by volume. The power is set to a value of 4, 6 or 8 kW per burner and the equivalence ratio, Φ is varied in the range of 0.4 to 0.7 by controlling the air mass flow rate, \dot{m} , which is measured at 3Hz. 12 dynamic pressure sensors are mounted at six azimuthal and two longitudinal positions flush with the inner wall of the injector tubes and recorded with a sampling frequency of 51.2 kHz.

Equivalence ratio sweeps were performed by first igniting the combustor and then keeping it at the initial equivalence ratio until the cooling water temperature settled. The air mass flow rate and thus Φ , was then ramped linearly over a ramp time of 20 or 60 s until the final operating condition was reached. The fuel mass flow was kept constant throughout. The initial and final operating points were chosen so that the combustor started in a state characterised by combustion noise and finished in a limit cycle state, with the onset of the instability occurring during the ramp. The instability encountered was an axial mode, with the mass flow rate of onset varying for different power conditions.

In total 30 data sets were generated: five for each ramp rate at three different power conditions. At all power conditions there exists an uncertainty in the exact mass flow rate at which the system encounters instability, as defined by a threshold on the peak-to-peak pressure fluctuations. This uncertainty is reduced using the knowledge of the ramp rate or the wall temperature measurements. These parameters are obviously related by the fact that the system has less time to reach a thermal equilibrium on the faster ramp and thus the temperature at a given point in the ramp is different. We use the wall temperature measurements, T_w , in this case as an example of a hidden variable that depends on the earlier state of the system and to which the information from the dynamic pressure measurements will be compared.

The second data set was taken from tests carried out on full-scale prototype combustion systems (FPCS) at sea level conditions. The FPCS were equipped with low frequency measurement instrumentation, that was representative of what might be found on practical engines, with the addition of two dynamic pressure sensors on the cold side of the combustor, at two azimuthal locations. The pressure measurements were taken at a sampling frequency of 25kHz while the other measurements were taken at a sampling rate of 20 Hz. The low frequency measurements relevant to the control of the combustor were used to describe the state of the combustor. Since the instrumentation reflected that of a real practical however, this did not include any direct measurement of the combustor, but instead consisted of measured and control variables including: compressor exit temperature and pressure, fuel flow rate, primary/secondary fuel split and core speed, ω_{HP} . For the purpose of this study, FPCS were intentionally run at operating points where thermoacoustic instability above a certain acoustic amplitude level is expected. On a subset of these points, the peak-to-peak pressure exceeded a given threshold, indicating thermoacoustic instability. The exact operating condition where each system encounters instability is not identical and varies within a limited range of the low frequency measurements. Within this range however, there is an uncertainty in the operating condition at which the instability will be triggered. We again show that this uncertainty, partly due to hidden variables, can be reduced by using dynamic pressure measurements, where in this case the hidden variable represents slight differences between systems (e.g. geometries) and differences in trajectories through the operating space that are not captured by the low frequency parameters.

BAYESIAN NEURAL NETWORKS FOR PREDICTING THERMOACOUSTIC INSTABILITIES

Overconfident point estimates make ordinary neural networks unsuitable for use in high-risk domains, such as a power plants or jet engines, and Bayesian statistics can provide a natural framework for estimating predictive uncertainties. A Bayesian Neural Network places a prior probability distribution over net-

work parameters (weights and biases) which is updated using observational data by applying Bayes' rule. The distribution over outputs, which results from having distributions over all network parameters rather than discrete values, makes the model robust to overconfident extrapolations on input data which is substantially different from training data. Bayesian inference is expensive to perform in neural networks because of the high dimensional nature of the parameter space. In this study, we use Randomized MAP sampling, an approximate inference technique introduced by Pearce et al. [17], which is a computationally efficient approach. In this method, m neural networks with identical architecture are created, with weights randomly initialised from a prior distribution. The neural networks are then trained using standard backpropagation methods with an additional regularisation term that penalises the deviation of each weight from an anchor, which is also selected from the prior distribution over parameters. The resulting ensemble will have an updated distribution over parameters with predictions that converge when well supported by training data and diverge when making predictions outside of the training domain. The standard deviation of the predictions can therefore be used as a measure of the confidence of the model's prediction, which is a major advantage over standard neural networks.




FIGURE 2. During BNN training, individual NNs are created and trained with anchoring terms on all weights sampled from untrained BNN weight distributions. The ensemble of the trained NNs makes up the trained BNN.

In the context of predicting the onset of thermoacoustic instabilities, we are interested in the probability that a system will exhibit instability at a time in the future, $P(U_t)$. We assume that this depends on the current state of the system, the future operating parameters and how fast it will reach them. The model will therefore predict the state of the combustor given this information, i.e.

$$P(U_{t+\Delta t} | OP_{t+\Delta t}, \Delta t, OP_t, f(P'_t)) \quad (1)$$

where t is the current time, from which we make the prediction, Δt is the difference between t at the time for which we are pre-

dicting the combustor state, OP are the operating parameters and $f(P'_t)$ is a transformation of the pressure measurements made at time t .

In order to make a prediction, we concatenate the vector of variables describing the current state, the time at which we would like to predict and a vector of the operating parameters corresponding to that time. We then pass this vector as an input to the BNN, allowing us to make predictions of the system exhibiting instability at any time and operating point in the future. This is shown in Figure 3 for predictions made at several times in the future.




FIGURE 3. The structure of BNN inputs and outputs when predicting the system state at different time in the future.

We can therefore evaluate the probability of the system encountering instability along multiple candidate trajectories and then select the safest trajectory, using the standard deviation of the prediction as an indicator of whether we can trust the prediction. Since we do not pass the intermediate operating points to the BNN it is implicitly assumed that the system will move between operating points in the same way as it does in the training dataset.

The number of hidden layers and the number of neurons in each layer were hyperparameters that required tuning. The tuning was carried out using a random search over the hyperparameters, selecting those which gave the lowest negative log-likelihood on the validation data set. The neuron activation function was set as ReLU and the networks were trained using the Adam optimisation algorithm, until the negative log-likelihood of the validation data stopped decreasing [18].

The training, tuning and testing of the BNNs are carried out in different phases. Each phase uses a different and non-overlapping subset of the data in order to ensure that no 'data leakage' occurs. This is required as the prediction of models made on the training and tuning data sets will be better than on an 'unseen' data due to steps taken during the training and tuning processes to optimise the model on the training and tuning data sets respectively. Crucially the models are not run on the test data set until the final results are generated to ensure they are not biased by the training and tuning process. In the tuning process the hyperparameters of the BNN, namely the number of hidden layers (varied between 2 and 10) and the number of neurons in each layer (varied between 10 and 100) are optimised. This is done using a random search over the hyperparameters, for each

combination of which a BNN is trained using the training data and evaluated on the tuning data set.

In the training phase, the BNN is set up with an architecture determined by a given set of hyperparameters. As described previously, an ensemble of neural networks are created with their weights initialised from the prior distribution. An anchor parameter that is used in the training regularisation is also sampled from this distribution. The prior of the network was chosen so that predicted probabilities made by the untrained ensemble fell between 0 and 1 with a mean of around 0.5 when passed the training data. For the first layer this required the variables to have a variance equal to the number of independent samples in our training data set. Since all of the input features are time traces, there exists significant correlation between the data samples. The number of independent samples was therefore estimated by taking the number of samples in the training data and dividing it by the mean time-shift required for the auto-correlation of each feature to fall to zero. The variance of the prior for the following layers was set equal to the inverse of the number of neurons in that layer. Once set up, the NNs are trained on the training data set using back-propagation with the Adam optimisation algorithm [18] to minimise the negative log-likelihood (a.k.a. cross-entropy) and the regularisation term for randomised MAP sampling from Pearce et al. [17]

$$\frac{1}{N} \|\mathbf{\Gamma}^{1/2}(\boldsymbol{\theta} - \boldsymbol{\theta}_{anc})\|_2^2 \quad (2)$$

where $\boldsymbol{\theta}$ are the NN parameters, $\boldsymbol{\theta}_{anc}$ are the anchors for each parameter sampled from the prior distribution and $\mathbf{\Gamma}$ is a regularisation matrix for which $diag(\mathbf{\Gamma})_i = \frac{1}{2}\sigma_{prior,i}^2$ for the distribution over each parameter, θ_i . The number of training epochs run in the training process is another hyperparameter that must be tuned. To do this, at the end of each training epoch, the loss function is evaluated on the tuning data set and the training is halted once the loss stops decreasing.

For classification tasks, such as the prediction of the onset of instability shown later in the study, a threshold probability must be set, above which the model is considered to predict the sample as stable or exhibiting instability. This is determined for the final model by making prediction of the probability of onset of instability on the validation data and choosing the threshold that minimises the error in the onset prediction. Sample code showing how this approach can be implemented can be found at <https://github.com/mccartney-ge/GT2021-60283>.

PRESSURE DATA TRANSFORMATION

McCartney et al. [15] showed that one of the advantages of using machine learning to predict the onset of instability is the

ability to control the sources of information available to the algorithms and compare the resulting changes in predictive performance in order to evaluate the relative importance of the information source. In order to test the comparative advantage that is provided by the information carried in the dynamic pressure signal, BNNs were trained using the framework described previously (the future operating condition was omitted in the Annular Rig case, as will be discussed later), using different sets of variables to describe the current combustor state. In theory, any number of transformations from multiple pressure measurements could be included as inputs to the BNN; however, in this study data from a single pressure measurement is taken and transformations are considered in isolation. As all the instability modes were axial, there was no difference in results when pressure signals from different sensors were used. The current operating conditions were used as a baseline description of the current state and this was compared with the current operating conditions concatenated with a transformation of a sample of the dynamic pressure measurement. Three transformations were considered in this study and are described in more detail below: $insta(P')$, $FFT_W(P')$ and $DFA(P')$. In all cases the dynamic pressure signal was down sampled to 25kHz and a signal segment length of 4096 data points, corresponding to approximately 160ms, was used for all of the transformations.

Binary Indication of Instability ($insta(P')$) is a binary indication of whether the peak-to-peak pressure threshold has been exceeded in the signal sample. This gives the model information about the current state of the system so that the state prediction made using the operating parameters alone can be updated to avoid false positives and false negatives.

Welch's Method ($FFT_W(P')$) is a spectral density estimation method that windows the signal into overlapping segments and applies the Fast Fourier Transform (FFT) before averaging the result [19]. The resulting periodogram contains a lower frequency resolution than the standard FFT, but is more robust to noise due to the averaging over multiple windows. This transformation gives indications of periodic behaviour appearing across the spectrum which is indicative of ordered behaviour that precedes an instability. It has previously been applied in the context of thermoacoustics by Sengupta et al. [20] where it was shown that the frequency spectrum of a thermoacoustic system varies across its operating envelope and can be used to estimate the current operating state of the system. In this study a Hann window of length 256 was used, which resulted in a frequency resolution of approximately 100Hz.

Detrended Fluctuation Analysis ($DFA(P')$) is a transformation used in dynamical systems analysis to estimate

the Hurst exponent by calculating the average integral of the systems over different time scales after the fluctuating signal has been integrated and detrended.

$$DFA(\tau) = \frac{1}{M} \sum_{m=0}^M \sum_{t=m\tau}^{(m+1)\tau} P'_{detrend}(t) \quad (3)$$

where τ is a length scale, M is the number of times τ fits into the signal segment and $P'_{detrend}$ is the fluctuating signal that has been integrated to create a random walk and then detrended. It was first applied to thermoacoustic systems by Gotoda et al. who used it to detect intermittency by estimating the amount of periodic behaviour in the signal. More recently, McCartney et al. showed that using the entire output of the $DFA(P')$ transformation is more useful in predicting the onset of instability as it also carries information about the amplitude of the signal and shorter term correlations, which can also respond to precursors of instability [15]. In this study 30 values were used for τ , log spaced between 10 and 1000 and the signal was detrended using quadratic detrending.

The data samples were labelled according to whether the dynamic pressure signal exceeded a peak-to-peak pressure threshold, set relative to the P2P pressure seen during stable combustion. The labels were then shifted forwards in steps of 100ms up to 1s, so that the label for a sample taken at time t corresponds to the state of the system Δt_{pred} seconds in the future. The Δt_{pred} and future operating parameters in the input vector were also shifted forwards to match the time of the label. For both cases, three data sets were generated: a training data set; a validation data set for tuning the hyperparameters and number of epochs; and a testing data set, for comparing the performance of the different algorithms. In the Annular Rig the training data set contained three 20s and 60s ramps from the 4kW and 8kW power condition, the validation set contained the remaining data from the 4kW and 8kW power conditions and the test data set contained data from the 6kW power condition. For the FPCS, the systems were randomly allocated into three groups, from which each data set was generated.

METHOD

ANNULAR RIG

The objective of the tests in the Annular Rig case is to show that using information from the dynamic pressure measurement allows the system to recover the information (hence predictive ability) that is lost by removing the wall temperature sensors. Since the wall temperature differences arise from the difference in ramp rates, the future operating parameters were omitted from the input and just the Δt_{pred} feature was kept to define the point to which the label corresponds to. Three sets of models were

trained on different input features: 1) only the future operating parameters (i.e. \dot{m} and burner power), which represents the minimum information model; 2) the current operating parameters and the wall temperatures, T_w , which represents a system with all the information required to predict the instability onset; 3) the current operating parameter and $DFA(P')$. After training and tuning on the train and validation data sets, the log-likelihood of the data given the predicted probabilities is evaluated for predictions at different times in the future, Δt_{pred} , and used to compare the models. The log-likelihood (LL), given in Eqn. 4, is a statistical measure of the quality of the fit and is calculated by estimating the likelihood of observing the data, given the mean probability predicted by the BNN.

$$LL = \frac{1}{N} \sum_{i=1}^N \hat{y}_i \cdot \log(P(y_i)) - (1 - \hat{y}_i) \cdot \log(1 - P(y_i)) \quad (4)$$

where \hat{y}_i is the observed state for sample i and $P(y_i)$ is the predicted probability of sample i exhibiting instability. The absolute value depends on the dataset and so is not important, but the relative values can be used to compare the model predictions. The models are also evaluated in terms of the difference in the time, t , between the predicted and the actual onset of the instability for each run. The error is calculated by generating an ensemble of predictions of the probability of instability, from a given point in the test data, at all times in the future along the planned trajectory through the operating space. For each individual system run, the value of t where the mean probability predicted by the ensemble exceeds the classification threshold (the setting of which is detailed earlier) is taken and compared with the value of t when the system first encountered an instability on that run. The error in the prediction is averaged over all of the test runs and then this is repeated for decreasing time to instability, t_{TTI} (from 1000ms to 100ms)

Full-Scale Prototype Combustion System (FPCS)

The tests in the second case aim to show how the models perform when the hidden variable is due to differences between samples of the same system type, rather than differences in the history of the ramp profile of the system. It also aims to show how the predictions would evolve when using BNNs to inform the controller in a practical scenario. All of the input feature sets contain the future operating parameters and the time to which they correspond. However, since all of the systems carry out the same ramp, this information is redundant in the feature set containing only the operating parameters and the predictions are expected to be independent of Δt_{pred} . Additionally, all tests were conducted at sea-level conditions, meaning that the compressor exit conditions are correlated more strongly with the control vari-

ables than would be seen at multiple altitudes. The log-likelihood across the whole of the test data set and the error in the predicted onset of instability, in terms of ω_{HP} , are used to compare the performance of the different algorithms. As in the Annular case, models are also evaluated in terms of the error in predicting the onset of instability. In this case the error is given in terms of the difference in the corespeed, ω_{HP} , between the predicted and the actual onset and calculated in the same way as the Annular case.

RESULTS

The predicted probabilities were compared for the BNNs trained with different inputs in terms of the quality of the predicted probabilities and their ability to predict the onset of an instability.

ANNULAR RIG

Figure 4A shows the LL of the predicted probabilities for each of the models. It can be seen that the model trained using T_w predicts much likelier probabilities than the OP model, which confirms the hypothesis that the wall temperatures carry a lot of information about the run type and hence when the system is exhibiting instabilities. It can also be seen that the $OP \& DFA(P')$ model performs similarly for timescales over 1 second, and even better at Δt_{pred} smaller than 700ms, showing that the pressure signal can compensate for the missing temperature information.

In Figure 4B the error in the predicted time at which the system will exhibit instabilities compared to the time when the system actually encounters an instability, (t_{error}), is shown for predictions made within 1s of the instability. It can be seen that for the OP and $OP \& T_w$ models, there is little variation in the error of the prediction as the instability is approached. The $OP \& DFA(P')$ model however, shows a significant decrease in the error as the system approaches the instability, reaching an error similar to that of the OP model at 700ms and decreasing to lower than the other models at t_{TTI} smaller than 300ms.

FPCS

The first evaluation on the FPCS data is shown in Figure 5A and compares the prediction across the whole of the test dataset for the different BNNs and across different prediction horizons up to 1s in the future, in terms of the log-likelihood. As can be seen in Figure 5A, the prediction from all models that utilise P' in their predictions deteriorate when they are predicting further into the future, tending towards the performance of using just the operating parameters at a prediction horizon of 500ms. This is consistent with our expectation that the information contained in the pressure signal is relevant over short timescales. The BNN utilising the $DFA(P')$ performs the best over the shortest timescales, approaching the predictive ability of the $OP \& FFT_w(P')$

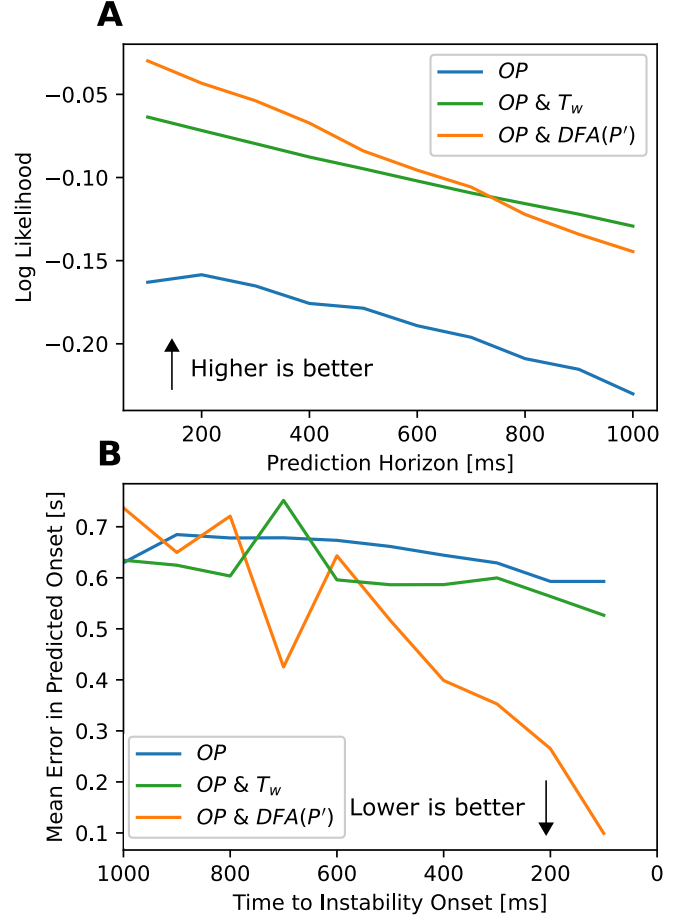


FIGURE 4. A) The log-likelihood for probabilities of instability predicted at different times in the future on the Annular Rig test data set (higher is better) B) The error in the predicted time at which the system will encounter an instability as the system carries out a ramp

and $OP \& insta(P')$ models at a horizon of 300ms. The performances diverge at horizons larger than 500ms with $FFT_w(P')$ and $OP \& DFA(P')$ performing worse than the operating parameters alone.

In the second evaluation on the FPCS data, shown in Figure 5B, the mean absolute error in the value of ω_{HP} when the system is predicted to encounters instability is calculated for decreasing time to instability, which gives an indication of how the predictions evolve as the system approaches an instability, in a way that might be done in a practical scenario. The behaviour is similar to that seen in the previous experiment, with the error for the models trained on $OP \& FFT_w(P')$ and $OP \& DFA(P')$ reducing with decreasing t_{TTI} and being greater than OP at large t_{TTI} . The plot shows that, on average, the $OP \& FFT_w(P')$ and $OP \& DFA(P')$ models achieve a lower error than the OP and $OP \& insta(P')$ models at t_{TTI} smaller than 500ms, decreasing to approximately 45% lower error than the OP model and 25%

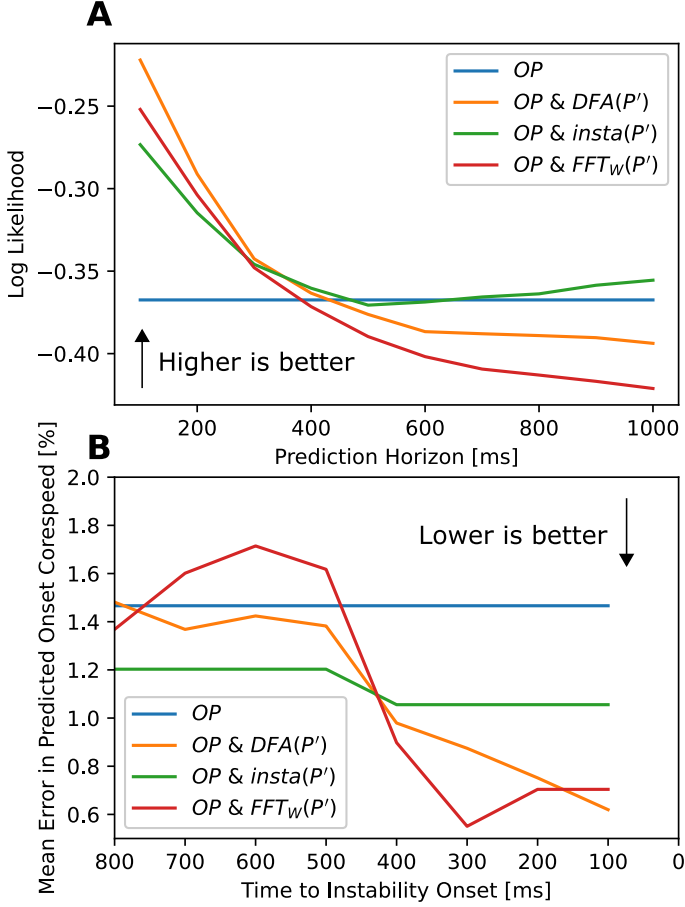


FIGURE 5. A) The log-likelihood for probabilities of instability predicted at different times in the future on the FPCS test data set (higher is better) B) The error in the predicted corespeed where the system will encounter an instability as the system carries out a ramp

lower error than the OP & $insta(P')$ model at times of 200ms away from the instability.

In order to examine the independence of information relevant to prediction that is carried by the different inputs a comparison of models trained with different subsets of features in the FPCS case is shown in Figure 6. Figure 6A shows predictions made with different levels of information about the pressure signal. The dotted red line shows the LL achieved with an average value prediction i.e. predicting all probabilities to be equal to the portion of data points exhibiting instability in the data set. This represents a predictor with no information about the current state (i.e. from OP or P') and shows the lower bound for the LL. The dashed green line shows the LL achieved with a model that only knows if the system is currently exhibiting instability, which represents a model with the minimal amount of P' information. This model converges to an LL less negative than the no information rate at large Δt_{pred} due to the correlation between

the current and future states and becomes less negative at shorter Δt_{pred} through avoiding false positives/negatives and the increasing correlations. The orange dashed line shows the performance of a model trained with only the information from $DFA(P')$ and the orange solid line performance of a model trained with OP & $DFA(P')$. It can be seen the $DFA(P')$ model converges towards the minimum P' information model ($insta(P')$) at large Δt_{pred} but LL is less negative at smaller Δt_{pred} due to it having more information relevant to the prediction of stability over short timescales. The OP & $DFA(P')$ model follows the same trend as the $DFA(P')$ model, achieving the same LL at short Δt_{pred} . It converges, however, to a less negative LL due to the information provided by the OP variables which, when used alone, provide a better prediction as was seen in Figure 5.

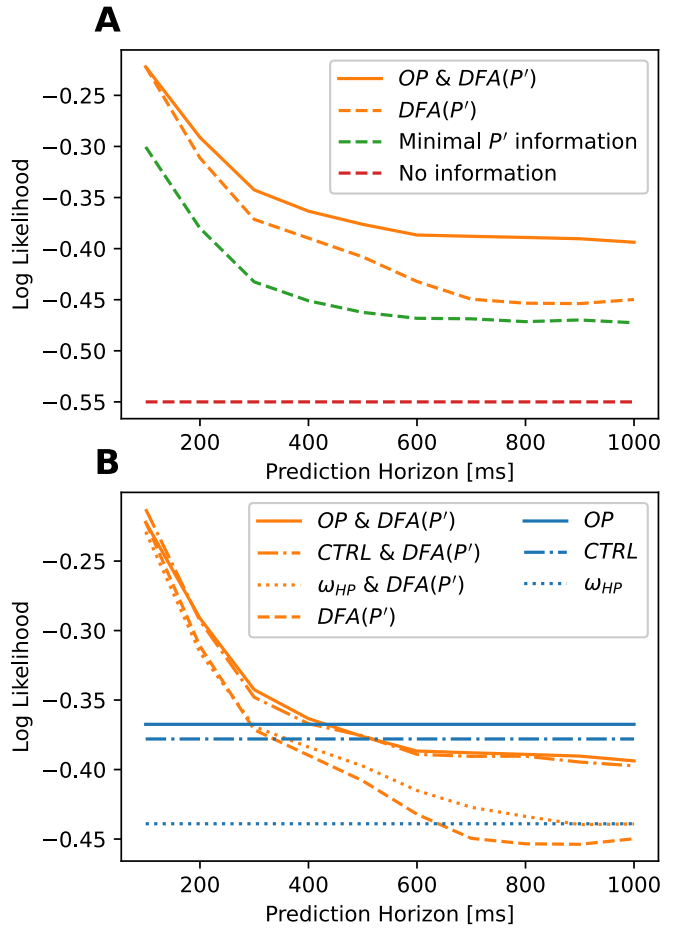


FIGURE 6. A) The log-likelihood for probabilities of instability predicted at different times in the future on the FPCS test data set for different amounts of pressure information B) The log-likelihood of predictions made with varying low frequency data sources: OP contains all LF features, $cont.$ contains fuel flow split, fuel flow rate and core speed and ω_{HP} contains only the core speed.

Figure 6B shows the relevance of the information provided by the different features in the OP variables and how they affect the predictions. The OP features are split into two subsets of low frequency (LF) variables: $CTRL$, which removes the measured variables (compressor exit and pressure and FAR) leaving just the system control inputs (pilot/secondary fuel split and fuel flow rate) and the control target variable, ω_{HP} ; and ω_{HP} alone. Models trained on just these features are shown in blue and models trained with these features combined with $DFA(P')$ are shown in orange. It can be seen that there is a small increase in LL for the LF models when the measured variables are removed; however, the impact on the $DFA(P')$ model performance is negligible. This suggests that the information contained in these variables that is relevant for predicting the state is also carried by the $DFA(P')$ features. Keeping only ω_{HP} as an LF feature has a more significant impact on the models, but ω_{HP} only provides a small improvement in LL for the $DFA(P')$ models at large Δt_{pres} . The ω_{HP} & $DFA(P')$ model LL exceeds the ω_{HP} model LL at Δt_{pred} smaller than 800ms, achieving the same LL as the OP and $DFA(P')$ model at a Δt_{pred} of 300ms, below which it performs better. This shows that the majority of the relevant information carried by the OP features is also carried by the $DFA(P')$ features at 300ms, although the difference between the ω_{HP} & $DFA(P')$ and OP & $DFA(P')$ at his Δt_{pred} show that the other control variables do contain some independent information.

DISCUSSION

For both cases presented above, there is a clear trend in the accuracy of predictions made by BNNs leveraging the dynamic pressure information, with respect to Δt_{pred} . This trend is consistent with our expectations that the pressure information provides a good description of the current state of the combustor, but is only useful over short timescales. Over larger timescales the high dimensional nature of the pressure signal and compounding stochastic effects make the pressure signal less informative of the system state. This is most obvious in the FPCS case, where the predictive performance of the measures utilising the pressure signal tends towards the OP model performance as the prediction horizon approaches 500ms into the future. Beyond 500ms, the performance deteriorates beyond the OP model performance as, at these time scales, the pressure signal carries no information about the instability and so the inputs from the features act as additional noise to the BNN. While the BNN will learn to attach a lower weight to the pressure signal inputs when Δt_{pred} is large, the regularisation of the NNs, described earlier, penalises NNs for discounting them completely, as the prior distribution weights all inputs equally. Furthermore, Pearce et al. observed that BNNs trained using the randomised MAP sampling approach performed worse than other methods when the noise in the data was large [17]. This means that a larger number of input features that are calculated from the pressure signal will

result in a worse performance compared to the operating parameter BNN at large Δt_{pred} . This effect diminishes with increasing training data set size and is driven by the selection of the prior distributions. It is possible to reduce the effect by using a prior that embeds our domain knowledge that the pressure signal carries no extra information at large timescales, by increasing the amount of training data, which will reduce the impact of the prior on the posterior predictions or by using a different type of BNN. However this was not investigated in this study.

The performance of the BNNs using P' relative to the other models show that, over short timescales (lower than 500ms), the pressure signal can be used as an additional source of information to make up for missing information from unmeasured (in the Annular Rig case) and unmeasurable (in the FPCS case) hidden variables. The performance of the OP & $insta(P')$ model shows that this ability partly comes from the knowledge of the stability of the combustor, which allows the model to avoid making false positive (predict exhibiting instability when stable) or false negative (predict stable when exhibiting instability) predictions. Despite the similarity of the performance of the OP & $insta(P')$ and OP & $FFT_W(P')$ models in terms of LL, the OP & $FFT_W(P')$ model performs better at short t_{TTI} in the onset prediction. In the onset prediction, the OP & $insta(P')$ model benefits only from its ability to avoid false positives made using OP information, whereas OP & $FFT_W(P')$ is also able to recognise changes in the frequency spectrum of P' and use this to inform its predictions. At short timescales, the OP & $DFA(P')$ model performs the best in terms of LL and also performs better than OP & $insta(P')$ in the prediction of the instability onset. The OP & $DFA(P')$ model responds to periodic behaviour in the signal, which is also detected by the OP & $FFT_W(P')$ model, as well as changes in the short term correlations of the signal. Additional testing would be required to determine if the better LL performance was due to the model leveraging different information or simply being able to find a better solution given the data. Furthermore, the plots in Figure 6 demonstrate that a significant amount of the information in the OP features relevant for predicting the system state is also contained in the pressure information at smaller Δt_{pred} , although the relative performance shows that there is some independence in the information and so the pressure information should be thought of as an extra information source rather than something that can act as a replacement for unmeasured variables.

Given the minor differences in performance between the OP & $FFT_W(P')$ and OP & $DFA(P')$ models, and the fact that the behaviour of P' before the onset of an instability is particular to that system and instability, one model will not outperform the other in general. So, when creating a model for predicting instabilities, many transformations should be tested before using the best method for the given system.

The training time on a 16 core Intel Xeon E5-2620 CPU varied between 5 and 90 minutes depending on the number of

features used and the BNN architecture when trained in parallel. The inference time of each NN in the ensemble varied between 1-25ms depending on the number of batches which suggests that, if properly optimised for inference speed, the BNN could be executed in $\mathcal{O}(10\text{ms})$, but this does not include the signal transformation and should be investigated in further work. The computation time for predicting using the BNN and the timescales over which the predictions using P' give lower errors ($\lesssim 500\text{ms}$) suggest that there is potential for using this type of model for ‘in the loop’ engine control. Improvements over larger timescales could potentially be achieved using different or combinations of different transformations of P' and using methods of inference or prior selection in the BNNs.

BNN prediction up to 1.4s in the future is shown. The standard deviation of the prediction at the onset is 0.04. In orange, two predictions that simulate out-of-distribution (OOD) predictions (i.e. predictions made outside of the training domain) are shown. Figure 7A shows a prediction along a trajectory that takes the system to a combination of operating parameters outside of those seen during testing and Figure 7B shows a predictions where broadband noise has been added to the dynamic pressure signal, so that the signal transformed with $FFT_W(P')$ has different characteristics to those observed during testing. In both cases the uncertainty in the OOD predictions is significantly larger (0.11 at onset for A and 0.17 at onset for B) than the ‘in distribution’ case (0.04).

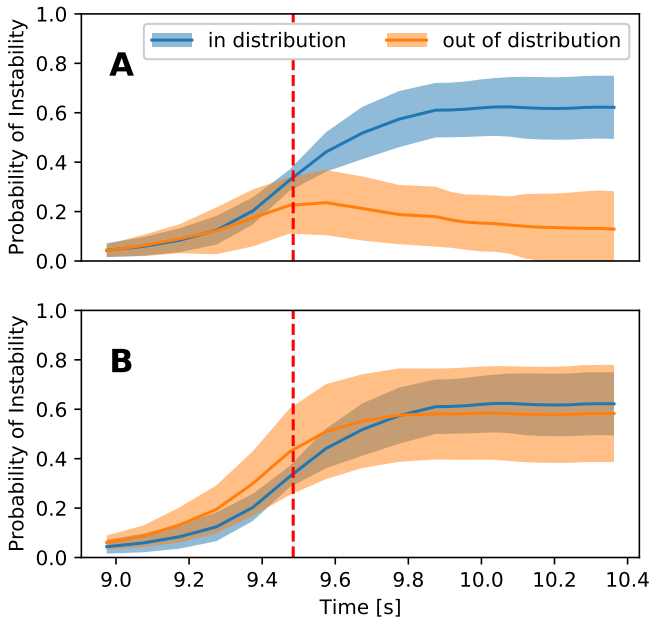


FIGURE 7. The predicted probability of an FPCS encountering an instability for an example run using the $OP \& FFT_W(P')$ model, with the prediction made 500ms before the onset of instability (dashed red line). $\bar{y}_{pred} \pm \sigma_{pred}$ is shown in blue for an example run from the test dataset and in orange for A) a trajectory not in the training data and B) broadband noise added to the pressure signal.

An additional property of BNNs that is not leveraged in this study is the confidence that is provided with their predictions. This is useful for detecting when the BNN is making predictions that are not well supported by training data, and so should not be trusted. To demonstrate this, Figure 7 shows predictions made for a run from the FPCS case using the $OP \& FFT_W(P')$ model, when the system is 500 ms away from onset of instability (red dashed line). In blue, the mean and standard deviation of the

The uncertainty provided by the Bayesian framework can be used in two ways: to monitor how a system is operating relative to the training domain to predict when maintenance may be required and in control of the system. The best way to leverage uncertainty in control is an ongoing area of research, largely taking place in the context of medical diagnoses, where it has been shown that referring to experts when the model uncertainty is high significantly improves overall performance [21] and reinforcement learning, where the uncertainty is used to indicate areas that require exploration [22]. The model discussed in this study estimates the probability of an instability occurring in the near future and is designed to act as an input to the control system. This input should work in tandem with existing control logic, which keeps the system operating in regions where there is zero probability of an instability occurring. Where, during operation, a more favourable trajectory lies outside of the normally-allowed operating envelope, the BNN should be used to decide whether to leave the normally-allowed operating space and move along that trajectory. The controller should decide to follow that trajectory if two criteria are met: the probability of instability is below acceptable limits and the uncertainty indicates that the BNN is not predicting ‘out of distribution’. The first criteria will depend on the severity of the instability (e.g. an instability with a low limit cycle amplitude might cause excessive degradation if it were to continue for a long time but poses no extra risk if present for a short duration, whereas an instability with an extremely large limit cycle amplitude could cause components failure) and so should be set on a case by case basis. The second criteria will depend on the data, the model used and the settings of the prior and therefore should be set based on the uncertainty observed in test predictions. The controller must also be able to return to ‘safety’ in the event that the uncertainty increases beyond allowed limits once the controller has already started following an alternative trajectory. All of these items will require further investigation of tests with the BNN framework operating as part of the control loop.

CONCLUSION

We have demonstrated in this study that Bayesian Neural Networks can be used to combine information from the dynamic pressure measurements from the combustor, with system operating parameters in order to reduce uncertainty, caused by hidden variables, in the point at which an instability will occur. We demonstrated this first using data from an annular, atmospheric test rig where a hidden variable was created by removing information from wall temperature measurements, which significantly reduces the ability of the BNNs to predict the probability of the system exhibiting instability. Information from the dynamic pressure was then added as an input to the BNN and it was shown that this allowed the BNN to recover the predictive performance over short timescales. We also showed, on a full-scale prototype combustion system at realistic operating conditions, that the error in predicting at which point in the operating space an instability will occur reduces when the dynamic pressure is considered and the system is less than 500ms from an instability. We found that at times larger than this, the pressure information acted as a source of noise to the model which leads to poorer predictions, although this effect could be reduced with a more advanced method for selecting the prior distribution of the BNN parameters. We found that training on the DFA transformation of the signal gave the best predictions in terms of the log-likelihood and at short timescales its error in prediction of the onset of the instability was matched by models trained on a signal transformed with Welch's method. These models achieved a 30-45% reduction in the error of predicted onset of the instability in terms of the co-speed, at times 400-200ms before the instability compared to predictions made using the operating parameters alone.

These results were created using a novel framework for evaluating the probability of an instability occurring at different times in the future along multiple potential trajectories, along with a confidence in the prediction that indicates whether the prediction is well supported by training data or is extrapolating and provided suggestions for future work on leveraging this effectively. This study has shown that Bayesian Neural Networks trained using the proposed framework provide a promising approach for incorporating dynamic pressure information and precursors to combustion instability into engine control logic, for expansion of the operating envelope and more confident operation of the engine close to instability.

ACKNOWLEDGMENT

This project has received funding through the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 766264.

REFERENCES

- [1] Smith, K., and Blust, J., 2005. "Combustion instabilities in industrial gas turbines: Solar turbines' experience". In *Combustion Instabilities In Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modeling*. pp. 29–41.
- [2] Foust, M., Thomsen, D., Stickles, R., Cooper, C., and Dodds, W., 2012. "Development of the GE Aviation Low Emissions TAPS Combustor for Next Generation Aircraft Engines". In 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition.
- [3] Juniper, M. P., and Sujith, R., 2018. "Sensitivity and non-linearity of thermoacoustic oscillations". *Annual Review of Fluid Mechanics*, **50**(1), pp. 661–689.
- [4] Dakos, V., Carpenter, S. R., Brock, W. A., Ellison, A. M., Guttal, V., Ives, A. R., Kéfi, S., Livina, V., Seekell, D. A., van Nes, E. H., and Scheffer, M., 2012. "Methods for Detecting Early Warnings of Critical Transitions in Time Series Illustrated Using Simulated Ecological Data". *PLoS ONE*, **7**(7), jul, p. e41010.
- [5] Gotoda, H., Miyano, T., and Shepherd, I. G., 2010. "Dynamic properties of unstable motion of swirling premixed flames generated by a change in gravitational orientation". *Phys. Rev. E*, **81**, Feb, p. 026211.
- [6] Gotoda, H., Ikawa, T., Maki, K., and Miyano, T., 2012. "Short-term prediction of dynamical behavior of flame front instability induced by radiative heat loss". *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **22**(3), p. 033106.
- [7] Sarkar, S., Chakravarthy, S. R., Ramanan, V., and Ray, A., 2016. "Dynamic data-driven prediction of instability in a swirl-stabilized combustor". *International Journal of Spray and Combustion Dynamics*, **8**(4), pp. 235–253.
- [8] Murugesan, M., and Sujith, R. I., 2016. "Detecting the onset of an impending thermoacoustic instability using complex networks". *Journal of Propulsion and Power*, **32**(3), pp. 707–712.
- [9] Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., Van Nes, E. H., Rieterkerk, M., and Sugihara, G., 2009. "Early-warning signals for critical transitions". *Nature*, **461**, pp. 53–59.
- [10] Nair, V., and Sujith, R., 2014. "Multifractality in combustion noise: predicting an impending combustion instability". *Journal of Fluid Mechanics*, **747**, p. 635–655.
- [11] Gotoda, H., Nikimoto, H., Miyano, T., and Tachibana, S., 2011. "Dynamic properties of combustion instability in a lean premixed gas-turbine combustor". *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **21**(1), p. 013124.
- [12] Jha, D. K., Virani, N., and Ray, A., 2018. "Markov modeling of time series via spectral analysis for detection of combustion instabilities". In *Handbook of Dynamic Data*

Driven Applications Systems. Springer International Publishing, pp. 123–138.

- [13] Mondal, S., Ghalyan, N. F., Ray, A., and Mukhopadhyay, A., 2019. “Early detection of thermoacoustic instabilities using hidden markov models”. *Combustion Science and Technology*, **191**(8), pp. 1309–1336.
- [14] Hachijo, T., Masuda, S., Kurosaka, T., and Gotoda, H., 2019. “Early detection of thermoacoustic combustion oscillations using a methodology combining statistical complexity and machine learning”. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **29**(10), p. 103123.
- [15] McCartney, M., Indlekofer, T., and Polifke, W., 2020. “Online prediction of combustion instabilities using machine learning”. In Proceedings of the ASME Turbo Expo 2020: Turbomachinery Technical Conference and Exposition. Volume 4A: Combustion, Fuels, and Emissions., ASME.
- [16] Indlekofer, T., Faure-Beaulieu, A., Noiray, N., and Dawson, J., 2021. “The effect of dynamic operating conditions on the thermoacoustic response of hydrogen rich flames in an annular combustor”. *Combustion and Flame*, **223**, pp. 284 – 294.
- [17] Pearce, T., Leibfried, F., and Brintrup, A., 2020. “Uncertainty in neural networks: Approximately bayesian ensembling”. In Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, S. Chappa and R. Calandra, eds., Vol. 108 of *Proceedings of Machine Learning Research*, PMLR, pp. 234–244.
- [18] Goodfellow, I., Bengio, Y., and Courville, A., 2016. *Deep Learning*. MIT Press, ch. 8-11.
- [19] Welch, P., 1967. “The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms”. *IEEE Transactions on Audio and Electroacoustics*, **15**(2), pp. 70–73.
- [20] Sengupta, U., Rasmussen, C., and Juniper, M., 2021. “Bayesian machine learning for the prognosis of combustion instabilities from noise”. *Journal of Engineering for Gas Turbines and Power*.
- [21] Leibig, C., Allken, V., Ayhan, M. S., Berens, P., and Wahl, S., 2017. “Leveraging uncertainty information from deep neural networks for disease detection”. *Scientific Reports*.
- [22] Deisenroth, M. P., and Rasmussen, C. E., 2011. “Pilco: A model-based and data-efficient approach to policy search”. In Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11, Omnipress, p. 465–472.

APPENDIX

A BNN ARCHITECTURES

Tables 1 and 2 shows the selected hyperparameters for all of the models presented in this study. Also shown are the average

log-likelihoods for the training and tuning data sets. All BNNs contained an ensemble of 12 networks and the hyperparameters were selected after a random search over 50 trials with the number of hidden layers varied between 2 and 10 and the number of neurons in each layer varied between 10 and 100.

Model	# Hidden Layers	# Neurons per layer	Train LL	Tune LL
OP	10	50	-0.169	-0.426
$OP \& DFA(P')$	5	50	-0.227	-0.397
$OP \& insta(P')$	10	15	-0.192	-0.423
$OP \& FFT_w(P')$	6	55	-0.203	-0.422

TABLE 1. Hyperparameters for models trained on Annular Rig data

Model	# Hidden Layers	# Neurons per layer	Train LL	Tune LL
OP	3	80	-0.193	-0.188
$OP \& T_w$	8	34	-0.064	-0.121
$OP \& DFA(P')$	9	52	-0.082	-0.088

TABLE 2. Hyperparameters for models trained on FPCS data