

Probabilistic machine learning for the elimination of thermoacoustic instabilities



Ushnish Sengupta

Supervisor: Professor Carl Rasmussen

Advisor: Professor Matthew Juniper

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Clare College

September 2022

Dedicated to my parents.

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this thesis are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This report is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This thesis contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations.

Ushnish Sengupta
September 2022

Acknowledgements

I'm deeply grateful to my advisor Professor Matthew Juniper, who, for the last four years, made sure to meet me every week and received the support that I needed. Whether I got stuck on a specific technical issue, needed to discuss the direction or significance of our work or wanted help with mundane administrative matters, Matthew was always there. Research is typically a task with sparse and underspecified rewards and it is hard to overstate the importance of his regular feedback and enthusiasm. I would also like to thank my supervisor Professor Carl Rasmussen for his encouragement and advice. I'd like to acknowledge the Hopkinson lab technicians Roy Slater and Mark Garner because it would have been impossible to set up the experiments in Chapters 3 and 5 without their aid.

I am also thankful for the funding I received from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement number 766264.

Much of the work presented in this thesis is an outcome of collaborations. The parameter inference project described in Chapter 3 was devised together with Max Croci who worked on the flame image processing, G-equation simulations and ensemble Kalman filter while I performed the Bunsen flame experiments, wrote the Bayesian neural network code and ran the neural network inference using Max's processed data. Chapter 6 used data from the BKD rocket engine thrust chamber collected by experimental collaborators in DLR Lampoldshausen, among whom Wolfgang Armbruster, Jan Martin, Justin Hardi and Günther Waxenegger-Wilfing deserve special mention for kindly hosting me in Lampoldshausen, arranging the data transfer, helping me understand the data and providing their scientific input. Chapter 7 also used data from colleagues in Rolls-Royce aircraft engines and we are grateful for their generosity.

I would also like to thank my division A/F workmates, in particular, Francesco, Yubiao and Matthew, for the many discussions, topically relevant or irrelevant, that we've had. Their company has always been a source of joy and inspiration. A special acknowledgement goes out to my remote friends Ankit and Manosij, who were always ready to listen to my nonsense. And finally, a well-deserved thank you to my mom and dad for instilling a love for science and technology in me from a very young age and supporting me through thick and thin.

Abstract

Thermoacoustic instabilities have hindered the development of high energy density combustors, such as those in jet engines, gas turbines or rockets, for decades. They are notoriously difficult to model and despite efforts to eliminate them, can show up unexpectedly. The current study demonstrates how Bayesian machine learning techniques may be leveraged to exploit noise radiation, and combustor sensor data in general, for forecasting such instabilities. We use Bayesian neural networks to learn a functional relationship between sensor data and measures of thermoacoustic stability or operating parameters. First, we demonstrate on a laboratory-scale Rijke tube driven by a turbulent flame that it is possible to predict decay rates of acoustic pulses from the spectrum of 100 millisecond combustion noise samples, thus enabling us to monitor the combustor's thermoacoustic stability in real time. We then apply these ideas to predict instabilities in an experimental rocket chamber, where we use the history of sensor data, including measurements of noise (dynamic pressure), temperature, static pressure, etc. to predict instabilities upto 250 milliseconds before they occur. The Bayesian nature of our algorithms allows principled estimates of uncertainty to accompany each prediction while the technique of Integrated Gradients lets us interpret our models.

It is hoped that this study will serve as a first step towards establishing Bayesian machine learning techniques as a tool to help us discover interesting relationships in combustor data and build trustworthy, robust and reliable combustion prognostics or sensor validation.

Table of contents

List of figures	xiii
1 Introduction	1
1.1 The problem of thermoacoustic instabilities	1
1.2 Mitigation strategies for thermoacoustic instabilities	3
1.3 Opportunities for Probabilistic Machine Learning	4
2 Bayesian machine learning	7
2.1 Introduction	7
2.2 Uncertainty Quantification in Machine Learning Models	8
2.3 Bayesian neural networks	9
2.4 Neural network interpretation	10
2.5 Gaussian Process Modelling	11
2.6 Global optimization using Gaussian Processes	13
3 Parameter inference in reduced order flame models	17
3.1 Introduction	17
3.2 Bunsen flame experiment and simulations	18
3.2.1 Bunsen flame experiment	18
3.2.2 Flame front model	19
3.2.3 Simulated flame front library	20
3.3 Inference using heteroscedastic Bayesian neural network ensembles	20
3.4 Inference using the Ensemble Kalman Filter	22
3.5 Results	23
4 Global optimization of combustor models using Bayesian optimization and adjoint methods	27
4.1 Introduction	27
4.2 Methods	29

4.2.1	Thermoacoustic models	29
4.2.2	Gradient-augmented Bayesian optimization	31
4.3	Results	31
5	Prognosing instabilities from noise in a Rijke burner	35
5.1	Introduction	35
5.1.1	Combustion noise as a diagnostic	35
5.1.2	The case for interpretable, Bayesian machine learning	37
5.2	Experimental Setup	38
5.3	Statistical tools	42
5.3.1	Precursors of thermoacoustic instability from the literature	42
5.3.2	Bayesian neural network ensembles	43
5.3.3	Interpretation using Integrated Gradients and Permutation Importance	43
5.4	Results	44
6	Prognosing instabilities in a rocket engine combustor	51
6.1	Introduction	51
6.2	Experimental Setup	52
6.3	Methods	55
6.3.1	Nonlinear autoregressive models with exogenous variables (NARX) for timeseries modeling	55
6.3.2	Interpretation using Integrated Gradients	56
6.3.3	Transforming sensor signals into neural network inputs	56
6.3.4	Train-test-validation split, hyperparameters, performance metrics .	57
6.4	Results	58
7	Prognosing instabilities in an intermediate pressure industrial fuel spray nozzle	65
7.1	Introduction	65
7.2	Results	66
8	Conclusions and Outlook	69
8.1	Summary of contributions	69
8.2	Outlook	71
	References	75

List of figures

2.1	Approximate Bayesian inference using ensembles of neural networks. . . .	10
2.2	A 1D Gaussian Process metamodel with and without derivatives. Including derivative information results in a GP which has a higher confidence in its predictions. The gray area represents 3 s.d. uncertainty bounds in the Gaussian process metamodel, the red crosses are function evaluations, the red lines denote function and derivative evaluations and the dotted blue lines are randomly sampled functions from the Gaussian process metamodel. Figure adapted from Williams and Rasmussen [1996]	14
3.1	(<i>Left</i>) Image of a Bunsen flame. (<i>Middle</i>) <i>G</i> -equation model of the flame front. (<i>Right</i>) Simulated flame in LSGEN2D.	20
3.2	Architecture of a single neural network. The input and hidden layers have 900 nodes each, while the output layers have 6 nodes each. All layers are fully connected (FC). Rectified Linear Unit (ReLU) activation functions are used for the hidden layers and sigmoid and exponential (Exp) activation functions are used for the mean and variance output layers respectively. . . .	22
3.3	Scatter plots of true parameters (scaled) versus predicted parameter values for simulated test data.	23
3.4	Results of inference using the BayNNE method compared to the EnKF method for the simulation of two different Bunsen flames. The flame images are preprocessed to find the flame front and then used for inference. The green and orange bands in the EnKF and BayNNE predicted flame shape plots are regions of high likelihood. The normalised area variations over one period of the simulated flames show good agreement with the experiments.	25

4.1	Contour plots of non-dimensionalized growth rate and frequency of the Rijke tube fundamental mode plotted against heater position and iris diameter, reproduced from Juniper [2018]. The numbered black contour lines correspond to frequencies and the colored contour plot represents the growth rate. White arrows indicate the direction and magnitude of the growth rate gradients, as computed by the adjoint model.	30
4.2	Geometric parameters of the network model from Aguilar Perez [2019] used in the optimization routine. The parameters in black are allowed to be updated. The parameters in red are kept constant.	30
4.3	Surrogate Gaussian process model of the Rijke tube cost function at different iterations. Black dots indicate where the Bayesian optimization has evaluated the function and its gradients. The bottom right plot is a contour plot of the true cost function, with the black dot indicating the location of the true global minimum.	32
4.4	Best objective function values achieved by the BFGS, BO without gradient and gradient-augmented BO routines in the longitudinal combustor test case, averaged across 10 trials, plotted against the number of cost function evaluations. Error bars indicate 1 s.d. variation across trials.	33
4.5	A stable combustor configuration, shown in red, found by a gradient-augmented Bayesian optimization run after 50 evaluations. Original unstable configuration in black.	34
5.1	Schematic of experimental setup, consisting of a 1 KW turbulent swirl flame inside a steel tube of length 800 mm and internal diameter 80 mm	39
5.2	Left: Turbulent combustor excited by an acoustic pulse. Right: Filter signal with best-fit line in red.	39
5.3	Equivalence ratios and flow rates of datapoints for methane:ethene ratio of 3:4 and outlet diameter 80 mm	40
5.4	Contour plot of decay timescale (seconds) for methane:ethene ratio of 3:4 and outlet diameter 80 mm	41
5.5	Plot of measured decay timescales vs decay timescales predicted by the neural network ensemble on the test data	45
5.6	Plot of measured equivalence ratio vs equivalence ratio predicted by the neural network ensemble on the test data	46
5.7	Plot of measured burner power vs burner power predicted by the neural network ensemble on the test data	46

5.8	Top: Min-max normalized power spectrum of a noise sample corresponding to a 0.43 second decay timescale (close to instability) Bottom: Integrated gradients attribution plot for this power spectrum	48
5.9	Permutation feature importance for decay timescale prediction	49
5.10	Plot of generalized hurst exponent H vs decay timescale	49
5.11	Plot of autocorrelation decay vs decay timescale	50
6.1	Experimental thrust chamber BKD.	53
6.2	BKD test sequence and p' spectrogram showing self-excited instability of the first tangential mode [Gröning et al., 2016]. p_{cc} is the chamber static pressure, ROF is the mixture ratio of oxidizer to fuel, and T_{H_2}, T_{O_2} are the LOX and hydrogen injection temperatures respectively.	54
6.3	Forecast p'_{rms} values for the four experimental runs based on sensor data prior to $t - 0.5$ sec, superimposed on observed p'_{rms} values at time t . The model was trained on the other three runs in each case. This model used p'_{FFT} and p'_{rms} as state variables and order $h = 2$. Actual observations are shown as black dots, 3 standard deviation (S.D.) total uncertainty bounds for the forecast are in light blue and 3 standard deviation (S.D.) epistemic uncertainty bounds are in deep blue.	60
6.4	Bayesian Neural Network epistemic uncertainty (top) and p'_{rms} forecasts (bottom) for experimental run no. 4 where a fibre optical probe failure was simulated between $t = 10$ sec and $t = 40$ sec (marked by two grey dotted lines). The model used I'_{FFT}, p'_{rms} as input state variables and order = 0. Actual observations are shown as black dots, 3 standard deviation (S.D.) total uncertainty bounds for the forecast are represented by light blue bars and 3 standard deviation (S.D.) epistemic uncertainty bounds are in deep blue. The light blue bars are barely discernible from the deep blue bars in the failure region because the total uncertainty is almost entirely dominated by the large epistemic uncertainty.	61

-
- 6.5 Integrated gradient attributions and input feature values (normalized to have zero mean and unit standard deviation) for two datapoints from experimental run no. 4 at $t = 16.0$ seconds (top, immediately preceding an instability) and $t = 60.0$ seconds (bottom, stable). The model used p'_{FFT} and p'_{rms} as state variables and $\text{order} = 0$. The blue dots show the normalized values of the features for this particular datapoint and the bars represent the integrated gradient attributions for each feature. For example, feature 6 (spectral power around the 16.66 kHz frequency) in the top plot has a large negative value and a positive integrated gradients attribution, meaning that the network has learned that low power in this frequency band is a potential indicator of instability. 62
- 7.1 The black line shows the thermoacoustic instability threshold as a function of air-fuel ratio (AFR) and exit temperature T_{30} for three aeroplane engine fuel injectors in an intermediate pressure rig. The coloured lines show the distance to the black line. Injectors $1a$ and $1b$ are nominally identical. . . . 66
- 7.2 The black line shows the thermoacoustic instability threshold as a function of air-fuel ratio (AFR) and exit temperature T_{30} for three aeroplane engine fuel injectors (a) $1a$ (b) $1b$ and (c) 2 in the intermediate pressure rig. The coloured lines show the distance to the black line. Injectors $1a$ and $1b$ are nominally identical. 67
- 7.3 Predicted distance to the instability threshold as a function of measured distance to instability threshold for (a) injector $1a$ (b) injector $1b$ (c) injector 2 when the prediction is obtained from a BayNNE trained on injector $1a$. Injectors $1a$ and $1b$ are nominally identical. 68

Chapter 1

Introduction

1.1 The problem of thermoacoustic instabilities

Thermoacoustic instabilities are a potentially catastrophic phenomenon arising in combustion chambers from the coupling between unsteady heat release rate and acoustic waves. There is no alternative to combustion when it comes to long distance flight or space travel, because of the unmatched power to weight ratios that jet and rocket engines possess. But it is due to this very high energy density that designers of rocket thrust chambers, jet engines or gas turbines remain worried about thermoacoustic instabilities. Even if 0.1% of the energy within the chamber is dissipated via the thermoacoustic mechanism, it would result in unacceptably large pressure oscillations. These can cause flame blow-off, extensive structural damage, high rates of heat transfer and component failure.

The earliest reported observation of this phenomenon by Higgins dates back to 1777, while the first qualitative explanation was proposed decades later by Rayleigh [1878], who realized that the driving mechanism of these oscillations is similar to that of a piston engine. In thermoacoustics, a continuous flame behaves like the periodically ignited gas while an acoustic wave is the piston-analogue which periodically perturbs the flame. Acoustic velocity fluctuations perturb the base of the flame creating ripples that convect downstream and cause heat release rate fluctuations after a certain delay, which in turn gives rise to more acoustic waves (either directly or indirectly via the acceleration of entropy fluctuations). If the two-way interaction is such that more heat is released during moments of high pressure, as occurs in a piston engine, then heat is converted to work per cycle. If this excess work cannot be dissipated via viscous damping or acoustic radiation, it increases the pressure oscillation amplitude, leading to instability. This understanding was mathematically quantified by Chu

[1965], who formulated the Rayleigh criterion for instability as follows:

$$\int_0^T \int_V \frac{\gamma-1}{\gamma\bar{p}} p' q' dv dt > \int_0^T \int_S p' \mathbf{u}' \cdot \mathbf{n} |dS| dt \quad (1.1)$$

where p' and q' , are the perturbations in pressure and the heat release rate, respectively, and T , V and \mathbf{n} are the period of oscillation, the combustor volume and the normal vector to the surfaces represented by S , respectively.

Thermoacoustics made the transition from scientific curiosity to industrial headache during the development of liquid-propellant rocket engines in the 1930s and subsequently in the 1960s. For example, Oefelein and Yang [1993] report that achieving a stable design of the Saturn V Rocket F-1 engine required approximately 2000 full-scale tests to eliminate instabilities, which caused pressure oscillations as large as the mean pressure of the combustion chamber. Then, in the 1950s, they were observed in the afterburners and ramjets used in turbojet engines [Zinn et al., 1997]. More recently, stricter emission regulations have motivated the development of aero-engine and gas turbine combustors that operate in lean, low temperature regimes to prevent the dissociation of nitrogen molecules and reduce the formation of NO_x gases. However, lean operating conditions are also more prone to thermoacoustic instabilities [Keller, 1995] and this has made the elimination of thermoacoustic oscillations even more of a priority for manufacturers.

Despite decades of research on this topic, quantitatively accurate models of this phenomenon eludes researchers even today. The phase lag between pressure and heat release rate fluctuations, which governs the thermoacoustic stability of a system, is exquisitely sensitive to system parameters and dependent on multiple mechanisms with different scaling behaviours—hydrodynamic, acoustic and combustion. This makes accurate computational modeling very challenging. The complex, sensitive dependence of thermoacoustic stability on various system parameters also means that instabilities can recur in the later developmental stages of an engine. It is not uncommon for component tests or even full combustor tests to proceed without a problem but then have instabilities crop up in full engine tests. This makes it difficult for us to predict at the design stage whether an engine will suffer from thermoacoustic instabilities and leads to expensive redesigns. This thesis explores some of the ways in which data-driven tools from probabilistic machine learning can help us go beyond traditional physics-based modeling and allow us to predict and avoid thermoacoustic instabilities better.

1.2 Mitigation strategies for thermoacoustic instabilities

Designers can suppress thermoacoustic oscillations in four ways. The first, and most preferred way is to find design modifications that will stabilise the combustor. Owing to the high sensitivity of thermoacoustic growth rates to system parameters, it is often possible to stabilise unstable modes by making minor modifications to the geometry [Aguilar and Juniper, 2020] or flame behaviour

The second mitigation strategy is tuned passive control, where thermoacoustic oscillations are mitigated by adjusting a passive damping device in response to changing operating conditions Kobayashi et al. [2019a]. For example, the resonant frequencies of Helmholtz resonators can be tuned by changing the cavity volume, neck area or neck length. These devices, however, can be bulky and affect the performance of the aircraft or rocket engine. The range over which their bandwidth can be tuned is also limited.

The third mitigation strategy is avoiding unstable regions in the operating parameter space. This can be achieved through extensive experimentation and setting a margin of safety around dangerous islands of instability. This is not only expensive, but differences in behaviour between experimental rigs and real operating conditions can also force designers to be too conservative. The construction of instability precursors from pressure measurements and optical measurements has therefore attracted considerable attention from the combustion research community [Juniper and Sujith, 2018]. Lieuwen [2005] used the autocorrelation decay of combustion noise, filtered around an acoustic eigenfrequency, to obtain an effective damping coefficient for the corresponding instability mode. Other researchers have used tools from nonlinear time series analysis, which capture the transition from the chaotic behaviour displayed by stable turbulent combustors to the deterministic acoustics during instability, such as Gottwald's 0-1 test [Nair et al., 2013] and the Wayland test for nonlinear determinism [Gotoda et al., 2011] to assess the stability margin of a combustor. Nair et al. [2014] reported that instability is often presaged by intermittent bursts of high-amplitude oscillations and used recurrence quantification analysis (RQA) to detect these. In a later paper, Nair and Sujith [2014] noted that combustion noise tends to lose its multifractality as the system transitions to instability. They proposed the Hurst exponent as an indicator of impending instability. Measures derived from symbolic time series analysis (STSA) [Sarkar et al., 2016] and complex networks [Murugesan and Sujith, 2016] are also able to capture the onset of instability. Recent studies have explored machine learning techniques to learn precursors of instability from data. These promise greater accuracy than physics-based precursors of instability, though at the cost of robustness and generalizability. Hidden Markov models constructed from the output of STSA [Jha et al., 2018] or directly from pressure measurements [Mondal et al., 2018] have been used to classify the state of combustors.

Hachijo et al. [2019] have projected pressure time series onto the entropy-complexity plane and used support vector machines (SVMs) to predict thermoacoustic instability. SVMs were also employed by Kobayashi et al. [2019b], who used them in combination with principal component analysis and ordinal pattern transition networks to build precursors from simultaneous pressure and chemiluminescence measurements. Related work by McCartney et al. [2020] used the detrended fluctuation analysis (DFA) spectrum of the pressure signal as input to a random forest and found that this approach compared favorably to precursors from the literature. A recent study by Gangopadhyay et al. [2021] trained a 3D Convolutional Selective Autoencoder on flame videos to detect instabilities. Transfer learning is also being explored [Mondal et al., 2021] as a means of efficiently transferring knowledge of precursors across machines.

Finally, there is active feedback control [Dowling and Morgans, 2005] where an actuator such as loudspeaker or a fuel control valve is operated so that the perturbations generated destructively interfere with the developing instability. These systems may be operated in an open-loop configuration, e.g. a simple phase shifted strategy, or in a closed loop configuration. Unlike the passive approach, the actuator in such a system must act on a very short, millisecond timescale. Despite their laboratory successes, challenging demands of high power output, reliability and robustness placed on an active feedback control system means that feedback control is not used in rocket engines and aircraft engines.

1.3 Opportunities for Probabilistic Machine Learning

Machine learning techniques seek to capture functional relationships between quantities of interest. Unlike traditional physics-based approaches, they can have a sizable number of parameters and this flexibility allows them to represent the data more faithfully. However, the flexibility also has downsides: a particular dataset can be evidence for an infinite number of different functional relationships and we do not know which one is correct. Bayesian machine learning provides a solution to this problem by defining a probability distributions over these possible functions and estimating the resulting uncertainty in its predictions. A Bayesian model starts with an appropriately vague prior belief about what the output of our model should be and, as we observe more data, we update this belief to obtain progressively tighter posterior distributions in accordance with Bayes' rule. This ensures that the model does not make overconfident predictions from out-of-distribution inputs, which are entirely different from those which it was trained on. These uncertainty estimates are particularly important when using a machine learning model in a critical device such as an aircraft engine. Bayesian machine learning techniques with correctly specified priors can also learn

successfully from smaller amounts of data and are resistant to overfitting [Ghahramani, 2005]. They can also be used in continual learning without catastrophic forgetting [Li et al., 2019], which is particularly important if we want to keep learning from data throughout the operating lifetime of a device. Chapter 2 discusses the Bayesian machine learning tools used in this thesis in greater detail.

The following chapters describe how these tools can augment instability mitigation strategies. We know that for thermoacoustic models to be quantitatively accurate, their parameters and other inputs need to be highly accurate, because of thermoacoustic systems' high sensitivity to changes. Chapter 3 describes how we can use a Bayesian neural network to assimilate flame imaging data from experiments or DNS into a simplified combustion model known as the G-equation [Williams, 1985]. We find that our approach is as effective as previous Kalman filter based approaches [Yu et al., 2019a], but roughly 10^6 times faster.

Once we have a model, we need to find an optimally stable configuration. Gradient-descent based approaches coupled with an adjoint model work well but they can become stuck in local optima. The stability landscape of a combustor design is typically multimodal. Global optimization heuristics such as genetic algorithms [Jones et al., 1998], on the other hand, are inefficient and require too many simulations to be run, which may be unfeasible if the underlying model is expensive. High-dimensional, gradient-augmented Bayesian optimization [Eriksson et al., 2018] offers an effective compromise. In Chapter 4, we show how we can construct a Gaussian process surrogate of the combustor's stability in the design space, use the uncertainty in the surrogate to evaluate the adjoint model strategically and find a global optima within a constrained computational budget.

In Chapter 5, we investigate how Bayesian machine learning can use sensor data to help with the prognosis and avoidance of instabilities. We perform experiments on a toy system, a vertical tube with a swirling turbulent flame, and evaluate its thermoacoustic stability in thousands of different operating conditions using pings from a loudspeaker. We then train a Bayesian neural network to recognize how close the system is to instability from its noise spectrum. We find that this network is able to generalize to unseen operating conditions, provides uncertainty estimates in its predictions and even an explanation as to which features in the spectrum influenced its prediction the most. We also find that it compares favorably with previous physics-inspired approaches such as the Hurst exponent [Nair and Sujith, 2014], autocorrelation decay [Lieuwen, 2005] and permutation entropy [Gotoda et al., 2011].

Chapter 6 takes the techniques used in the previous chapter from the laboratory into industrially relevant systems. First, we train a Bayesian neural network on multimodal data from BKD, a rocket engine combustor which is operated by DLR and equipped with pressure, temperature and optical sensors, to predict the amplitude of impending instabilities. The

model is able to forecast instabilities on unseen experimental runs quite accurately and its uncertainty estimates are well-calibrated. We also demonstrate the robustness of the Bayesian model by simulating a sensor failure and observing how this affects predictions.

Chapter 7 applies the same tools to data from an industrial intermediate pressure combustion test rig equipped with three pressure transducers. Experiments were performed on three different fuel injectors, two of which are nominally identical but display somewhat different thermoacoustic behaviour. A Bayesian neural network is trained on data from one of the injectors to predict an operating point's distance from the instability threshold and applied to all three datasets. The performance of the network is also compared to physics-inspired precursors.

Chapter 8 summarizes the contributions made in this present work and suggests some directions for future research.

Chapter 2

Bayesian machine learning

2.1 Introduction

A central theme of science is learning how to predict a quantity of interest from known inputs. Traditionally, this mapping from inputs to outputs is mediated by simple algebraic relations or differential equations. However, we often encounter problems where the formulation of an analytical model is non-obvious or where simulating the governing equations may be computationally expensive. In such cases, we need to resort to learning from large sets of data using iterative, physics-agnostic processing algorithms to extract patterns and features. In computer science, this task of automatically learning a mapping between inputs $\mathbf{x} \in \mathbb{R}^D$ and outputs $\mathbf{y} \in \mathbb{R}^m$ is known as supervised machine learning. The main difference between these machine learning algorithms and physical models is in the number of parameters. While a physical model may only need a handful of tunable parameters, the broad applicability that we desire in a machine learning algorithm means that they need to be flexible enough to represent a wide variety of functional relationships. The number of parameters can therefore range from thousands (e.g. small neural networks, decision trees) to billions (e.g. large transformer-based language models like GPT-3 [Dale, 2021]) or potentially even infinite (e.g. non-parametric models like Gaussian processes). A caveat of this flexibility is well-illustrated by the apocryphal von Neumann quote: "With four parameters I can fit an elephant, with five I can make him wiggle his trunk." A highly flexible model can fit many functional relationships that are consistent with the limited available data but most of these will fail to generalize, i.e., provide reasonable predictions for inputs that were not seen in the dataset. Therefore, we need to constrain the range of the possible functional relationships that are consistent with the data and estimate the uncertainty in our predictions. Bayesian inference provides a systematic way of doing this.

2.2 Uncertainty Quantification in Machine Learning Models

In Bayesian machine learning, uncertainty over the model parameters or hidden variables is quantified by placing a prior distribution over the hidden variables $p(\theta)$. This represents our initial naive belief of what the model's predictions should look like. Bayes' theorem then lets us update the prior distribution $p(\theta)$ using the observed data $D = \{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$ and gives us the posterior distribution of the hidden variables given the observed data: $p(\theta|D)$ in terms of the likelihood distribution $p(D|\theta)$ and the prior.

$$p(\theta|D) = \frac{p(\theta)p(D|\theta)}{\int_{\theta} p(\theta)p(D|\theta)d\theta} \quad (2.1)$$

There are two main types of uncertainty one can model [Huang et al., 2006]. Aleatoric uncertainty captures noise inherent in the observations, e.g., sensor noise and it is not possible to reduce this even if more data were to be collected. Mathematically it appears in the likelihood term $p(D|\theta)$. For example, in a regression problem, this could be the noise standard deviation σ of a Gaussian likelihood distribution. Aleatoric uncertainty can further be subdivided into homoscedastic uncertainty, uncertainty which is constant for different inputs, and heteroscedastic uncertainty, where some inputs potentially have more noisy outputs than others. Epistemic uncertainty accounts for uncertainty in the model parameters θ – uncertainty which captures our ignorance about which model generated our collected data. This uncertainty can be explained away given enough data, and is often referred to as knowledge uncertainty.

With a few exceptions such as Gaussian processes for which an analytical solution is possible, the computation of the intractable high dimensional integral $\int_{\theta} p(\theta)p(D|\theta)d\theta$, known as the marginal likelihood, poses computational challenges. The most widely used family of numerical methods to integrate over the posterior, Markov Chain Monte Carlo, is too computationally expensive to be practical as it requires us to sample from the entire parameter space with a random walk. Variational inference is often used [Blundell et al., 2015], in which the posterior is approximated using a convenient parametrization and the Kullback-Leibler divergence between this variational distribution and the true posterior is minimized using backpropagation. However, although computationally cheap, mean-field variational inference has drawbacks such as not capturing correlations between parameters. Similarly, the Laplace approximation [Ritter et al., 2018], which uses the Hessian of the probability distribution around the Maximum A-Posteriori (MAP) estimate to deduce the posterior distribution, also restricts the posterior distribution to be unimodal and Gaussian. Recently, a different

approximate inference method, based on ensembling, has been proposed. This is called the anchored ensembling algorithm. It is cheap, simple and scalable but manages to outperform variational inference in several neural network uncertainty quantification benchmarks [Pearce et al., 2018]. This is the technique we use in this thesis to perform inference in neural networks and will be discussed in greater detail in the following section.

2.3 Bayesian neural networks

In this thesis, we will work with simple feedforward neural networks. These are a repeated composition of the function $f(\mathbf{W}\mathbf{x} + \mathbf{b})$, where f is a non-linear function known as the activation function, \mathbf{W} is the matrix of weights and \mathbf{b} is the vector of biases. Each application of the function is known as a "layer" and the size of W and b is referred to as the number of nodes in that layer of the network. The weight and bias matrices are the tunable parameters θ of the neural network that we need to infer.

Anchored ensembling [Pearce et al., 2018] is a cheap, simple, scalable approximate Bayesian technique for neural networks. Consider a data set (\mathbf{x}_n, y_n) , where each data point consists of features $\mathbf{x}_n \in \mathbb{R}^D$ and output $y_n \in \mathbb{R}$. Define the likelihood for each data point as $p(y_n | \theta, \mathbf{x}_n, \sigma_\epsilon^2) = \mathcal{N}(y_n | \text{NN}(\mathbf{x}_n; \theta), \sigma_\epsilon^2)$, where NN is a neural network whose weights and biases form the latent variables θ and σ_ϵ^2 is the data noise. Define the prior on the weights and biases θ to be the standard normal $p(\theta) = \mathcal{N}(\theta | \mu_{\text{prior}}, \Sigma_{\text{prior}})$. The anchored ensembling algorithm then simply does the following:

1. The parameters $\theta_{0,j}$ of each j -th ensemble member of our neural network ensemble are initialized by drawing from the prior distribution $\mathcal{N}(\mu_{\text{prior}}, \Sigma_{\text{prior}})$.
2. Each ensemble member is trained ordinarily (e.g. using stochastic gradient descent optimizer ADAM [Kingma and Ba, 2015]) but with a slightly modified loss function that "anchors" the parameter values to their initial ones. The loss function for the j -th ensemble member is given by $Loss_{\text{anchor},j} = \frac{1}{N} \|\mathbf{y} - \hat{\mathbf{y}}_j\|_2^2 + \|\Sigma^{1/2}\theta - \theta_{0,j}\|_2^2$, where the i -th diagonal element of Σ is the ratio of data noise to the prior variance of the i -th parameter.

Pearce *et al* Pearce et al. [2018] prove that this procedure approximates the true posterior distribution for wide neural networks and that the trained neural networks in the ensemble may be treated as samples from an approximate posterior distribution. The resulting ensemble has predictions that converge when they are well-supported by the training data and diverge when they are not. The variance of the ensemble predictions is an estimate of the Bayesian model's epistemic uncertainty. The total uncertainty of each prediction is obtained by adding




Fig. 2.1 Approximate Bayesian inference using ensembles of neural networks.

the epistemic uncertainty and the irreducible aleatoric uncertainty σ_ϵ which stems from observation noise:

$$\sigma_{\text{total}}^2 = \sigma_\epsilon^2 + \frac{1}{N_e} \sum_j \hat{y}_j^2 - \left(\frac{1}{N_e} \sum_j \hat{y}_j \right)^2 \quad (2.2)$$

The tunable hyperparameters such as the learning rate, the data noise, the number of nodes in each layer are optimized by minimizing the log-likelihood of data in a validation set.

2.4 Neural network interpretation

Neural networks have a reputation as black boxes, which disincentivizes their application to problems in which practitioners must understand why the algorithm is or is not working. We use the technique of integrated gradients (IG) [Sundararajan et al., 2017a] to attribute the predictions of our network ensemble to the input features. This is a simple scalable method that only requires access to the gradient operation. IG has several desirable properties that other attribution methods lack, such as implementation invariance, sensitivity, linearity, completeness and preservation of symmetry.

For deep neural networks, the gradients of the output with respect to the input are an analog of the linear regression coefficients. In a linear model, the regression coefficients characterize the relationship between input and output variables globally. In a nonlinear neural network, however, the gradient of the output at a point merely characterizes the local relationship between a predictor variable and the output. IG computes the path integral of the gradient of the outputs with respect to the inputs around a baseline to the input under consideration. For an image recognition algorithm, a completely black image could be a reasonable choice of baseline, while for a regression problem like ours, where the input variables were normalized to have zero mean and unit standard deviation, the average input of all zeros is a sensible baseline. We consider the straight-line path (in the feature space

\mathbb{R}^D) from the baseline \mathbf{x}' to the input \mathbf{x} being considered. The IG attribution $attr_j$ for the j -th feature is then defined as follows:

$$attr_j(\mathbf{x}) = (\mathbf{x}_j - \mathbf{x}'_j) \int_{\alpha=0}^1 \frac{\partial f(\mathbf{x}' + \alpha(\mathbf{x} - \mathbf{x}'))}{\partial \mathbf{x}_j} d\alpha \quad (2.3)$$

This integral is computed numerically by sampling the gradient evenly along the path. For our neural network ensemble, we average the integrated gradient attributions obtained from individual ensemble members to obtain a mean $attr_j$ for each feature.

Understanding network predictions via Integrated Gradients involves looking at individual examples and their attribution plots. To gain a global overview of the relative importance of input features, we also use the permutation importance technique Breiman [2001]. To compute permutation feature importances, we shuffle the values of a feature between samples in the dataset and measure the impact of randomizing them on the accuracy of a trained model. The average error of model predictions should increase significantly if important features are randomized in this way and the decrease in accuracy can therefore be understood as a measure of a feature's criticality.

2.5 Gaussian Process Modelling

The underlying theory of Gaussian process (GP) regression is presented briefly below. A more detailed exposition may be found in the book by Rasmussen and Williams Rasmussen [2006].

Unlike neural networks which work with a fixed, finite set of parameters, GP models are non-parametric stochastic processes which can grow in learning capacity with the size of the training data. One can derive a Gaussian process model from a single layer neural network in the limit of infinitely many parameters [Neal, 1996]. GPs are superior to neural networks in their analytic tractability and interpretability, but the downside is their poor scaling behaviour w.r.t. dataset size and input dimension.

A GP assumes a multivariate, jointly Gaussian probability distribution for any finite collection of random variables. For example, if the function $f(\mathbf{x})$ for an input \mathbf{x} is considered a random variable, then, for a finite sub-collection $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T$, the corresponding function outputs $f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_T)$ are assumed to have a multivariate jointly Gaussian distribution.

$$\begin{bmatrix} f(\mathbf{x}_1) \\ \vdots \\ f(\mathbf{x}_T) \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} m(\mathbf{x}_1) \\ \vdots \\ m(\mathbf{x}_T) \end{bmatrix}, \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_T) \\ \vdots & & \vdots \\ k(\mathbf{x}_T, \mathbf{x}_1) & \cdots & k(\mathbf{x}_T, \mathbf{x}_T) \end{bmatrix} \right) \quad (2.4)$$

where the underlying Gaussian process is completely characterized by a mean function $m(\mathbf{x})$, $E[f(\mathbf{x})]$, and a covariance function: $k(\mathbf{x}, \mathbf{x}') = E[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$.

Let $D^{trn} = (\mathbf{x}_i^{trn}, \mathbf{y}_i^{trn})$, $i = 1, \dots, N$, be the set of observed data points that we will use to train the GP model (also referred to as the training set) and $D^{tst} = (\mathbf{x}^{tst}, \mathbf{y}^{tst})$ the test set in which \mathbf{x}^{tst} are the desired inputs for which we wish to estimate the unknown \mathbf{y}^{tst} . In the regression model, it is assumed that

$$\mathbf{y}_i = f(\mathbf{x}_i) + \varepsilon_i \quad (2.5)$$

where ε_i are independent and identically distributed (i.i.d) additive noise variables $\sim \mathcal{N}(0, \sigma^2)$. ε_i takes into account the aleatoric uncertainty of the output. Then, a zero-mean Gaussian process prior $GP(\mathbf{0}, \mathbf{K})$ is assumed for the function $f(\cdot)$. By incorporating the noise term, the joint distribution of the observed values and the functional values at the test locations are:

$$\begin{bmatrix} f(\mathbf{y}^{trn}) \\ f(\mathbf{y}^{tst}) \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \begin{bmatrix} k(\mathbf{x}^{trn}, \mathbf{x}^{trn}) + \sigma^2 I & k(\mathbf{x}^{trn}, \mathbf{x}^{tst}) \\ k(\mathbf{x}^{tst}, \mathbf{x}^{trn}) & k(\mathbf{x}^{tst}, \mathbf{x}^{tst}) + \sigma^2 I \end{bmatrix}) \quad (2.6)$$

A crucial ingredient in a GP model is the choice of covariance kernel, which encodes the correlation between points in the feature space. In this thesis, we will work with the automatic relevance determination (ARD) Matern 5/2 covariance function $k(\mathbf{x}, \mathbf{x}'; \theta) = \sigma_f^2 (1 + \sqrt{5}r + \frac{5}{3}r^2) e^{-\sqrt{5}r}$ where $r = \sqrt{\sum_{m=1}^D \frac{(x_m - x'_m)^2}{\sigma_m^2}}$. The Matern 5/2 function generates smooth, twice differentiable functions and the ARD formulation facilitates learning a length-scale for each input dimension to deal with directional anisotropies in the data set. Learning the noise σ and the kernel hyperparameters $\sigma_f, (\sigma_m)_{k=1}^D$ from the data involves maximizing the marginal likelihood or equivalently, minimizing the negative log marginal likelihood (NLML) using gradient descent methods like L-BFGS:

$$\log p(\mathbf{y}^{trn} | \mathbf{x}^{trn}) = -\frac{1}{2} \log |K + \sigma^2 I| - \frac{1}{2} \mathbf{y}^{trn T} (K + \sigma^2 I)^{-1} \mathbf{y}^{trn} - \frac{N}{2} \log(2\pi) \quad (2.7)$$

Finally, we can obtain the posterior distribution of the outputs at the test inputs. By the property of normal distribution, these are also Gaussians.

$$\mathbf{y}^{tst} | \mathbf{y}^{trn}, \mathbf{x}^{trn}, \mathbf{x}^{tst} \sim \mathcal{N}(\boldsymbol{\mu}^{tst}, \boldsymbol{\Sigma}^{tst}) \quad (2.8)$$

where the posterior means for the test outputs are given by

$$\boldsymbol{\mu}^{tst} = K(\mathbf{x}^{tst}, \mathbf{x}^{trn}) [K(\mathbf{x}^{trn}, \mathbf{x}^{trn}) + \sigma^2 I]^{-1} \mathbf{y}^{trn} \quad (2.9)$$

and the posterior covariances are given by

$$\boldsymbol{\Sigma}^{tst} = K(\mathbf{x}^{tst}, \mathbf{x}^{tst}) - K(\mathbf{x}^{tst}, \mathbf{x}^{trn}) [K(\mathbf{x}^{trn}, \mathbf{x}^{trn}) + \sigma^2 I]^{-1} K(\mathbf{x}^{tst}, \mathbf{x}^{trn}) + \sigma^2 I \quad (2.10)$$

The matrix inversions in the formulae impose a cubic computational cost penalty w.r.t. the size of the dataset N . However, there are ways to work around this, as we will see in the next section.

2.6 Global optimization using Gaussian Processes

In this thesis, we propose using gradient-augmented Bayesian optimization (BO) to optimize functions derived from adjoint models whose derivatives we have access to. BO is able to find global optima in expensive, multi-modal functions with strikingly few function evaluations. It does so by building a Gaussian Process (GP) metamodel of the objective function at every iteration. To decide where to evaluate the objective function next, an acquisition function such as expected improvement [Huang et al., 2006], upper confidence bound [Srinivas et al., 2009] or knowledge gradient [Scott et al., 2011] is used. This balances exploration (i.e., moving to places where the GP has high predictive uncertainty) and exploitation (i.e., moving to where the GP has a high predictive mean). Vanilla BO, however, does not utilize derivative information. Nevertheless, since the derivative of a GP with a twice-differentiable kernel is also a GP [Williams and Rasmussen, 1996], it is possible to model derivatives in a GP metamodel. The incorporation of derivative observations in the GP model leads to lower predictive variances (Figure 2.2 shows a simple example of a 1D GP with and without derivatives) and makes our exploration of the search space much more efficient. The main challenge is the impractical $\mathcal{O}(n^3 D^3)$ computational cost of building a GP model with n function evaluations and D derivatives. In this thesis, we use a scalable gradient-augmented Bayesian optimization algorithm loosely based on Eriksson et al. [2018].

Bayesian optimization is a powerful tool but the inclusion of D extra pieces of gradient information per point causes severe scaling issues, particularly when a large number of




Fig. 2.2 A 1D Gaussian Process metamodel with and without derivatives. Including derivative information results in a GP which has a higher confidence in its predictions. The gray area represents 3 s.d. uncertainty bounds in the Gaussian process metamodel, the red crosses are function evaluations, the red lines denote function and derivative evaluations and the dotted blue lines are randomly sampled functions from the Gaussian process metamodel. Figure adapted from Williams and Rasmussen [1996]

iterations are run. In this paper, we use the algorithm from Eriksson et al. [2018] for scaling GPs with derivatives. The cubic scaling with the number of functional evaluations for kernel learning in GPs stems from the fact that a linear system must be solved, which necessitates a Cholesky factorization of the kernel matrix. To work around this, we approximate the true kernel with a structured kernel interpolation for products (D-SKIP) approximation from Gardner et al. [2018] which allows fast matrix-vector multiplication. The structured kernel interpolation (SKI) approach uses local polynomial interpolation on an induced grid with sparse weights to approximate the kernel matrix in each dimension. Instead of computing kernel values between data points directly, SKI computes kernel values between inducing points and interpolates these kernel values to approximate the true data kernel values $k(\mathbf{x}, \mathbf{x}') \approx \sum_i w_i(\mathbf{x})k(\mathbf{x}_i, \mathbf{x}')$. The final kernel matrix can be written as a Hadamard product (\odot) of approximate one-dimensional matrices; for example, in two dimensions, the kernel matrix with derivatives is approximated as follows:

$$K^\nabla \approx \begin{bmatrix} W_1 K_1 W_1^T & W_1 K_1 \partial W_1^T & W_1 K_1 W_1^T \\ \partial W_1 K_1 W_1^T & \partial W_1 K_1 \partial W_1^T & \partial W_1 K_1 W_1^T \\ W_1 K_1 W_1^T & W_1 K_1 \partial W_1^T & W_1 K_1 W_1^T \end{bmatrix} \odot \begin{bmatrix} W_2 K_2 W_2^T & W_2 K_2 W_2^T & W_2 K_2 \partial W_2^T \\ W_2 K_2 W_2^T & W_2 K_2 W_2^T & W_2 K_2 \partial W_2^T \\ \partial W_2 K_2 W_2^T & \partial W_2 K_2 W_2^T & \partial W_2 K_2 \partial W_2^T \end{bmatrix} \quad (2.11)$$

where W_j and K_j denote the structured kernel interpolation (SKI) and inducing point grid matrices in the j -th coordinate direction. Additionally, we use the iterative conjugate gradient

method with pre-conditioning for solving the linear system instead of Cholesky factorization. These tricks enable $\mathcal{O}(nD)$ complexity kernel learning for the GP, where n is the number of data points and $\mathcal{O}(1)$ prediction per test points.

Once we have the GP surrogate model of our function, we use its mean and uncertainty to compute the expected improvement acquisition function (Jones et al. [1998]) which helps us decide where best to sample the function and its derivative in each iteration. Expected improvement is defined as $\text{EI}(\mathbf{x}) = \mathbb{E}[\max(f(\mathbf{x}^+) - f(\mathbf{x}), 0)]$ where $f(\mathbf{x}^+)$ is the value of the best sample so far and \mathbf{x}^+ is the location of that sample i.e. $\mathbf{x}^+ = \operatorname{argmax}_{\mathbf{x}_i \in \mathbf{x}_{1:t}} f(\mathbf{x}_i)$. The expected improvement can be evaluated analytically for a Gaussian process:

$$\text{EI}(\mathbf{x}) = \begin{cases} (f(\mathbf{x}^+) - f(\mathbf{x}) - \xi)\Phi(Z) + \sigma(\mathbf{x})\phi(Z) & \text{if } \sigma(\mathbf{x}) > 0 \\ 0 & \text{if } \sigma(\mathbf{x}) = 0 \end{cases} \quad (2.12)$$

where

$$Z = \begin{cases} \frac{f(\mathbf{x}^+) - \mu(\mathbf{x}) - \xi}{\sigma(\mathbf{x})} & \text{if } \sigma(\mathbf{x}) > 0 \\ 0 & \text{if } \sigma(\mathbf{x}) = 0 \end{cases} \quad (2.13)$$

$\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ are the mean and the standard deviation of the GP posterior predictive at \mathbf{x} , respectively. Φ and ϕ are the CDF and PDF of the standard normal distribution, respectively. The first summation term in Equation (3) is the exploitation term which encourages the algorithm to explore where the surrogate model has a lower predictive mean than the current best sample and the second term is the exploration term which encourages sampling new points in regions of high uncertainty.

The parameter ξ in Equation (3) determines the amount of exploration during optimization and higher ξ values lead to more exploration. With increasing ξ values, the importance of improvements predicted by the GP posterior mean $\mu(\mathbf{x})$ decreases relative to the importance of potential improvements in regions of high prediction uncertainty, represented by large $\sigma(\mathbf{x})$ values. We use the default value $\xi = 0.01$ recommended by Lizotte [2008].

In most high-dimensional optimization problems of practical significance, there exist only a few relevant directions that capture most of the variation in the function. This is exploited by the active subspace dimension reduction method [Constantine et al., 2014]. The optimal subspace is given by the dominant eigenvectors of the covariance matrix $C = \int_{\Omega} \nabla f(\mathbf{x}) \nabla f(\mathbf{x})^T$, estimated by Monte Carlo integration. Because we already have access to gradient information, it is natural to combine active subspaces with BO to reduce the computational overhead of building our GP metamodel at every iteration. Random projections (without gradients) have already been used in the BO literature for high-dimensional functions

Wang et al. [2013]. We learn the active subspace that captures 95% of the variation at every iteration, choose d random directions of the active subspace onto which the function is to be projected and add a small amount of random noise to the projection matrix.

Algorithm 1: BO with derivatives and active subspace learning

sample cost function f and gradients ∇f at D initial points using simple Latin

Hypercube sampling, where D is the dimension of the design space.

while *computational budget remaining* **do**

Estimate active subspace of cost function using sampled gradients that capture 99% of the variation.

Pick d directions at random from estimated active subspace to compute projection $P \in \mathbb{R}^{D \times d}$.

Add a small amount of random Gaussian noise to obtain the final projection matrix P_{noisy} .

Update hyperparameters of GP with gradient defined by kernel $k(P_{noisy}^T \mathbf{x}, P_{noisy}^T \mathbf{x}')$ and construct the GP surrogate.

Use the GP surrogate to optimize the expected improvement

$\mathbf{u}_{n+1} = \operatorname{argmax} \operatorname{EI}(\mathbf{u})$ to obtain the next candidate point $\mathbf{x}_{n+1} = P_{noisy} \mathbf{u}_{n+1}$.

Sample point \mathbf{x}_{n+1} , value f_{n+1} , and gradient ∇f_{n+1}

Update data $D_{n+1} = D_n \cup \{x_{n+1}, f_{n+1}, \nabla f_{n+1}\}$

end

Chapter 3

Parameter inference in reduced order flame models

The research in this chapter is a product of collaboration with Maximilian Croci. The project was devised with him. I performed the Bunsen flame experiments, wrote the Bayesian Neural Network code and ran the parameter inference on processed flame images. Maximilian ran the LSGEN flame simulations, performed thresholding and interpolation on the flame image data and compared our results to the Ensemble Kalman Filter. This work was presented at the International Conference on Computational Science 2021 where it was selected for publication as an article in the Lecture Notes in Computer Science book series titled "Data Assimilation Using Heteroscedastic Bayesian Neural Network Ensembles for Reduced-Order Flame Models" (volume 12746, pp 408-419).

3.1 Introduction

Complex, nonlinear physical models of engineering systems give rise to challenging inverse problems which require the estimation of unknown model parameters from observations, along with uncertainty quantification. With the growing popularity of digital twins [Fuller et al., 2020], there is also a need to reduce the computational cost of parameter inference such that these models may be updated in real-time using the latest sensor observations. Traditional filtering-based techniques require the system to be simulated in parallel, which becomes unfeasible if the underlying model is computationally expensive to simulate. Amortized inference using neural networks [Cranmer et al., 2020] circumvents this problem by having an expensive offline training phase, where a surrogate of the approximate posterior $p(\theta|\mathbf{z})$ is learnt from a library of simulator-generated observations \mathbf{z}_i corresponding to input parameters

θ_i , typically using normalizing flow-based methods [Rezende and Mohamed, 2015] [Radev et al., 2020]. The surrogate can then be rapidly evaluated online to perform parameter inference on new observed data.

Thermoacoustic instabilities are driven by the coupling of the heat release rate of the flame and the acoustics in the combustor [Strutt, 1878]. It is therefore necessary to model the flame dynamics such that the heat release rate as a function of velocity perturbations is quantitatively accurate. The simplest physics-based model sets the heat release rate fluctuation to be a linear multiple of the velocity perturbation at the base of the flame some time earlier. This delay models the time taken for perturbations to travel down the flame. This is known as the $n - \tau$ model. It is too simple for our purposes because it cannot simulate flame dynamics. The G -equation [Williams, 1985] is the kinematic model used in LSGEN2D [Hemchandra, 2009] to simulate a ducted, premixed flame such as that from a Bunsen burner. By tuning the parameters of the G -equation model to fit a Bunsen flame experiment, a digital twin of the flame is created and its surface area variation in time, which when multiplied by the laminar flame speed gives us the heat release rate variation, can be calculated. The ensemble Kalman filter (EnKF) is the current state-of-the-art, which iteratively performs Bayesian inference of the G -equation parameters from G -equation model forecasts generated in LSGEN2D and observations of the flame edge [Yu et al., 2019b]. The model forecasting is expensive, however, which makes the EnKF too expensive to be used online. We propose to infer online the G -equation parameters of Bunsen flame experiments using a heteroscedastic Bayesian neural network ensemble [Sengupta et al., 2020a,b]. In an expensive offline step, the ensemble learns a surrogate for the Bayesian posterior distribution of the parameters given observations of the flame front from a library of pairs of G -equation parameters and corresponding flame front shapes generated in LSGEN2D. The ensemble can then be used to infer online the parameters of the model for simulation of the Bunsen flames.

3.2 Bunsen flame experiment and simulations

3.2.1 Bunsen flame experiment

A Bunsen burner is placed inside a transparent duct and a high-speed camera is used to take images of the Bunsen flame at a frame rate of $f_s = 2500$ frames per second and a resolution of 1200×800 pixels. Speakers force the flame at frequencies in the range 250 Hz to 450 Hz. The gas composition (methane, ethene and air) and flow rate are varied using mass flow controllers. By varying the forcing frequency and amplitude and gas composition and flow rate, flames with different aspect ratios, propagation speeds and degrees of cusping of the

flame front are observed. In some cases, the flame front cusping leads to pinch-off at the flame tip. For each of the 270 different flame operating conditions, 500 images are taken.

The flame images are processed to find a single-valued discretisation of the flame front, $x = f(y)$. First, the pixel intensities are thresholded and a position x for every vertical position y is found by weighted interpolation of the thresholded pixels, where the weights are the pixel intensities. Next, splines with 28 knots are used to smooth the (x, y) coordinates. Each flame image is therefore converted into a 90×1 vector of flame front x coordinates \mathbf{x} (as the y coordinates are the same for all flames, they are discarded). Observation vectors \mathbf{z} are created by stacking 10 subsequent \mathbf{x} vectors. These observation vectors are used for inference with the neural networks. All 500 images of each Bunsen flame are processed in this way.

3.2.2 Flame front model

In this chapter we use a kinematic model of the flame front as a boundary between reactants and products (see Figure 3.1). The flame front is defined to be the $G = 0$ contour of a scalar field $G(x, y, t)$. Regions of negative and positive G correspond to unburnt and burnt gases respectively (the magnitude of G does not have a useful meaning). The position of the flame front in space and time is governed by:

$$\frac{\partial G}{\partial t} + \mathbf{v} \cdot \nabla G = s_L |\nabla G|, \quad (3.1)$$

where \mathbf{v} is a prescribed velocity field and s_L is the laminar flame speed: the speed at which the flame front propagates normal to itself into the reactants. The flame speed $s_L = s_L^0 (1 - \mathcal{L} \kappa)$ is a function of the unstretched (adiabatic) flame speed s_L^0 , the flame curvature κ and the Markstein length \mathcal{L} , and is insensitive to pressure variations. The unstretched flame speed s_L^0 depends only on the flame chemistry. The velocity field $\mathbf{v} = u' \mathbf{i} + (V(x) + v') \mathbf{j}$ comprises a parabolic base flow profile $V(x) = V(1 + \alpha(1 - 2(\frac{x}{R})^2))$ and superimposed continuity-obeying velocity perturbations $u'(x, y, t) = V \varepsilon \sin(\text{St}(\frac{K}{R}y - t))$ and $v'(x, y, t) = -\frac{V \varepsilon K \text{St}}{R} x \cos(\text{St}(\frac{K}{R}y - t))$ where α determines the shape of the base flow profile ($\alpha = 0$ is uniform flow, $\alpha = 1$ is Poiseuille flow), ε is the amplitude of the vertical velocity perturbation with phase speed V/K , $\text{St} = 2\pi f R \beta / V$ is the Strouhal number with forcing frequency f and flame radius R , and β is the aspect ratio of the unperturbed flame. The parameters $K, \varepsilon, \mathcal{L}, \alpha, \text{St}$ and β are tuned to fit an observed flame shape.




Fig. 3.1 (Left) Image of a Bunsen flame. (Middle) G -equation model of the flame front. (Right) Simulated flame in LSGEN2D.

3.2.3 Simulated flame front library

A library of simulated flame fronts with known parameters $K, \varepsilon, \mathcal{L}, \alpha, St$ and β is created for neural network training. The parameter values are sampled using quasi-Monte Carlo sampling to ensure good coverage of the parameter space. The parameters are sampled from the following ranges: $0.0 < K \leq 1.5$, $0.0 < \varepsilon \leq 1.0$, $0.02 \leq \mathcal{L} \leq 0.08$, $0.0 \leq \alpha \leq 1.0$, $2.0 \leq \beta \leq 10.0$ and $0.08 \leq f/f_s \leq 0.20$. The values of St are calculated by additionally sampling $0.002 \leq R \leq 0.004$ m and $1 \leq V \leq 5$ m/s and calculating $St = 2\pi f R \beta / V$. The parameters are sampled 8500 times, normalised to between 0 and 1 and recorded in target vectors $\{\mathbf{t} = [K, \varepsilon, \mathcal{L}, \alpha, St, \beta]\}$.

For each of the 8500 unique parameter configurations, LSGEN2D iterates the G -equation model of the flame front until it converges to the corresponding forced cycle. For each of the 200 forced cycle G field states produced by LSGEN2D, the flame front is found by interpolating the G values. This results in a $x = f(y)$ discretisation and the vectors \mathbf{x} of x coordinates are recorded. Observation vectors \mathbf{z} are created by stacking 10 subsequent \mathbf{x} vectors. There are 200 observation vectors created from every cycle, resulting in a library of 1.7×10^6 observation-target parameter pairs $\{(\mathbf{z}, \mathbf{t})\}$. This library is split 80%-20% into training and testing data sets.

3.3 Inference using heteroscedastic Bayesian neural network ensembles

We assume the posterior probability distribution of the parameters given the observations can be modelled by a neural network: $p_\theta(\mathbf{t}|\mathbf{z})$ with its own parameters θ . We assume this

posterior distribution has the form:

$$p_{\theta}(\mathbf{t}|\mathbf{z}) = N(\mathbf{t}; \boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\Sigma}(\mathbf{z})), \quad (3.2)$$

where $\boldsymbol{\Sigma}(\mathbf{z}) = \text{diag}(\boldsymbol{\sigma}^2(\mathbf{z}))$. This encodes our assumption that the parameters are all mutually independent given the observations \mathbf{z} . The architecture of a single neural network is shown in Figure 3.2. Each neural network comprises an input layer, four hidden layers with ReLU activations and two output layers: one for the mean vector $\boldsymbol{\mu}(\mathbf{z})$ and one for the variance vector $\boldsymbol{\sigma}^2(\mathbf{z})$. The output layer for the mean uses a sigmoid activation to restrict outputs to the range $(0, 1)$. The output layer for the variance uses an exponential activation to ensure positivity. An ensemble of $M = 20$ such neural networks is initialised, each with unique initial weights $\boldsymbol{\theta}_{j,anc}$ sampled from a Gaussian prior distributions according to He initialisation [He et al., 2015].

For a single observation \mathbf{z} , the j -th neural network in the ensemble produces a sample of the posterior with an associated estimate of the aleatoric noise in the observations: $\boldsymbol{\mu}_j(\mathbf{z})$, $\boldsymbol{\sigma}_j^2(\mathbf{z})$. This is achieved by using the loss function \mathcal{L}_j :

$$\mathcal{L}_j = (\boldsymbol{\mu}_j(\mathbf{z}) - \mathbf{t})^T \boldsymbol{\Sigma}_j(\mathbf{z})^{-1} (\boldsymbol{\mu}_j(\mathbf{z}) - \mathbf{t}) + \log(|\boldsymbol{\Sigma}_j(\mathbf{z})|) + (\boldsymbol{\theta}_j - \boldsymbol{\theta}_{anc,j})^T \boldsymbol{\Sigma}_{prior}^{-1} (\boldsymbol{\theta}_j - \boldsymbol{\theta}_{anc,j}). \quad (3.3)$$

The loss function comprises the negative log of the Gaussian likelihood function (probability of the observations given the targets, first two terms) and a regularising (penalty) term. By regularising about parameter values drawn from a prior distribution, the NNs produce samples from the posterior distribution. This is called randomised maximum a-posteriori (MAP) sampling [Pearce et al., 2020].

Once converged, the prediction from the ensemble for an observation \mathbf{z} is therefore a mixture of M Gaussians each centered at their respective means $\boldsymbol{\mu}_j(\mathbf{z})$. This mixture is approximated by a single multivariate Gaussian posterior distribution $p(\mathbf{t}|\mathbf{z}) \approx \mathcal{N}(\mathbf{t}; \boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\Sigma}(\mathbf{z}))$ with mean $\boldsymbol{\mu}(\mathbf{z}) = \frac{\sum_j \boldsymbol{\mu}_j(\mathbf{z})}{M}$ and covariance $\boldsymbol{\Sigma}(\mathbf{z}) = \text{diag}(\boldsymbol{\sigma}^2(\mathbf{z}))$ where $\boldsymbol{\sigma}^2(\mathbf{z}) = \frac{\sum_j \boldsymbol{\sigma}_j^2(\mathbf{z})}{M} + \frac{\sum_j \boldsymbol{\mu}_j^2(\mathbf{z})}{M} - \left(\frac{\sum_j \boldsymbol{\mu}_j(\mathbf{z})}{M}\right)^2$ following similar treatment in Lakshminarayanan et al. [2017]. This is done for every observation vector \mathbf{z} . The posterior distribution $p(\mathbf{t}|\mathbf{z}_i)$ with the smallest total variance $\boldsymbol{\sigma}_{i,tot}^2 = \|\boldsymbol{\sigma}^2(\mathbf{z}_i)\|_1$ is chosen as the best guess to the true posterior. The M parameter samples from the chosen posterior are used for re-simulation, which allows us to check the predicted flame shapes and to calculate the normalised area variation over one cycle.

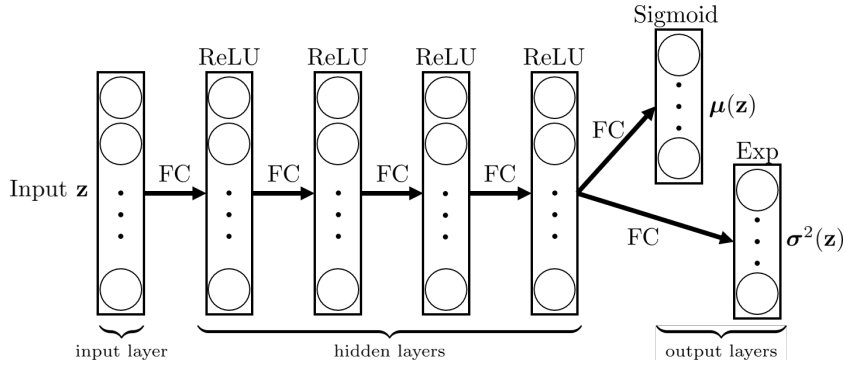


Fig. 3.2 Architecture of a single neural network. The input and hidden layers have 900 nodes each, while the output layers have 6 nodes each. All layers are fully connected (FC). Rectified Linear Unit (ReLU) activation functions are used for the hidden layers and sigmoid and exponential (Exp) activation functions are used for the mean and variance output layers respectively.

3.4 Inference using the Ensemble Kalman Filter

The Kalman filter iteratively performs Bayesian inference to find the probability distribution of the state of a system given noisy observations of the system and an imperfect model of the system dynamics. In this study, the state comprises the location of the flame edge and the parameters K and ε . These parameters are assumed to be independent given the observations of the flame edge and constant for each of the Bunsen flame experiments. The flame edge is modelled using the G-equation (3.1). The ensemble Kalman filter (EnKF) [Katzfuss et al., 2016] evolves an ensemble of simulations forward in time. The covariance matrix of these ensembles is assumed to approximate the covariance matrix of the state evolved over the same period of time. The EnKF is more practical when the state contains many variables and the evolution is nonlinear. In this study, the state contains $O(10^2)$ variables and the governing equation (3.1) is nonlinear.

The parameters \mathcal{L} , α , St , β are calculated by solving the G-equation when steady and do not need to be inferred with the EnKF. This reduces the cost of the EnKF but increases the number of steps compared with the BayNNE method. The forcing frequency f is manually set when running the Bunsen flame experiments. The unperturbed laminar flame speed s_L^0 is calculated using Cantera and knowledge of the methane and ethene flow rates. V is calculated from s_L^0 and β using the relation $V = s_L^0 \sqrt{\beta^2 - 1}$. The Strouhal number can then be calculated: $St = 2\pi f \beta \mathcal{L} / V$. Further details on the implementation of the EnKF may be found in the original paper by Yu et al. [2019b].

An ensemble size of 32 is used in this study. A multiplicative inflation factor of 1% is chosen to mitigate the underestimation of the error covariances due to the finite ensemble




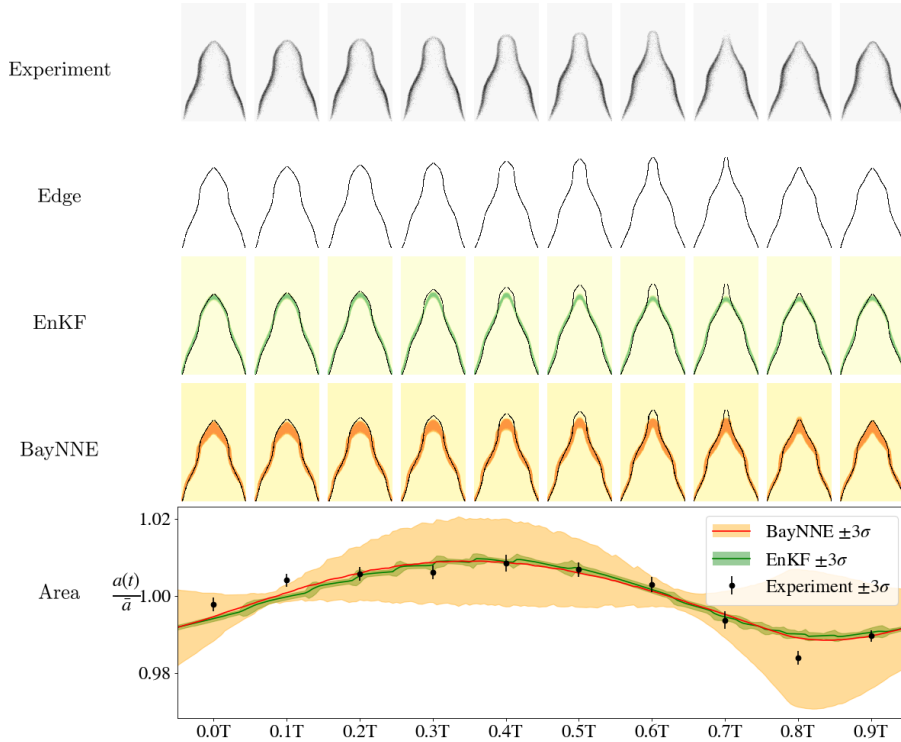
Fig. 3.3 Scatter plots of true parameters (scaled) versus predicted parameter values for simulated test data.

size [Yu et al., 2019a]. Once the EnKF has converged, the parameters K and ε calculated by each ensemble member are recorded. These are samples from the posterior distribution of the parameters given all the flame x-location vectors: $p(K, \varepsilon | \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)$. This differs from the BayNNE, which was given x-location vectors in groups of 10.

3.5 Results


The neural network is first evaluated on the test set of simulated flame observations to ensure that it recovers the true parameters (known in this case) correctly. The correlation coefficient ρ between true and predicted parameter values are 0.982, 0.994, 0.971, 0.993, 0.976 and 0.990 for K , ε , \mathcal{L} , α , St and β , respectively. Estimates of parameter uncertainty are also well-calibrated, with particularly high uncertainties in K and St predicted for datapoints with low ε_a (amplitude of perturbation), which is physically sensible, as it is not possible to recover these parameters from an unperturbed flame.

Next, the ensemble is evaluated on the Bunsen flame image data. Parameter predictions for an observation vector \mathbf{z} take $O(10^{-4})$ seconds on an Nvidia P100 GPU. Including the pre-processing time for thresholding, spline fitting and area calculation, predictions can therefore be made in less than a millisecond for every few seconds of acquired video data. Figure 3.4 show two different Bunsen flames simulated using the parameters inferred by the BayNNE and the EnKF, overlaid on the real experimental data. The BayNNE predicts flame shapes in good agreement with the experiments: the root mean square distances between the



simulated and Bunsen flame shapes of 5 different Bunsen flames ranges from 0.017 units to 0.024 units, with mean 0.019 units (where the flame radius is 1 unit). These results show that the BayNNE's parameter estimates match those of the EnKF and require $O(10^8)$ less computing power.

The uncertainty in the BayNNE's predictions is greater than that of the EnKF due to the calculated posterior being the probability of the parameters given 10 flame images, whereas the EnKF considers all 500 flame images of the Bunsen flame. Unfortunately, it is not possible to combine the BayNNE's parameter estimates without knowledge of the joint probability distribution of the observation vectors. Any two observation vectors are not independent as knowledge about the first restricts our expectation of the second to a likely set of forced cycle states. Future work will address this limitation by using alternative neural network architectures, such as long-short term memory networks [Hochreiter and Schmidhuber, 1997], that are better suited to time-series data.



Chapter 4

Global optimization of combustor models using Bayesian optimization and adjoint methods

This chapter has been published as an article in the International Journal of Spray and Combustion Dynamics titled "Thermoacoustic stabilization of combustors with gradient-augmented Bayesian optimization and adjoint models".

4.1 Introduction

Thermoacoustic stability is exceedingly sensitive to changes in the acoustic characteristics of a combustion chamber and the acoustic/ hydrodynamic response of the flame. As a result, small changes in the system geometry, boundary conditions or flame behavior can often stabilise an unstable combustor. However, coming up with a suitable design modification through pure trial-and-error experimentation is usually infeasible. A famous example is the F1 engine of the Saturn V rocket whose stabilization required 2000 full-scale engine tests; baffles were introduced in the injector plate to suppress thermoacoustic oscillations [Joos, 2006]. Adjoint-based sensitivity analysis in thermoacoustic models offers a systematic procedure for discovering appropriate design changes.

The advantage of adjoint methods lies in the efficient computation of derivatives of a given objective function at a cost that scales independently of the number of design variables. The sensitivity of eigenvalues (natural frequencies and growth rates) of a thermoacoustic model w.r.t. to design variables can therefore be computed cheaply using adjoint methods; this information is then utilized by an optimization routine from the gradient descent family

to arrive at a stabilized design. Adjoint-based sensitivity analysis has been applied to low-order network models by Magri and Juniper [2013], Mensah and Moeck [2017] and Silva et al. [2017], to 2D Helmholtz equation models by Falco and Juniper [2021] and to 3D Helmholtz models by Mensah et al. [2019]. A comprehensive review of the use of adjoints in thermoacoustic can be found in Magri [2019].

Other techniques have also been used to stabilize designs. Aguilar Perez [2019] used an exhaustive grid search while Jones et al. [2021] used a genetic algorithm to optimize the shape of a thermoacoustically unstable combustor. Compared to gradient descent using adjoints, these tools have the advantage of not getting trapped in local minima, but are much slower in comparison and require many expensive simulations to be run for each configuration they evaluate.

In this chapter, we propose using gradient-augmented Bayesian optimization (BO) to optimize adjoint models. BO is able to find global optima in expensive, multi-modal functions with strikingly few function evaluations. It does so by building a Gaussian Process (GP) metamodel of the objective function at every iteration. To decide where to evaluate the objective function next, an acquisition function such as expected improvement [Huang et al., 2006], upper confidence bound [Srinivas et al., 2009] or knowledge gradient [Scott et al., 2011] is used. This balances exploration (i.e., moving to places where the GP has high predictive uncertainty) and exploitation (i.e., moving to where the GP has a high predictive mean). Vanilla BO, however, does not utilize derivative information. Nevertheless, since the derivative of a GP with a twice-differentiable kernel is also a GP [Williams and Rasmussen, 1996], it is possible to model derivatives in a GP metamodel. The incorporation of derivative observations in the GP model leads to lower predictive variances (Figure 1 shows a simple example of a 1D GP with and without derivatives) and makes our exploration of the search space much more efficient. The main challenge is the impractical $\mathcal{O}(n^3 D^3)$ computational cost of building a GP model with n function evaluations and D derivatives. In this chapter, we use a scalable gradient-augmented Bayesian optimization algorithm loosely based on Eriksson et al. [2018].

To demonstrate the effectiveness of gradient-augmented Bayesian optimization, we consider two optimization problems from the literature on adjoint models in thermoacoustics. The first one is a simple toy problem from Juniper [2018] where the iris diameter and heater position in a 1D Helmholtz equation model of a Rijke tube must be optimized. The second involves geometry optimization in a low-order network model of a longitudinal combustor from Aguilar and Juniper [2020]. We compare gradient-augmented BO with BFGS, a standard quasi-Newton optimizer which uses an estimate of the inverse Hessian matrix to improve convergence. Compared to BFGS, we find that the gradient-augmented BO does

not get stuck in local optima and requires fewer evaluations of the solver to arrive at more thermoacoustically stable configurations. It is also able to efficiently explore the whole design space and locate multiple suitable combinations of design parameters, which is useful further downstream in the design cycle.

4.2 Methods

4.2.1 Thermoacoustic models

Our goal is to design systems in which all eigenmodes decay in time, so we use the summed exponential of the growth rates as our objective function. This strongly penalizes positive growth rates, but the rewards for stabilizing the system diminish progressively as growth rates become less than zero. If X is the vector that contains all the combustor parameters and λ_i -s are growth rates of the eigenmodes of interest, then our cost function is given by

$$\mathcal{J}(X) = \sum_i e^{\lambda_i} \quad (4.1)$$

One could alternatively optimize $\log(1 + \sum_i e^{\lambda_i})$, which approximates the above objective for highly negative values of λ_i , but increases linearly with the growth rates for large values. This ensures numerical stability since the function values and their gradients do not become arbitrarily large.

We only consider longitudinal thermoacoustic modes in this chapter, but we note that the above expression would result in double-counting of symmetric azimuthal modes in a system that has them. When optimizing such a system, the corresponding terms should be weighed by 0.5 to avoid penalizing them doubly.

The first test case considered in this chapter comes from Juniper [2018] where optimization using adjoints is demonstrated on a 1D Helmholtz model of a Rijke tube. A finite element discretization of the weak form is used and the discrete adjoint framework is used to compute eigenvalue sensitivities. The design parameters are the heater position and the diameter of a variable-diameter iris placed at the downstream boundary. The growth rate of the fundamental mode, which we seek to minimize, has two local minima (Figure 4.2) in the search space.

The second test case is taken from Aguilar and Juniper [2020] and relates to geometry optimization in a network model of Rama Balachandran's 10 kW longitudinal combustor built in Cambridge, originally intended for the experimental investigation of the response of turbulent premixed flames to acoustic oscillations [Balachandran et al., 2005]. The geometry

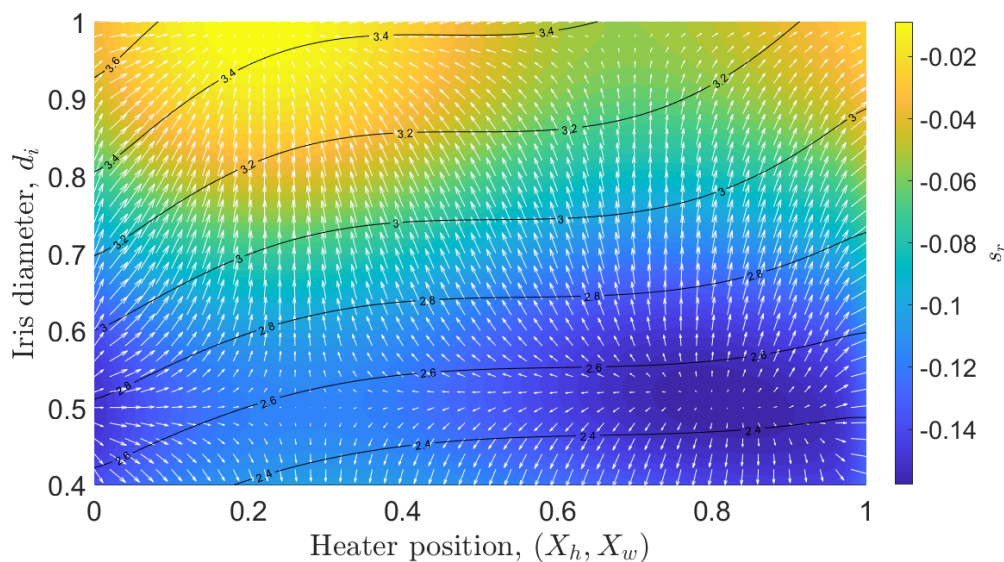


Fig. 4.1 Contour plots of non-dimensionalized growth rate and frequency of the Rijke tube fundamental mode plotted against heater position and iris diameter, reproduced from Juniper [2018]. The numbered black contour lines correspond to frequencies and the colored contour plot represents the growth rate. White arrows indicate the direction and magnitude of the growth rate gradients, as computed by the adjoint model.




Fig. 4.2 Geometric parameters of the network model from Aguilar Perez [2019] used in the optimization routine. The parameters in black are allowed to be updated. The parameters in red are kept constant.

consists of an inlet duct connected to a plenum with a linearly varying cross section on either end. This leads to the neck which contains the fuel injection plane and a centred bluff body used to stabilize the flame. The outlet is a cylindrical pipe which contains the flame. Figure 4.3 shows a schematic of the combustor model and the geometric parameters being optimized. The rig is modeled using a one-dimensional network model with 124 straight ducts. This model has 4 unstable eigenmodes between 0 and 1000 Hz that need to be stabilized with geometry modifications. We restrict the search space to a box bounded by $\pm 4.5\%$ of the areas and lengths in the original configuration.

4.2.2 Gradient-augmented Bayesian optimization

We use the gradient-augmented Bayesian optimization algorithm with active subspace dimension reduction described in Chapter 2. For the longitudinal combustor problem, we learn the active subspace that captures 95% of the variation at every iteration and we retain $d = 3$ random directions of the subspace at every iteration. The subspace projection is not used for the 1D Rijke tube problem, since the optimization problem is two-dimensional.

4.3 Results

For the 1D Rijke tube problem, the gradient-augmented BO consistently finds the known global optimum near $(0.82, 0.5)$ fairly rapidly across all trials, whereas a gradient-descent-based approach such as BFGS predictably converges to either of the two local optima (the other one is located at $(0.0, 0.5)$), depending on the starting point chosen. Figure 4.4 illustrates how the gradient-augmented Bayesian optimization refines its surrogate model with increasing number of iterations. After 7 evaluations, the surrogate is very inaccurate because the bottom right and top left quadrants are left unexplored. However, these regions also have higher uncertainty, so we observe that by 17 evaluations, these regions have also been explored. After this exploration phase, the BO deems that the model is accurate enough for the purpose of optimization and spends the next 10 evaluations trying to converge on the global optimum. BO consistently explores both local minima regardless of initialization. This is very useful in thermoacoustics because almost all design optimization problems have multiple stable configurations. Discovering more of them makes the designer's life easier, since thermoacoustic stability is typically just one of many conflicting design objectives. The MATLAB code for the 1D Rijke tube optimization may be found [here](#).

To compare between gradient-augmented BO and BFGS, we conduct 10 trials of both algorithms on the longitudinal combustor test problem. In each trial, the initial points are

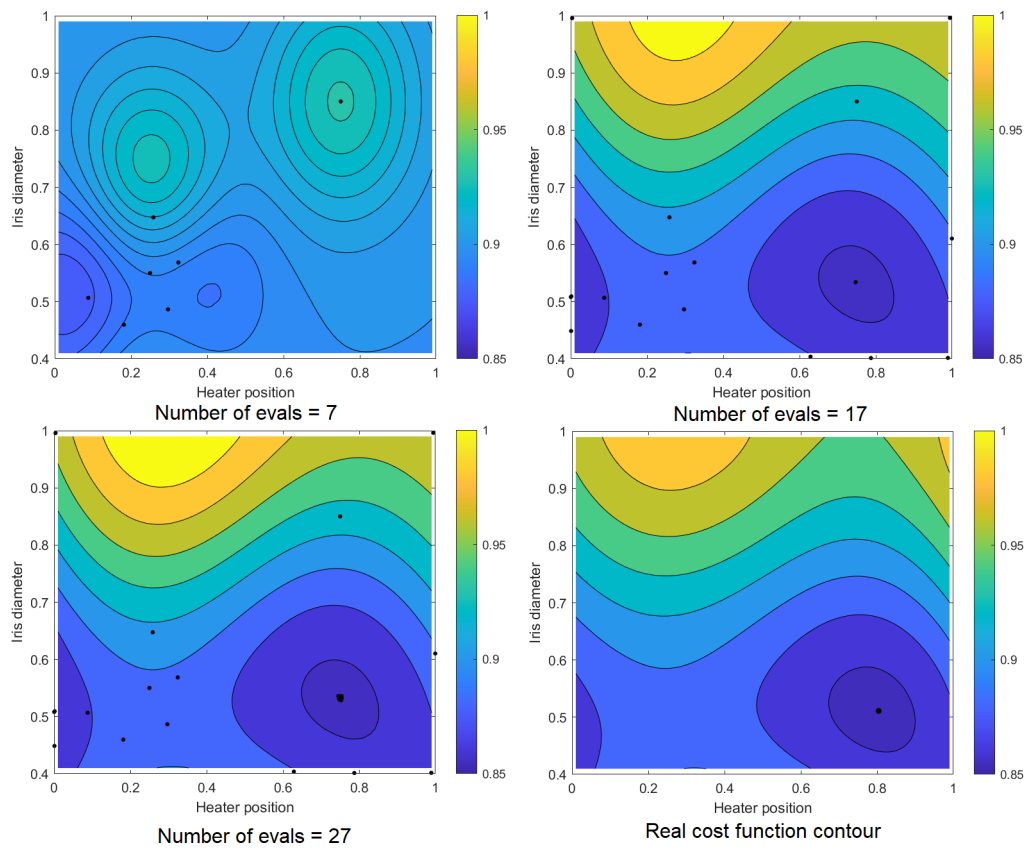


Fig. 4.3 Surrogate Gaussian process model of the Rijke tube cost function at different iterations. Black dots indicate where the Bayesian optimization has evaluated the function and its gradients. The bottom right plot is a contour plot of the true cost function, with the black dot indicating the location of the true global minimum.




Fig. 4.4 Best objective function values achieved by the BFGS, BO without gradient and gradient-augmented BO routines in the longitudinal combustor test case, averaged across 10 trials, plotted against the number of cost function evaluations. Error bars indicate 1 s.d. variation across trials.

sampled uniformly from the design space with a different seed for the pseudo-random number generator. When the number of iterations is very small, the BFGS appears to outperform Bayesian optimization, but this is because Bayesian optimization explores more early on due to the high initial uncertainty in the Gaussian Process surrogate. BO quickly overtakes the BFGS average, however, and by the 50th iteration, 9 out of 10 BO runs have found a configuration as stable as the best BFGS run. The final average cost function after 50 runs is 1.62 for BFGS and 1.01 for BO. As in the Rijke tube case, multiple distinct stable geometries are found in each run. In general, stable configurations have a reduced flame holder area and a reduced neck area, because the eigenvalues are quite sensitive to these two parameters. Figure 4.5 shows a stable geometry found by one of the BO runs.

Figure 4.4 also shows the comparison between the gradient-augmented BO to BO without gradient values. As expected, it converges to the global optimum more slowly. This shows the benefit of having derivative information from the adjoint models.

It should be pointed out the BO involves a modest computational overhead whereas BFGS has nearly none. For the longitudinal combustor model, it takes around 30 seconds per iteration for the gradient-augment BO on a Dell G7 7590 laptop with a 4 core 8th Generation Intel i5-8300HQ processor. However, when optimizing expensive models with this algorithm, the evaluation of the cost function and its gradient becomes the computational bottleneck while the cost per BO iteration becomes a non-issue.

This chapter demonstrated how scalable, gradient-augmented Bayesian optimization pairs perfectly with adjoint methods in thermoacoustics. The BO algorithm consistently




Fig. 4.5 A stable combustor configuration, shown in red, found by a gradient-augmented Bayesian optimization run after 50 evaluations. Original unstable configuration in black.

finds more stable parameter combinations using considerably fewer iterations than BFGS, a popular quasi-Newton optimizer.

Chapter 5

Prognosing instabilities from noise in a Rijke burner

The work in this chapter has been published as an article in the Journal of Engineering for Gas Turbines and Power titled "Bayesian machine learning for the prognosis of combustion instabilities from noise" (vol 143 issue 7).

5.1 Introduction

The phase lag between pressure and heat release rate fluctuations, which governs the thermoacoustic stability of a system, depends on acoustic, hydrodynamic and combustion mechanisms, which have different scaling behaviours. Accurate computational modeling is thus very challenging [Juniper and Sujith, 2018]. At the moment, designers accommodate thermoacoustic instabilities in their engines by avoiding unstable regions of the operating parameter space. This, however, conflicts with other design objectives such as reducing NOx emissions by operating at leaner fuel-air ratios. The aim of this chapter is to develop and test a machine learning algorithm that can learn how close a combustor is to instability and ensure its safe operation near unstable conditions.

5.1.1 Combustion noise as a diagnostic

The noise radiated by a turbulent combustor is generated by deterministic fluid dynamic phenomena such as unsteady dilatation due to fluctuating heat release rates or the acceleration of vorticity or entropy waves, and modified by acoustic reflections off boundaries [Dowling and Mahmoudi, 2015]. We therefore expect pressure measurements to contain some information about the state of the combustor. Inspired by Zelditch [2000] who answered Kac's whimsical

question “Can you hear the shape of a drum?” [Kac, 1966] by proving that the eigenfrequencies of vibration determine the shape of an analytic and convex membrane uniquely, our study seeks to extract useful information about the state of a turbulent combustor from the power spectra of noise samples. This would have important practical implications. For example, it would allow noise to serve as an early warning prognostic for thermoacoustic instability or blowoff. It would also enable pressure measurements to validate the readings of other sensors such as flowmeters, making the system more robust to sensor failures. Pressure and vibration measurements are easily accessible in fielded combustors so it makes sense to use them as extensively as possible.

Historically, the motivation to understand and model combustion noise stems from the desire to reduce noise pollution, such as that from an aircraft [Duran et al., 2014] or a factory furnace [Putnam, 2009]. Pioneering theoretical work by Lighthill in aeroacoustics [Lighthill, 1961] was extended to the analytical study of combustion noise by Strahle [Strahle, 1971], who used Lighthill’s acoustic analogy to derive a formula relating the far-field acoustic perturbation to heat release rate fluctuations by treating the turbulent pre-mixed flame as an assembly of monopole sound sources. It was also noted by Marble and Candel [1977] that noise can be generated by entropy inhomogeneities in regions of accelerating flow when combustion occurs in confined chambers. Aside from theoretical analysis, various empirical correlations that try to predict the overall noise level [Strahle and Shivashankara, 1975] or the spectral characteristics such as peak frequency, slope of rolloff in the high-frequency range, [Abugov and Obrezkov, 1978] etc. as a function of operating conditions were also obtained from experimental data. More recent studies have employed numerical simulations to predict combustion noise for open flames as well as complex geometries. A study by Ihme and Pitsch [2012] employs a model for predicting direct combustion noise generated by turbulent non-premixed flames where the Lighthill acoustic analogy is combined with a flamelet-based combustion model and incorporated into an LES simulation. Their predictions match well with experimental results although discrepancies were noted at high frequencies. The hybrid CHORUS method [Duran et al., 2014] predicts the noise output by performing LES of the combustion chamber, extracting the acoustic and entropy waves and then propagating these waves through the engine using analytical methods. Their results compare well with experiments.

The inverse problem of using noise to infer conditions inside the combustor is somewhat less well studied, although there has been a fair amount of research interest within the thermoacoustic community [Juniper and Sujith, 2018]. Simplifying the combustion noise generation process, Lieuwen [2005] uses the decay rate of autocorrelation to determine the stability margin of a combustor. Several subsequent studies apply tools from nonlinear

dynamics to combustion noise time-series and obtain useful precursors of instability. Gotoda et al. [2011] employ the Wayland test for non-linear determinism to show that when a system transitions to thermoacoustic instability, the combustion noise changes its character gradually from random and uncorrelated to completely deterministic. Similarly, Nair and Sujith [2014] show the disappearance of the multifractal signature of combustion noise as it transitions to instability and note that measures such as the Hurst exponent can serve as an early warning of thermoacoustic instabilities. A follow-up study by Godavarthi et al. [2017] looks at measures derived from recurrence networks as instability precursors. Kobayashi et al. [2019a] use a modified version of the permutation entropy to detect a precursor of the frequency-mode-shift in their staged aircraft engine model combustor before the amplification of pressure fluctuations. More recent work from the past year has also explored machine learning oriented approaches to the problem. Mondal et al. [2018] apply Hidden Markov Models to pressure timeseries for the early detection of instabilities in a Rijke tube, while Kobayashi et al. [2019b] and Hachijo et al. [2019] combine Support Vector Machines with complex networks and statistical complexity measures, respectively, to do the same in a swirl-stabilized combustor.

The utilization of combustion noise for diagnostic purposes has not been limited solely to the forecasting of thermoacoustic instabilities. Acoustic precursors from combustion noise have been identified for lean blowout in a pre-mixed flame by Nair and Lieuwen [2008] using the concentration of acoustic power in low-frequency bands, wavelet-filtered variance and thresholding techniques. Gotoda et al. [2011] study the dynamics of pressure fluctuations near lean blowout using permutation entropy, fractal dimensions, and short-term predictability. Murayama et al. [2019] have also used the weighted permutation entropy of combustion noise to develop precursors of blowout for their model gas turbine combustor.

5.1.2 The case for interpretable, Bayesian machine learning

A limitation of the approaches described above is that, by looking at the data through a handcrafted and predefined lens, one may miss other relevant information in the data. Machine learning techniques, on the other hand, find relevant functional relationships in the data without being influenced by researchers' preconceptions. Deployed correctly, they also use all available information in the data. The downside of a purely data-driven approach, however, is that it is only applicable to the specific system that generated the data.

To address this problem of limited portability and the fact that acoustic emissions are an imperfect source of information, the Bayesian machine learning technique we employ in this study provides principled measures of uncertainty in our predictions. We start with an appropriately vague prior belief about what the output of our model should be and as we

observe more data, we update this belief to obtain progressively tighter posterior distributions in accordance with Bayes' rule. This ensures that the model does not make overconfident predictions from out-of-distribution inputs which are entirely different from what it was trained on. These uncertainty estimates are particularly important when using a machine learning model in a critical device such as an aircraft engine. Bayesian machine learning techniques with correctly specified priors can also work with smaller amounts of data and are resistant to overfitting [Ghahramani, 2005]. They can also be used in continual learning without catastrophic forgetting [Li et al., 2019], which is particularly important if we want to keep learning from data throughout the operating lifetime of a device.

In this study, we use anchored ensembling, which is a simple and scalable way to train Bayesian neural networks [Pearce et al., 2018]. First, we perform a set of experiments on a Rijke tube driven by a swirling premixed turbulent flame. The power, equivalence ratio, fuel composition and the exit area of the tube are all varied so that noise data can be collected over a wide range of operating parameters. The decay rates of oscillations provoked by acoustic pulses are measured. The thermoacoustic behaviour of the combustor ranges from very stable to almost unstable. The challenge for our neural network ensembles is then to predict the power, the equivalence ratio and the measured decay rate using only a 300 ms sample of the (un-forced) combustion noise as their input. We select this challenge because conditions in an engine can change rapidly and decisions should be based on only the most recent sensor data history. This is a high-dimensional regression problem to which our neural network ensembles are perfectly suited.

A common criticism of machine learning techniques is that they are black-box models that are completely opaque to the user. To remedy this, we have used techniques known as Integrated Gradients [Sundararajan et al., 2017b] and permutation importance [Breiman, 2001] to reveal features in the acoustic spectrum that drive the predictions from our Bayesian neural network ensembles.

5.2 Experimental Setup

Figure 5.1 shows the experimental apparatus used in this study. An ordinary Bunsen burner is modified by attaching swirler vanes and a nozzle featuring a large central hole for the main flame and smaller surrounding holes for the pilot flames. A premixed mixture of methane and ethylene is used as fuel. This produces a noisy swirling premixed turbulent flame that is anchored over a wide range of operating conditions. The burner is placed inside a steel tube of length 800 mm and internal diameter 80 mm. Annular discs, with hole diameters 75,

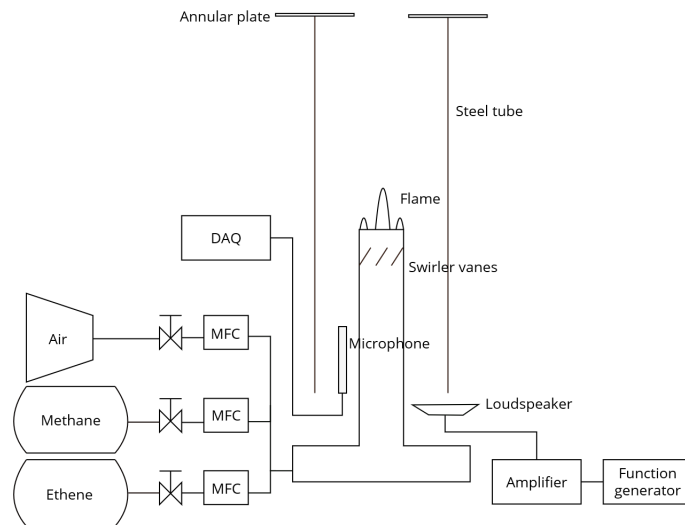


Fig. 5.1 Schematic of experimental setup, consisting of a 1 KW turbulent swirl flame inside a steel tube of length 800 mm and internal diameter 80 mm




Fig. 5.2 Left: Turbulent combustor excited by an acoustic pulse. Right: Filter signal with best-fit line in red.

65 and 55 mm, can be attached to the downstream end of the tube in order to change the acoustic boundary conditions.

The noise is recorded by a G.R.A.S. 26TK microphone placed near the bottom end of the tube. The raw pressure signal is sampled at 10000 Hz, which is considerably higher than the dominant frequencies in the typical noise spectra. Data acquisition is managed using a National Instruments BNC-2110 DAQ device and the software LabVIEW. Flow rates for fuel and air are controlled using Bronkhurst EL-FLOW® Select flowmeters. A 70 W VISATON 3020 BG loudspeaker is placed near the base of the burner to supply acoustic pulses.

The system is operated at 900 different combinations of operating parameters which form a grid in our 4-dimensional operating parameter space (volumetric flow rate, equivalence




Fig. 5.3 Equivalence ratios and flow rates of datapoints for methane:ethene ratio of 3:4 and outlet diameter 80 mm

ratio, fuel composition and outlet boundary condition). Experiments are performed at methane:ethene ratios of 3:4, 1:1 and 5:4 (v/v) and tube outlet diameters 80, 75 and 65 mm. For each of the 3 methane-ethene ratios and each of the 3 outlet boundary conditions, we methodically sweep through 100 different fuel and air mass flow rates. Figure 5.3 shows the 100 pairs of equivalence ratios and volumetric flow rates at which experiments are performed while the methane:ethene ratio and outlet diameter remained fixed at 0.75 and 80 mm, respectively. While this figure represents a 2D slice of the entire dataset— similar, but not identical, equivalence ratios and flow rates were achieved for the other fuel compositions and boundary conditions. These experiments mimic, in a laboratory setting, the multidimensional nature of the operating parameter space in a real jet engine where the boundary condition is typically dynamic and the engine controller has the authority to change multiple quantities such as fuel split, power, fuel inlet pressure, equivalence ratio, core speed and others. Any early warning signal needs to function over the whole range of operating parameters, not just when a single parameter is varied.

For each combination of operating parameters, the combustion noise is recorded and the decay rate of a 50 millisecond-long acoustic pulse at 230 Hz (the fundamental acoustic frequency of the system) is obtained. To extract the decay rate, the microphone signal is processed in a manner similar to Schumm et al. [1994]. First, a Butterworth filter, with a width of 20 Hz and centered at the excitation frequency 230 Hz, is used to filter out the undesired frequencies. Then a Hilbert transform is applied to obtain the instantaneous




Fig. 5.4 Contour plot of decay timescale (seconds) for methane:ethene ratio of 3:4 and outlet diameter 80 mm

amplitude, $A(t)$, of the pressure signal. When the logarithm of the obtained amplitude is plotted against time, it is possible to identify a linear region corresponding to exponential decay. To isolate the linear region, we ignore 50 ms of data immediately after the ping. The noise floor is computed from the RMS value of the pre-pulse signal and the decaying signal is cut-off when it decays to twice this value. The slope of this region then corresponds to the decay rate of the oscillations.

We use the measured decay rate (or its negative inverse, the decay timescale) at each operating point as a proxy for the thermoacoustic stability at that point. In general, decay rates or timescales quantify the linear stability of a system, which is a necessary but not sufficient condition for global thermoacoustic stability. However, for the high-amplitude instability we want to avoid, the linear stability boundary is observed to characterize the onset quite well. In Figure 7.1, we show a plot of the decay rate as a function of flow rate and equivalence ratio, for the same boundary condition and fuel composition as in Figure 5.3. (outlet diameter = 80 mm, methane-ethene ratio = 0.75) We observe that the decay rate only reaches values close to 0 and decay timescales reach their highest values (0.35 seconds for this particular subset of the data) in the vicinity of the high-amplitude 230 Hz instabilities. This holds true for the other 8 combinations of boundary conditions and fuel compositions as well. Therefore, if a diagnostic tool is well-correlated with the decay rate or timescale, it will be able to warn us when we are too close to the instability.

5.3 Statistical tools

5.3.1 Precursors of thermoacoustic instability from the literature

Nair and Sujith [2014] suggest that there is a loss of multifractality in combustion noise as combustors progress towards combustion instability, which is reflected in a decline of the signal's Hurst exponent H prior to an instability. The Hurst exponent is estimated using the Detrended Fluctuation Analysis technique introduced by Peng et al. [1995]. For some choice of window length τ , the signal is divided into segments of length τ from which linear trends are removed and the mean standard deviation of these segments, $\sigma(\tau)$, is computed. H is then defined by the scaling of this so-called fluctuation function, $\sigma(\tau)$ w.r.t. the window size τ . We obtain H by plotting the fluctuation function against window size on logarithmic axes and calculating its slope using least squares linear regression.

$$\sigma(\tau) \sim \tau^H \quad (5.1)$$

For our study, we calculate H using 1 second slices from the dynamic pressure sensor data (10000 data points) and 10 logarithmically spaced scaling window lengths between 0.01 and 0.02 seconds, which correspond to approximately two to four cycles of oscillations at combustion instability. Basic checks have been performed and it is found that a synthetic Gaussian white noise signal has a H close to 0.5 while a synthetic periodic signal at the instability frequency has a H close to 0, as expected.

Lieuwen [2005] derived an effective damping coefficient ξ_i for a thermoacoustic mode in terms of the decay rate ξ_i of the autocorrelation $C_i(\tau)$ of the i -th acoustic mode $\eta_i(t)$, assuming the combustor to be a second-order oscillator, the background noise to be spectrally flat and parametric disturbances to be absent:

$$C_i(\tau) = e^{-\omega_i \xi_i \tau} \left(\cos \left(\omega_i \tau \sqrt{1 - \xi_i^2} \right) + \frac{\xi_i}{\sqrt{1 - \xi_i^2}} \sin \left(\omega_i \tau \sqrt{1 - \xi_i^2} \right) \right) \quad (5.2)$$

The autocorrelation decay rate is calculated the same way as the decay rate of the acoustic pulses. A 1-second long sample of the raw pressure signal is first put through a Butterworth filter centered around the frequency of interest (here, 230 Hz) to obtain the signal $\eta_i(t)$. The autocorrelation $C_i(\tau)$ is then obtained as a function of the lag time τ and its envelope determined through the Hilbert transform. The slope of the least-squares fitted line through the logarithm of the autocorrelation amplitude then gives us the desired effective damping coefficient ξ_i . We expect this quantity to tend towards zero as our system approaches instability.

5.3.2 Bayesian neural network ensembles

In our study, we train a Bayesian neural network ensemble of 10 two-layer neural networks with 25 nodes in each layer and ReLU activation. The input to our network is the 51-dimensional power spectrum of 300 millisecond noise samples, computed using Welch's method by averaging the spectra of 1 millisecond segments from the sample with 0.95 milliseconds of overlap between segments. The outputs are the decay time scale (the negative inverse of the measured decay rate), the equivalence ratio, and power. Before training, all the input variables and outputs are normalized using a min-max strategy to lie between -1 and 1. Noise samples from 180 randomly chosen operating parameter combinations (20% of data) are held out in the test set for evaluating the performance of the model. 10-fold cross-validation is performed where 10 different models are trained using 10 random train-test splits. This ensures the stability of our algorithm's performance with respect to different train-test splits. Each ensemble member is trained using the stochastic gradient descent optimizer ADAM [Kingma and Ba, 2015]. The tunable hyperparameters of our model such as the learning rate, the data noise and the number of nodes in each layer are optimized by minimizing the negative log-likelihood of data in a validation set.

5.3.3 Interpretation using Integrated Gradients and Permutation Importance

To attribute the predictions of our network ensemble to the input features we use the technique of integrated gradients [Sundararajan et al., 2017b]. This is a simple, scalable method for attributing predictions to individual features that was described in Chapter 2. We use an implementation of Integrated Gradients from the DeepExplain library [Ancona et al., 2018].

Understanding network predictions via Integrated Gradients involves looking at individual examples and their attribution plots. To gain a global overview of the relative importance of input features, we also use the permutation importance technique ?. To compute permutation feature importances, we shuffle the values of a feature between samples in the dataset and measure the impact of randomizing them on the accuracy of a trained model. The average error of model predictions should increase significantly if important features are randomized in this way and the decrease in accuracy can therefore be understood as a measure of a feature's criticality. For our data, we shuffle each 100 Hz frequency block in the input spectra of the test data and measure the percentage increase of the Root Mean Squared Error.

The code and data used to produce the results in this chapter are available as a Google Colab notebook in the Github repository <https://github.com/Ushnish-Sengupta/FYR>. The notebook can be run in the web browser and does not require the installation of any software.

5.4 Results

We train an ensemble of neural networks that takes the power spectrum of a noise sample as input and predicts the negative inverse of the measured decay rate (the decay timescale). Figure 5.5 shows the performance of the ensemble on the test dataset, where we observe that the decay timescales are predicted reasonably accurately. In the course of our 10-fold cross-validation, the root mean squared prediction error ranged from 0.021 seconds to 0.024 seconds, indicating that our algorithm is stable to variations in the training-test split. It is particularly interesting to note how the grey errorbars (± 1 S.D.) widen for the operating points closer to instability, corresponding to larger decay timescales. This is because there are comparatively fewer data points close to instability (only 13 operating conditions in the training dataset have a decay timescale exceeding 0.3 seconds) making the ensemble less certain about its predictions in that region. This demonstrates how principled uncertainties prevent blind overconfidence in our machine learning models. Nevertheless, even with the slightly larger uncertainties and prediction errors for those points, this algorithm can clearly indicate when the system approaches thermoacoustic instability.

We also trained our ensembles to recognize the equivalence ratio and burner power from a noise sample. Figures 5.6 and 5.7 show the measured values of these state variables plotted against the predictions of our ensembles and reveal an even more accurate prediction. The root mean squared error in equivalence ratio prediction ranged between 0.031 to 0.033 ($\approx 3.5\%$), while the error for the power varied from 0.021 kW to 0.025 kW ($\approx 2\%$). The neural networks were thus able to predict these two important state variables quite accurately given a short noise sample. Each operating condition seem to have a unique acoustic signature which the machine learning algorithm can learn. In other words, one can indeed hear the state of a combustor.

For each ensemble of trained neural networks, we use the technique of integrated gradients to produce feature-level attribution plots, which can tell us how much a particular predictor influenced the prediction for a particular input. For example, the attribution plot in Figure 5.8 shows this technique applied to the decay rate prediction ensemble for an input with a decay timescale of 0.35 seconds. The most prominent features are the large positive attributions for the frequency component around the fundamental, which is marked in Figure 5.8 as a grey line labeled 1f, as well as the 3rd harmonic. The model has observed that an increased concentration of acoustic power around the fundamental frequency combined with lower powers around the third harmonic can be an indication that the system is close to instability. It is intriguing, however, that almost all of the predictor variables seem to make meaningful contributions to the final prediction and that if this information were removed, the accuracy of the predictions would deteriorate. This implies that for diagnostics to have higher predictive




Fig. 5.5 Plot of measured decay timescales vs decay timescales predicted by the neural network ensemble on the test data




Fig. 5.6 Plot of measured equivalence ratio vs equivalence ratio predicted by the neural network ensemble on the test data




Fig. 5.7 Plot of measured burner power vs burner power predicted by the neural network ensemble on the test data

powers, the data should be considered in its entirety and not examined through a pre-defined lens.

Our argument in favour of considering information from the entirety of the combustion noise spectrum is bolstered by the permutation feature importance plot (Figure 5.9). Here, a larger increase in the Root Mean Squared Error when a feature is randomized indicates a strong dependence of the model on this feature. We observe that the higher frequency portions of the input spectra, particularly near the third harmonic, have unique information that is independent of other parts of the spectra and randomizing them increases the RMSE on the test data by more than 20 percent, in some cases.

To compare our technique to those in the literature, the Hurst exponent and the decay of autocorrelation amplitude are computed for noise samples from each operating condition. While the Hurst exponent has a rough negative correlation with the measured decay rates, falling as the decay timescales grow, the relationship is very noisy (Figure 7.2). This means that, if the Hurst exponent were to be used to forecast instabilities for our combustors, there would be many false alarms, which is clearly undesirable.

The same is found for the autocorrelation decay (see Figure 5.11), which also becomes close to zero, as expected, at the edge of instability. However, it also approaches zero for several operating points at which the combustor is very stable, making it somewhat unreliable as a prognostic for instability. While these measures are attractive because they do not need to be trained on extensive experimental data, they also seem to have limited predictive power in our experiments.

To summarize, precursors of thermoacoustic instability are useful because they can be used by an engine controller to avoid triggering instabilities. In this chapter, we use Bayesian ensembles of neural networks to model relationships between measured combustion noise and the stability margins or operating conditions of a lab-scale turbulent combustor. We show that we can estimate the decay rate of acoustic pulses from a single 300 millisecond sample of the combustor's radiated noise over a wide range of operating conditions. We also show that the operating parameters (equivalence ratio and power) could be recovered from the noise data using the neural networks. This demonstrates that every operating point had a different sound and that a Neural Network could recognise the operating point just from that sound. A human may suspect this but would be unable to remember them all. Not only are our estimates reasonably accurate, but they contain principled estimates of uncertainty. This means that the model has the wisdom to "know what it doesn't know" and this is reflected in higher uncertainties when a provided input is too different from those on which it has been trained. With Integrated Gradients, we discern which features in the input spectra drive the networks towards a particular prediction, making our technique interpretable. We




Fig. 5.8 Top: Min-max normalized power spectrum of a noise sample corresponding to a 0.43 second decay timescale (close to instability) Bottom: Integrated gradients attribution plot for this power spectrum




Fig. 5.9 Permutation feature importance for decay timescale prediction




Fig. 5.10 Plot of generalized hurst exponent H vs decay timescale




Fig. 5.11 Plot of autocorrelation decay vs decay timescale

compare our approach with two precursors of thermoacoustic instability from the literature: the Hurst exponent and the autocorrelation decay. While these broadly behave as expected, their relationship to the measured decay rates was noisy.

Chapter 6

Prognosing instabilities in a rocket engine combustor

The work in this chapter has been accepted for publication in the International Journal of Spray Dynamics and Combustion conditional on implementing revisions suggested by reviewers.

6.1 Introduction

High-frequency thermoacoustic instabilities continue to hinder the development of liquid propellant rocket engines. This chapter proposes the diagnosis of impending instabilities using multiple sensor data-streams as input and Bayesian neural networks for forecasting. The forecast can be used by a rocket engine controller to perform a mass flow rate adjustment and avoid the unstable regions.

Most papers in the literature identify indicators of approaching instability. In this chapter, on the other hand, we use nonlinear autoregressive time series modeling with a Bayesian Neural Network to forecast, with uncertainties, the future amplitude of pressure fluctuations given the history of sensor signals and future flow rate control signals. This informs the engine controller of the expected timing and potential severity of a forthcoming instability. This approach also lets us take advantage of the rich instrumentation of our rig by integrating multimodal sensor data into our model: high frequency dynamic pressure ($p'(t)$) and OH* chemiluminescence ($I'(t)$) measurements, injector pressures (p_{H_2} , p_{O_2}) and temperatures (T_{H_2} , T_{O_2}), chamber static pressure (p_{cc}) and fuel flow rates (\dot{m}_{O_2} , \dot{m}_{H_2}). We compare models with different combinations of sensors, different sensor-derived features, and different lengths of signal histories in the inputs. We find that, although operating parameters and mass flow

rate control signals can anticipate instabilities on their own, inclusion of dynamic pressure signal history, in particular, results in a measurable increase in the forecast accuracy. It is not possible to predict triggered thermoacoustic instabilities from operating parameters and control signals alone, so we need to integrate multimodal sensor data with operating parameters and control signals in an instability forecasting framework.

The size of our dataset is limited by the high cost and effort required for each additional test run on the rocket thrust chamber. Fortunately, the uncertainty-aware nature of the Bayesian Neural Network allows us to work in this medium-data regime without overconfident extrapolation. We compute log-likelihoods on the test dataset and find that the uncertainty is well-quantified. We also simulate a sensor failure to assess the robustness of our Bayesian forecasting tool. As expected, the network exhibits large epistemic uncertainties during the failure, signalling the unfamiliar out-of-distribution nature of the inputs. Finally we use the interpretation technique of integrated gradients (IG) to understand how the different input features in our model influence the predictions.

6.2 Experimental Setup

The dataset used in this study is derived from experiments performed on the research combustor BKD [Gröning et al., 2016] operated at the P8 test facility (Fröhlke et al. [1997], Koschel et al. [1996]) of the DLR Institute of Space Propulsion in Lampoldshausen. It has three main components: an injector head, a cylindrical combustion chamber and a convergent-divergent nozzle, as shown in Figure 1. The combustion chamber is water-cooled and is designed to deal with the high thermal loads that are expected during instability events. The L42 injector head has 42 shear coaxial injectors and is operated with a liquid oxygen (LOX)/hydrogen (LH2) propellant combination. The cylindrical combustion chamber is 80 mm in diameter and the nozzle throat diameter is 50 mm, resulting in a contraction ratio of 2.56. The chamber static pressure (p_{cc}) is varied between 50 – 80 bar and mixture ratio of oxidizer to fuel ($ROF = \dot{m}_{O_2}/\dot{m}_{H_2}$) between 2 – 6. The experiments considered here have a LOX injection temperature T_{O_2} around 110 K and a hydrogen injection temperature T_{H_2} around 100 K.

For the operating point with $p_{cc} = 80$ bar, $ROF = 6$, the total propellant mass flow rate is 6.7 kg/s, the theoretical thermal power is 90 MW and the thrust achieved is about 24 kN. These specifications place BKD at the lower end of small upper stage engines.

A representative BKD test sequence, along with a spectrogram of dynamic pressure oscillations inside the combustion chamber, is shown in Figure 6.2. Stable and unstable operating conditions can be identified in the spectrogram. Strong high-frequency combustion




Fig. 6.1 Experimental thrust chamber BKD.

instabilities of the first tangential (1T) mode at about 10 kHz were excited consistently when approaching the operating point at $p_{cc} = 80$ bar, $ROF = 6$. The coupling mechanism for this instability is injection-driven [Gröning et al., 2016]. The flame dynamics are modulated by the LOX post acoustics and combustion instabilities emerge when the frequency of the 1T chamber mode matches the second longitudinal eigenmodes of the LOX posts. This mechanism has been confirmed using high-speed flame imaging by Armbruster et al. [2019].

The extreme conditions within the thrust chamber, where temperatures reach 3600 K and the pressure reaches 80 bar, limit the diagnostic instruments for our instability investigations. A specially designed measurement ring is placed between the injector head and the cylindrical combustion chamber segment, as shown in Figure 6.1. At this location, the temperatures are moderated by the injection of cryogenic propellants, which allows the mounted sensors to survive several test runs. Eight Kistler type 6043A water-cooled high-frequency piezoelectric pressure sensors are flush-mounted in the ring with an even circumferential distribution, in order to measure the chamber pressure oscillations $p'(t)$. The high-frequency pressure sensors have a measurement range set to ± 30 bar and a sampling rate of 100 kHz. An anti-aliasing filter with a cutoff frequency of 30 kHz is applied. Three fibre-optical probes are used to record the OH* radiation intensity $I'(t)$ of selected individual flames. The full acceptance angle of the optical probes is approximately 2° . In order to capture the OH* radiation signal, the probes are equipped with interference filters with a centre wavelength of 310 nm. The sampling frequency of the $I'(t)$ signals is also 100 kHz. There are also several




Fig. 6.2 BKD test sequence and p' spectrogram showing self-excited instability of the first tangential mode [Gröning et al., 2016]. p_{cc} is the chamber static pressure, ROF is the mixture ratio of oxidizer to fuel, and T_{H_2}, T_{O_2} are the LOX and hydrogen injection temperatures respectively.

low sampling frequency (100 Hz) sensors that measure the chamber static pressure p_{cc} , the mass flow rates of fuel \dot{m}_{H_2} and oxidizer \dot{m}_{O_2} and injector temperatures T_{H_2}, T_{O_2} .

6.3 Methods

6.3.1 Nonlinear autoregressive models with exogenous variables (NARX) for timeseries modeling

To forecast instabilities, we train NARX-style models Lin et al. [1996] to capture the functional relationship between the history of the sensor data, control signals and future pressure fluctuations. These model the evolution of pressure fluctuations as follows:

$$\log(p'_{\text{rms}}(t_m + \Delta t)) = f(\mathbf{x}_m, \mathbf{x}_{m-1}, \dots, \mathbf{x}_{m-h}, \mathbf{u}_{m+f}) + \sigma_\epsilon \quad (6.1)$$

Here, $p'_{\text{rms}}(t_m + \Delta t)$ is the root mean square amplitude of the dynamic pressure fluctuation signal in a 50 ms window centered around $t_m + \Delta t$. We choose $\Delta t = 500$ ms for our models because we consider this to be sufficient warning for a rocket engine controller to react and take appropriate action. The state variables $\mathbf{x}_m, \mathbf{x}_{m-1}, \dots, \mathbf{x}_{m-h}$ consist of various sensor data (or transformations thereof) at times $t_m - 50, t_m - 100, \dots, t_m - 50h$ ms. The order of dependence h determines how far in the past we look. In this study, we explore models with $h = 0$ and $h = 2$. The elements of \mathbf{u}_{m+f} are the two mass flux control signals $\dot{m}_{H_2, \text{set}}$ and $\dot{m}_{O_2, \text{set}}$ averaged over the window between t_m and $t_m + \Delta t$. These are the exogenous variables whose future values are known and can therefore be included in the inputs. It is possible to train models that capture higher order dependence on instantaneous values of the control signals instead of reducing them to an average, but we find that this does not improve predictive accuracy in our problem. σ_ϵ is the homoskedastic (constant) observation noise or aleatoric uncertainty term, which is assumed to be Gaussian. We forecast the logarithm of $p'_{\text{rms}}(t_m + \Delta t)$ instead of $p'_{\text{rms}}(t_m + \Delta t)$ itself because the noise in the pressure fluctuations increases with the amplitude, and so the simplifying assumption of constant homoskedastic noise only holds approximately once this logarithmic transform has been applied. It also forces p'_{rms} predictions to be positive.

We model the nonlinear function f in Equation (2) as a feedforward neural network with a simple multi-layer perceptron (MLP) architecture. In a standard neural network, one obtains a point estimate of the network parameters. Here, however, we train a machine learning algorithm on a dataset whose size is limited by the cost of performing experiments on large rocket thrust chambers. With a standard neural network this size limit could lead to naive and inaccurate extrapolations when the network encounters out-of-distribution inputs that

are different from those which were experienced during training. In safety-critical systems such as rockets, control decisions based on overconfident predictions could be catastrophic. Bayesian Neural Networks (BNNs), described in Chapter 2, are an elegant solution to this problem because they provide a confidence interval as well as a prediction.

6.3.2 Interpretation using Integrated Gradients

Neural networks have a reputation as black boxes, which disincentivizes their application to cases in which practitioners must understand why the algorithm is or is not working. We use the technique of integrated gradients (IG) [Sundararajan et al., 2017a] to attribute the predictions of our network ensemble to the input features, as described in Chapter 2. We use an implementation of Integrated Gradients from the DeepExplain library [Ancona et al., 2018].

6.3.3 Transforming sensor signals into neural network inputs

We compare NARX models with different combinations and transformations of the sensor data in their state variables \mathbf{x}_{m-i} (Equation 6.1) and evaluate their accuracy.

Models that include the low frequency signals $p_{cc}, \dot{m}_{H_2}, \dot{m}_{O_2}, T_{H_2}$ and T_{O_2} in their \mathbf{x}_{m-i} vectors use the average value of the signal between $t_m - 50i$ and $t_m - 50(i + 1)$. For the high frequency pressure and OH* signals, p' and I' , we use transformations from the literature that are known to capture instability-relevant information as described next. The current p'_{rms} is a natural candidate which is included in all models that incorporate pressure information. Sengupta et al. [2020c] and McCartney et al. [2021] have shown that, before a combustor transitions to instability, the power spectral density shows measurable changes which can be used as a precursor. We estimate the spectral density of the p' (p'_{FFT}) and I' (I'_{FFT}) time series by applying Welch's method Welch [1967] to a segment of the signals between $t_m - 50i$ ms and $t_m - 50(i + 1)$ ms. In this study, a Hann window of length 30 with 50% overlap is used for both p' and I' .

Kobayashi et al. [2019b] applied machine learning to the the ordinal partition transition network of the p' and I' timeseries to detect instabilities early. The ordinal partition transition networks for an m -dimensional time series are expressed by a weighted adjacency matrix W consisting of $W_{ij} = P(\pi_i \rightarrow \pi_j)$, $i, j \in [1, 2^m]$, $\sum W_{ij} = 1$, where $P(\pi_i \rightarrow \pi_j)$ is the observed probability from the i th- to j th-order transition patterns. In this study, the order patterns are π_1, π_2, π_3 , and π_4 which account for all combinations of the signs of the increments for our two-dimensional time series $[p'(t), I'(t)]$. The probability distribution of the corresponding

16-dimensional transition patterns W computed from the segment of the signals between $t_m - 50i$ and $t_m - 50(i + 1)$ are used in the state variable vector.

We also include a naive baseline model $p'_{\text{rms}}(t_m + \Delta t) = p'_{\text{rms}}(t_m)$ for comparison. This model, by design, is incapable of forecasting any future instabilities but, because large portions of the run consist of steady operation, its predictions have a reasonably good root mean squared error (RMSE). This model exists to provide a baseline RMSE and any useful forecast needs to be significantly more accurate than this.

6.3.4 Train-test-validation split, hyperparameters, performance metrics

We use a dataset of five experimental runs: three that are 50 seconds long and two that are 90 seconds long. Each run has two occurrences of the 1T instability, where an instability event is defined as being when the peak-to-peak amplitude of the pressure fluctuations exceeds 5% of the chamber static pressure.

To avoid data leakage, one experimental run is exclusively reserved for hyperparameter tuning. The timeseries is split into contiguous 250 millisecond blocks and every fifth block is used to evaluate the negative log-likelihood, which is the metric we use for tuning the BNN hyperparameters. The hyperparameters are the network depth, layer width, noise σ_ϵ , the learning rate for the ADAM optimizer, the number of epochs and the number of neural networks in the ensemble. For simplicity, and to facilitate comparison between models, the same hyperparameters are chosen for all models. This does not compromise the performance of any particular model, since increasing the key hyperparameters, width and depth, benefits all models. We choose a 3 hidden layer MLP with 100 units in each layer because larger networks show only minimal gains in performance. A learning rate of 2×10^{-4} is found to be optimal for convergence. The aleatoric noise parameter σ_ϵ is set to 0.07. In anchored ensembling, we train all members of the ensemble until convergence so the number of epochs was set to 512, which is found sufficient for all models. Converged estimates of the test data log-likelihood are obtained using 25 neural networks per ensemble.

To evaluate model performance, leave-one-out cross validation (LOOCV) is performed on the remaining four runs. Time series data involves strong temporal correlations, so assigning training and test data points randomly is dishonest because then the two sets would be highly similar. The test set must therefore be a completely independent experimental run. Additionally, LOOCV is necessary because with only four runs, firm conclusions cannot be drawn based on performance metrics for one particular test-train split, since any claimed

efficacy or differences between models could be caused by chance. We report average performance metrics across the four possible splits.

To evaluate model accuracy, we compute the root mean squared error (RMSE) of our pressure amplitude forecast for the entire test run (full RMSE). We also calculate this on a subset of the test run comprising the two segments of the timeseries one second before and after a transition to instability event (transition RMSE). The RMSE is computed separately on this subset because our primary interest is in forecasting instabilities and the relative merits of different models become clearer on this subset where there is a dramatic change in the amplitude. As a measure of the quality of uncertainty, we report mean negative log-likelihood per test data point.

6.4 Results

The models were trained on a laptop with 16 GB RAM, an Intel i7-10870H CPU and an NVIDIA RTX 2070 GPU. Each NARX model took ~ 2.5 hours to train. The inference time of the ensemble averaged around 10 milliseconds per sample, indicating that the algorithm is suitable as a real-time diagnostic tool.

Table 1 shows the average prediction RMSEs for the NARX models computed during cross-validation across the four runs. We note, firstly, that the NARX models outperform the naive baseline, which means that the data contains a predictable signal that can be extracted. We also observe that the simple models, which use only the two operating point parameters p_{cc} and ROF as inputs, perform well. This is unsurprising because we know that the 1T instabilities are consistently triggered when the operating point at $p_{cc} = 80$ bar and $ROF = 6$ is approached. This means that knowing the current operating point and future flow rate control signals is informative. All transitions to instability in the test runs differ from each other, however, even though they have identical control trajectories. The cross-validation shows unambiguously that the data from the additional instrumentation improves the prediction of future pressure fluctuations. Models that include the dynamic pressure spectrum p'_{FFT} , radiation intensity spectrum I'_{FFT} , current amplitude p'_{rms} , ordinal network transition probabilities, W , T_{O_2} , T_{H_2} , p_{O_2} or p_{H_2} , have better prediction accuracies, especially on the transition subset. Including higher order dependencies on state variables also provides a slight boost in accuracy. Interestingly, the dynamic pressure spectrum appears more informative than the spectrum of the OH* radiation intensity. This could be due to the fact that a single optical probe observes only a small volume of the chamber, while the pressure field integrates information across the whole chamber.

State variable inputs \mathbf{x}	Order	Full RMSE	Transition RMSE	NLL
Baseline	–	0.1552	0.2134	–
p_{cc} , ROF	0	0.1289	0.1630	-1.622
p_{cc} , ROF	2	0.1202	0.1576	-1.631
p'_{rms} , p'_{FFT}	0	0.1197	0.1546	-1.5752
p'_{rms} , p'_{FFT}	2	0.1169	0.1508	-1.5467
p'_{rms} , I'_{FFT}	0	0.1255	0.1612	-1.5154
p'_{rms} , W	0	0.1241	0.1566	-1.5801
T_{O2} , T_{H2} , p_{O2} , p_{H2}	0	0.1243	0.1578	-1.5948
p_{cc} , ROF, p'_{rms} , p'_{FFT} , T_{O2} , T_{H2} , p_{O2} , p_{H2}	0	0.1048	0.1407	-1.5632
p_{cc} , ROF, p'_{rms} , p'_{FFT} , T_{O2} , T_{H2} , p_{O2} , p_{H2}	2	0.1005	0.1311	-1.5350

Table 6.1 Cross-validation RMSEs (in bar) for entire runs (Full RMSE) and the transition subset of the run (Transition RMSE) and mean negative log-likelihoods (NLL) for different NARX models. Lower RMSEs and NLLs are better.

Table 1 also reports the mean negative log-likelihood (NLL) per datapoint, averaged across the four cross-validations. The mean negative log-likelihoods are fairly low for most of our models, implying that very few observations lie far outside the uncertainty bounds of the forecast (Figure 4). This is also confirmed by the percentage of data points in the test set within the ± 1 standard deviation (s.d.) and ± 3 s.d. bounds, which are roughly in line with our Gaussian assumptions (68% within 1 s.d., 99.7% within 3 s.d.). Additionally, though models using more features have higher prediction accuracies, they have slightly worse NLLs in some cases because the presence of additional features results in higher epistemic uncertainties.

To test the robustness of the Bayesian model against out-of-distribution inputs, we simulate a sensor failure event. This is shown in Figure 6.4. To simulate an optical probing event between 10 sec and 40 sec, the original photomultiplier signal is replaced by white noise with mean 0.05 and amplitude 0.015. During this simulated probe failure, the epistemic uncertainty of the Bayesian neural network greatly increases, indicating that the inputs are highly dissimilar to conditions experienced during training. Predictions remain reasonable despite the uninformative optical signal, although they become much less accurate.

We use the technique of integrated gradients to produce feature-level attribution plots, which can tell us how much a particular predictor influenced the prediction for a particular input. For example, the attribution plots in Figure 6.5 shows this technique being applied to two data points. The first one is a data point just prior to the first instability event in a run, where a large jump in the dynamic pressure fluctuation amplitude is forecast by the model. This model uses p'_{FFT} , p'_{rms} and the flowrate control signals as input. The most prominent feature is the large positive attribution for the oxygen flowrate control signal \dot{m}_{O2} , whose

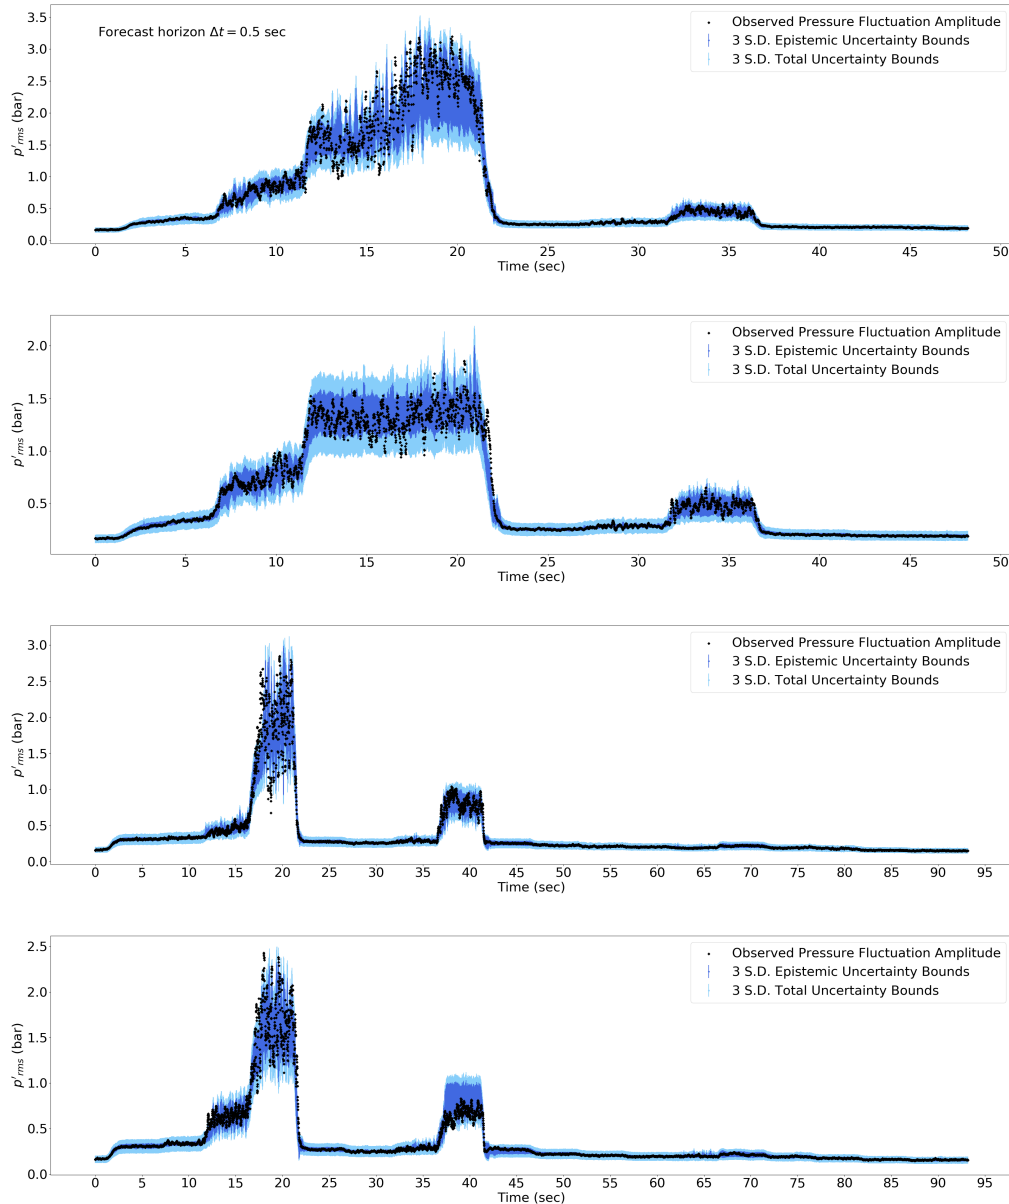


Fig. 6.3 Forecast p'_{rms} values for the four experimental runs based on sensor data prior to $t - 0.5$ sec, superimposed on observed p'_{rms} values at time t . The model was trained on the other three runs in each case. This model used p'_{FFT} and p'_{rms} as state variables and order $h = 2$. Actual observations are shown as black dots, 3 standard deviation (S.D.) total uncertainty bounds for the forecast are in light blue and 3 standard deviation (S.D.) epistemic uncertainty bounds are in deep blue.

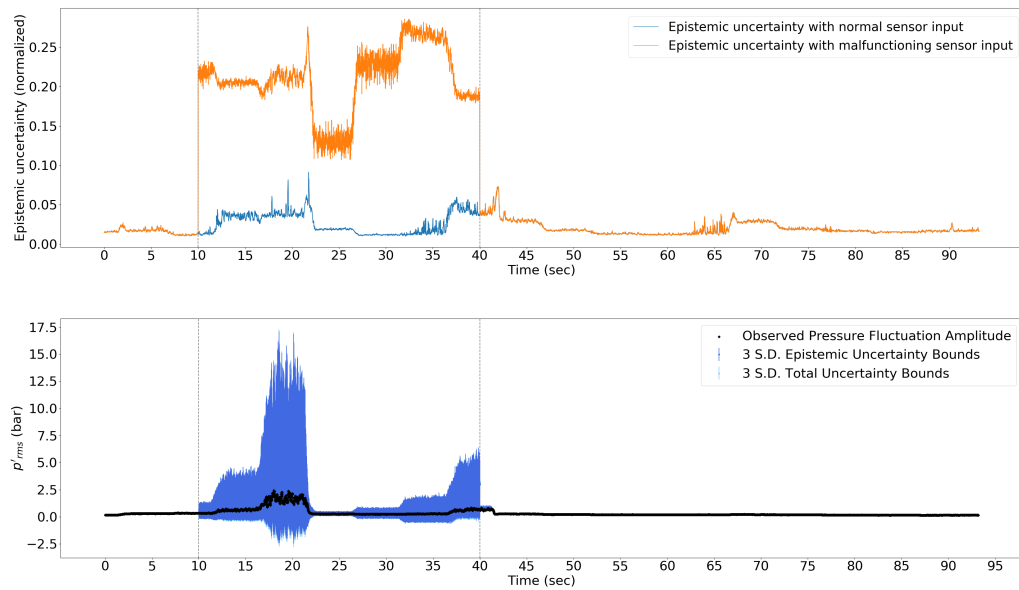


Fig. 6.4 Bayesian Neural Network epistemic uncertainty (top) and p'_{rms} forecasts (bottom) for experimental run no. 4 where a fibre optical probe failure was simulated between $t = 10$ sec and $t = 40$ sec (marked by two grey dotted lines). The model used I'_{FFT} , p'_{rms} as input state variables and order = 0. Actual observations are shown as black dots, 3 standard deviation (S.D.) total uncertainty bounds for the forecast are represented by light blue bars and 3 standard deviation (S.D.) epistemic uncertainty bounds are in deep blue. The light blue bars are barely discernible from the deep blue bars in the failure region because the total uncertainty is almost entirely dominated by the large epistemic uncertainty.

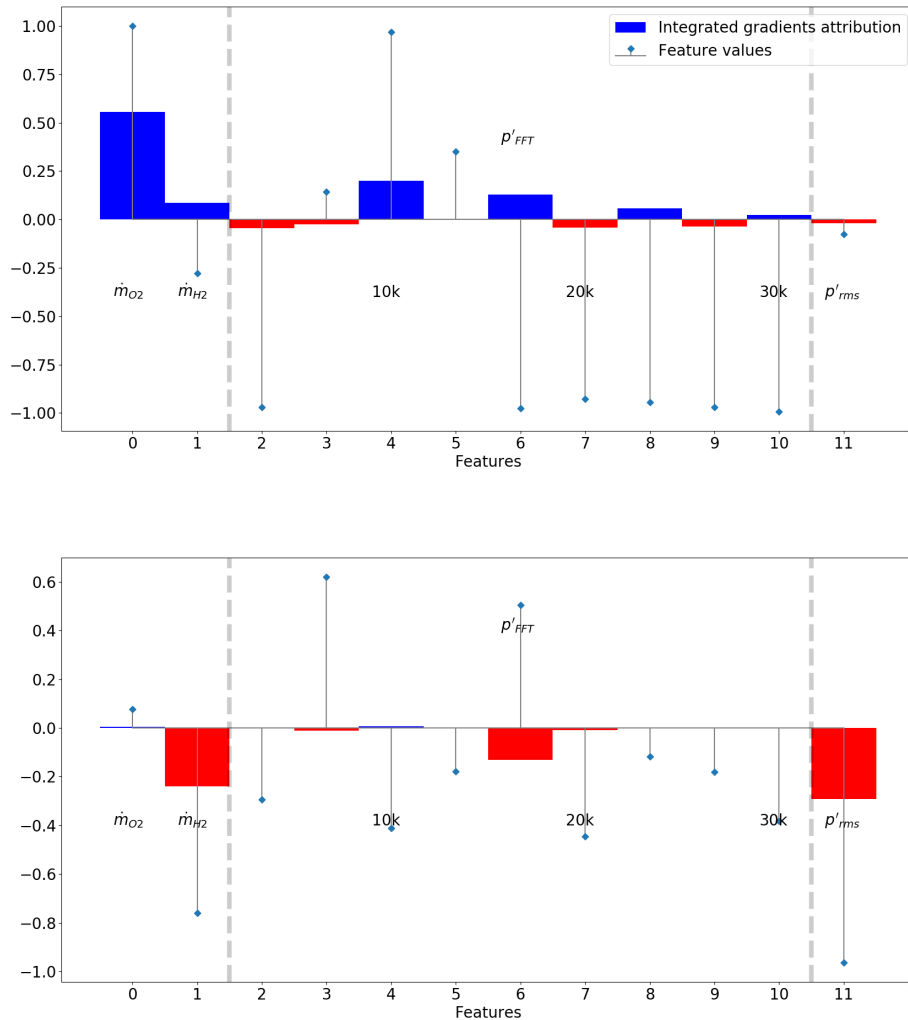


Fig. 6.5 Integrated gradient attributions and input feature values (normalized to have zero mean and unit standard deviation) for two datapoints from experimental run no. 4 at $t = 16.0$ seconds (top, immediately preceding an instability) and $t = 60.0$ seconds (bottom, stable). The model used p'_{FFT} and p'_{rms} as state variables and order = 0. The blue dots show the normalized values of the features for this particular datapoint and the bars represent the integrated gradient attributions for each feature. For example, feature 6 (spectral power around the 16.66 kHz frequency) in the top plot has a large negative value and a positive integrated gradients attribution, meaning that the network has learned that low power in this frequency band is a potential indicator of instability.

feature value is also large for this datapoint. This is expected, because a large \dot{m}_{O_2} would increase the p_{cc} and ROF, bringing the combustor closer to instability. The prediction is also modulated by contributions from the p'_{FFT} features. This datapoint has a concentration of power around the 10 kHz spectrum and low power in the mid to higher frequencies. There is a positive contribution from the 10 kHz frequency component of the spectrum, which is the frequency of the 1T mode, as well as from the 16.66k Hz frequency. This shows that the model has learned that an increased concentration of acoustic power around the frequency of the 1T mode and away from the mid-frequencies is an indication that the system is close to instability. The second data point is from stable operation. The attributions in this case have mostly negative values because the prediction of the model is strongly negative. The current pressure fluctuation amplitude p'_{rms} is close to -1.0 in this case and naturally has a large negative attribution. The \dot{m}_{H_2} input (here -0.78) also contributes negatively because low ROFs are stable. Finally, the high concentration of power in the mid-frequencies also tells the network that this operating point will stay stable.

To summarize, we presented a method that combines multiple sensory modalities in a rocket thrust chamber to predict impending thermoacoustic instabilities with uncertainties. This is accomplished by training an autoregressive Bayesian neural network model that forecasts the future amplitude of the dynamic pressure time series, using multiple sensor measurements (injector pressure/ temperature measurements, static chamber pressure, high-frequency dynamic pressure measurements, high-frequency OH* chemiluminescence measurements) and future flow rate control signals as input. The method is validated using experimental data from a representative cryogenic research thrust chamber. The Bayesian nature of our algorithms allows us to work with a dataset whose size is restricted by the expense of each experimental run, without making overconfident extrapolations. We find that the networks are able to accurately forecast the evolution of the pressure amplitude and anticipate instability events on unseen experimental runs 500 milliseconds in advance. We compare the predictive accuracy of multiple models using different combinations of sensor inputs. We find that the high-frequency dynamic pressure signal is particularly informative. We also use the technique of integrated gradients to interpret the influence of different sensor inputs on the model prediction. The negative log-likelihood of data points in the test dataset indicates that prediction uncertainties are well-characterized by our model and simulating a sensor failure event results in a dramatic increase in the epistemic component of the uncertainty, as would be expected when a Bayesian method encounters unfamiliar, out-of-distribution inputs.

Chapter 7

Prognosing instabilities in an intermediate pressure industrial fuel spray nozzle

The work in this chapter was done as part of a virtual industrial secondment with Rolls-Royce Aircraft Engines. The data was collected by Rolls-Royce engineers from their experimental combustion test rig and the analysis was performed by me.

7.1 Introduction

The third study on machine learning based instability precursors is on an industrial intermediate pressure combustion test rig, which is equipped with three pressure transducers, sampling at 50 kHz. Experiments are performed on three different fuel injectors over a range of operating points in order to identify operating points that are thermoacoustically unstable. The injectors are labelled $1a$, $1b$ and 2 . Injectors $1a$ and $1b$ are nominally identical. The operating points are identified by their air-fuel ratio (AFR) and their exit temperature (T_{30}). The threshold of thermoacoustic instability is defined as the operating points at which the acoustic amplitude exceeds 0.5% of the static pressure. The black lines in figure 10 show this threshold in (AFR, T_{30})-space. Despite being nominally identical, injectors $1a$ and $1b$ have instability thresholds at slightly different positions in (AFR, T_{30})-space.

The statistical tools used in this chapter have already been introduced in previous chapters, so we can proceed directly to a discussion of the results obtained.




Fig. 7.1 The black line shows the thermoacoustic instability threshold as a function of air-fuel ratio (AFR) and exit temperature T_{30} for three aeroplane engine fuel injectors in an intermediate pressure rig. The coloured lines show the distance to the black line. Injectors 1a and 1b are nominally identical.

7.2 Results

We normalize the ranges of AFR and T_{30} to run from 0 to 1 and then train a BayNNE to recognize the Euclidian distance to the instability threshold, based on 500 ms of normalized pressure measurements. Stable points are assigned positive distances and unstable points are assigned negative distances. We compare the predictions from the BayNNE with those from the autocorrelation decay, the permutation entropy, and the Hurst exponent. Figure 7.2 show the Hurst exponent, the autocorrelation decay, and the permutation entropy for injector 1a. The Hurst exponent reduces significantly as the system becomes unstable and this is a useful indicator of the instability threshold, albeit with significant unquantified uncertainty. The autocorrelation decay tends towards zero as the system becomes more unstable but, for this data, barely changes across the instability threshold and therefore does not provide a useful indicator for the threshold. The permutation entropy drops after the system has crossed the threshold from stable to unstable operation, meaning that it is not useful as a precursor in these experiments.

The BayNNE is trained on the training points of 1a and then applied to test data points of 1a, 1b and 2. Figure 7.3 shows the distance from instability threshold predicted by the BayNNE compared with the true distance. Uncertainty bands of the BayNNE are shown in grey. The BayNNE provides a remarkably accurate prediction of the distance to instability from the pressure signal alone. Figure 7.3(a) shows the distance from the instability threshold predicted by the BayNNE trained on injector 1a when applied to pressure data from the same injector and figure 7.3(b) shows the distance predicted by the BayNNE trained on

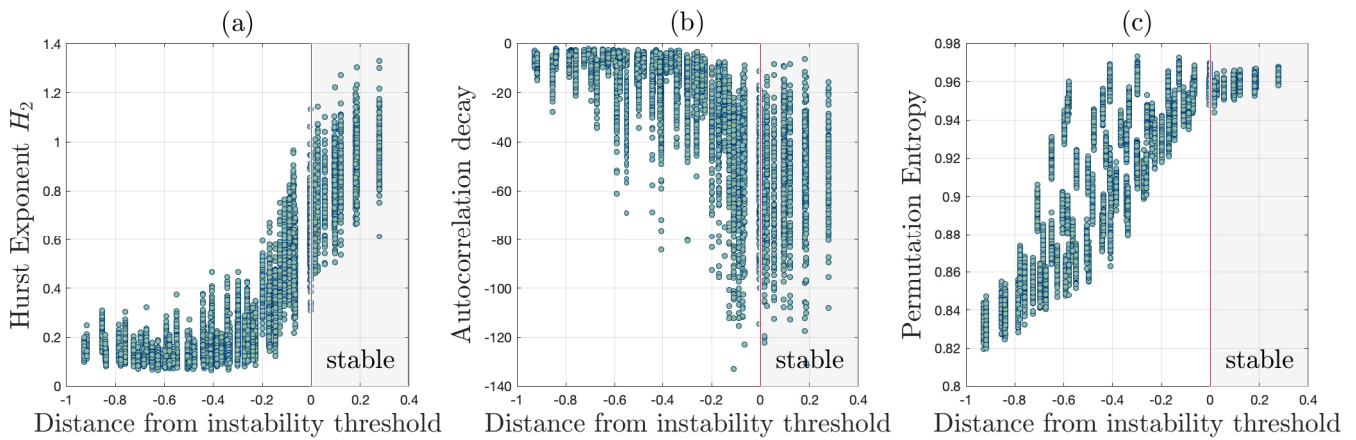


Fig. 7.2 The black line shows the thermoacoustic instability threshold as a function of air-fuel ratio (AFR) and exit temperature T_{30} for three aeroplane engine fuel injectors (a) $1a$ (b) $1b$ and (c) 2 in the intermediate pressure rig. The coloured lines show the distance to the black line. Injectors $1a$ and $1b$ are nominally identical.

injector $1a$ when applied to the nominally identical injector $1b$. The BayNNE predictably performs best when the test data originates from the same rig (prediction RMSE: 0.015) but is also fairly accurate when predicting distances for the $1b$ injector (RMSE: 0.034). As mentioned above, $1a$ is unstable over a different range of (AFR, T_{30})-space than $1b$ and, despite this difference, the BayNNE has successfully identified the distance to instability on the new injector and can therefore prognose instability. Figure 7.3 shows the distance from the instability threshold predicted by the BayNNE trained on $1a$ when applied to the pressure data from injector 2 . The BayNNE performs badly, particularly when the system is stable. However, the uncertainty in its predictions is also greater, reflecting the out-of-distribution nature of the input. This confirms that a BayNNE trained on one thermoacoustic system is a good indicator of thermoacoustic precursors on nominally-identical thermoacoustic systems, outperforming statistical measures, but is not useful for different systems. This is not surprising, given that the BayNNE is using all available information from the pressure signal of this particular system, while the statistical methods are quantifying general features in all systems.

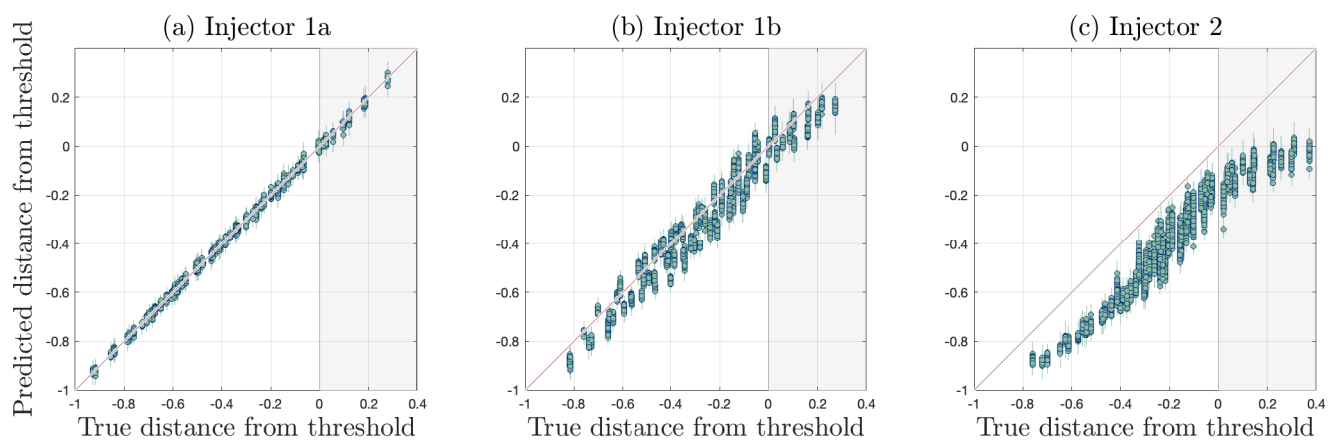


Fig. 7.3 Predicted distance to the instability threshold as a function of measured distance to instability threshold for (a) injector 1a (b) injector 1b (c) injector 2 when the prediction is obtained from a BayNNE trained on injector 1a. Injectors 1a and 1b are nominally identical.

Chapter 8

Conclusions and Outlook

8.1 Summary of contributions

Qualitatively accurate physical models can be augmented with data-driven probabilistic machine learning tools such as Gaussian processes and Bayesian neural networks to improve their quantitative accuracy and compute predictive uncertainties. This dissertation investigates several ways in which such tools can be of benefit when modeling, designing against and trying to avoid thermoacoustic instabilities.

The most influential element of any thermoacoustic system is the response of the flame to acoustic forcing. This is also the hardest element to model. In Chapter 3, we use a heteroscedastic Bayesian neural network ensemble to calibrate the parameters of the G-equation, a reduced-order model for predicting the acoustic response of premixed flames, based on observations. The ensemble is trained on 2.1 million simulated flame observations and results on the test set of simulated flames show that parameters are recovered accurately by the network from only 10 frames of flame data. Reasonable estimates of parameter uncertainties are also obtained. We then use the network to infer parameters from high-speed video footage of acoustically excited Bunsen flames in our lab. Re-simulating the flames using inferred parameters reveals a close agreement between the model output and experimental data, validating our data assimilation. Compared to the Ensemble Kalman filter used for this task in the literature, our neural network based technique can assimilate flame videos much more rapidly and obtain accurate parameter estimates using fewer observations.

Once a quantitatively accurate model of a thermoacoustic system is at hand, designers need to explore the parameter space of the model to find an optimally stable design. Our optimization algorithm should explore the design space efficiently using as few evaluations of the underlying model as possible, as simulating the model can be expensive. In Chapter 4, we demonstrate how scalable, gradient-augmented Bayesian optimization (BO) pairs

perfectly with adjoint methods in thermoacoustics. We use two optimization problems from the literature as our test cases: a 1D Helmholtz equation model of a Rijke tube and a low-order network model of a longitudinal combustor. We show that the BO algorithm consistently finds more stable parameter combinations using fewer iterations than BFGS, a popular quasi-Newton optimizer. It also builds a metamodel of the objective function in the entire design space and explores multiple feasible designs. The metamodel can naturally be used for uncertainty quantification as well.

Despite our best efforts to eliminate thermoacoustic instabilities in our combustor, there are often particular ranges of operating conditions which need to be avoided because they trigger instabilities. In Chapter 5, we use Bayesian ensembles of neural networks to model relationships between measured combustion noise and the stability margins or operating conditions of a lab-scale turbulent combustor. We show that we can estimate the decay rate of acoustic pulses, the power and the equivalence ratio from a single 300 millisecond sample of the combustor's radiated noise over a wide range of operating conditions. Not only are these estimates reasonably accurate, but they contain principled estimates of uncertainty. This means that the model has the wisdom to “know what it doesn't know” and this is reflected in higher uncertainties when a provided input is too different from those on which it has been trained. With Integrated Gradients, we discern which features in the input spectra drive the networks towards a particular prediction, making our technique interpretable. We compare our approach with two precursors of thermoacoustic instability from the literature: the Hurst exponent and the autocorrelation decay. While these broadly behave as expected, their relationship to the measured decay rates was noisy.

This study shows that a Bayesian ensemble of neural networks trained on a particular combustor is better at discerning the onset of thermoacoustic instability than are traditional measures based on the Hurst exponent and the autocorrelation decay. This shows that important information is lost when data is filtered through the traditional methods. The good agreement between predicted and measured values shown in Figures 5.5, 5.6 and 5.7 shows that each operating point has a distinctive noise, which can be learned by the ensemble. It is worth mentioning, of course, that this model is machine-specific and will not generalize to a different machine. In an industrial setting, however, this drawback is mitigated by the fact that there are only a few combustor designs and a lot of data for each combustor.

This work highlights the promise of building robust real-time early warning systems for instabilities such as thermoacoustic oscillations using combustion noise data.

Chapter 6 takes the work done in our lab and applies it to a real world system. We apply our uncertainty-aware and interpretable machine learning framework to forecast instabilities using multimodal sensor data in a realistic, thermoacoustically unstable rocket

thrust chamber. We predict the amplitude of dynamic pressure fluctuations 500 ms in advance using temperature, static pressure, high-frequency dynamic pressure and OH* radiation intensity data recorded during test runs performed with the cryogenic rocket thrust chamber BKD. We perform a rigorous cross-validation of the autoregressive Bayesian neural network and find that our models are able to anticipate all the instability events with varying accuracy (25 – 40% lower RMSEs on the transition subset compared to baseline), depending on the combination of inputs chosen. The inclusion of longer histories and additional sensor information, in general, boosts predictive accuracy. Models that use the dynamic pressure signal, in particular, show a measurable improvement in the forecast, confirming the findings of the literature on this topic.

Once again, a key benefit of our Bayesian approach is the predictive uncertainty. In safety-critical machines such as rocket engines, a diagnostic tool must be robust to unexpected events such as sensor failures. We compute the log-likelihood of the test data given model predictions and find that uncertainties are well-characterized. We also simulate a sensor failure event in a model whose inputs included the spectrum of the fibre-optical probe signal and demonstrate the robustness of Bayesian algorithms. Integrated Gradients helps us identify important features and how they influence a particular prediction of the neural network.

Chapter 7 applies these machine learning based instability precursors to pressure data from an industrial intermediate pressure combustion test rig. We used data from rigs with three different injectors, two of which were nominally identical. We found that a neural network ensemble trained on one of the identical injectors can predict the distance from instability when applied to data obtained from its twin, even though these systems display somewhat different thermoacoustic behaviour. The same network however fails to generalize to the different injector, confirming the machine-specificity of these statistical, physics-agnostic instability precursors. We also find that permutation entropy and the Hurst exponent fail to adequately prognose instabilities on this data, although the permutation entropy does a relatively better job.

8.2 Outlook

There are several promising avenues for future work in each of the applications considered.

For parameter inference using Bayesian neural networks, one can look at assimilating data from more realistic industrial flame configurations as well as integrating our calibrated reduced-order flame models into higher-level thermoacoustic network models of an entire combustor.

While Chapter 4 illustrates the untapped potential of gradient-augmented BO in thermoacoustics and other engineering disciplines where adjoint models are commonly used, the models we used in our test cases were relatively cheap to evaluate. BO really shines when the objective is expensive to evaluate, so a truly convincing use case would involve more expensive models, e.g. 3D adjoint Helmholtz solvers and a complex combustor geometry with many hundreds of parameters. Another interesting line of research would be looking at different acquisition functions in the BO routine. One acquisition function of particular interest is Noisy-Input Entropy Search [Fröhlich et al., 2020] which searches for robust optima that are less sensitive to noise in their input. Since thermoacoustic stability is highly sensitive to small perturbations in model parameters and geometry, finding more robust optima may lead to better designs.

Chapters 5, 6 and 7 highlight the promise of building robust real-time early warning systems for instabilities such as thermoacoustic oscillations using combustion noise data as well as other sensor data such as photomultipliers and thermocouples. It also shows how combustion noise can predict the power and equivalence ratio of the combustor and thus validate sensor measurements without any additional investment in hardware. We can seek to address the issue of machine-specificity in our machine learning models by training on a wider variety of datasets from different sources and using techniques from transfer learning and domain-invariant representation learning to build machine-invariant diagnostic tools, which will generalize better to new devices. It also remains to be seen how effective these diagnostic tools are at avoiding instabilities when implemented online. Hardware upgrades are currently being undertaken at the DLR test facility that will allow the implementation of these methods into the engine control algorithm.

Papers

These papers were produced as an outcome of the research I did during my PhD, but some of this work has not found its way into the thesis. This is either because the work was done in collaboration and I was not the primary contributor; or because the subject matter did not fit in with the rest of the thesis.

1. Sengupta, U., Amos, M., Hosking, S., Rasmussen, C. E., Juniper, M., Young, P. (2020). Ensembling geophysical models with Bayesian neural networks. *Advances in Neural Information Processing Systems*, 33, 1205-1217.
2. Waxenegger-Wilfing, G., Sengupta, U., Martin, J., Armbruster, W., Hardi, J., Juniper, M., Oswald, M. (2021). Early detection of thermoacoustic instabilities in a cryogenic rocket thrust chamber using combustion noise features and machine learning. *Chaos*, 31(6), 063128.
3. Croci, M. L., Sengupta, U., Juniper, M. P. (2021). Data Assimilation Using Heteroscedastic Bayesian Neural Network Ensembles for Reduced-Order Flame Models. In *International Conference on Computational Science* (pp. 408-419). Springer, Cham.
4. Sengupta, U., Rasmussen, C. E., Juniper, M. P. (2021). Bayesian machine learning for the prognosis of combustion instabilities from noise. *Journal of Engineering for Gas Turbines and Power*, 143(7).
5. Sengupta, U., Juniper, M. P. (2022). Thermoacoustic stabilization of combustors with gradient-augmented Bayesian optimization and adjoint models. *International Journal of Spray and Combustion Dynamics*. <https://doi.org/10.1177/17568277221109118>
6. McCartney, M., Sengupta, U., Juniper, M. P. (2022). Reducing uncertainty in the onset of combustion instabilities using dynamic pressure information and bayesian neural networks. *Journal of Engineering for Gas Turbines and Power*, 144(1).

7. Amos, M., Sengupta, U., Young, P., Hosking, J. S. (2022, accepted). A continuous vertically resolved ozone dataset from the fusion of chemistry climate models with observations using a Bayesian neural network. *Environmental Data Science*.
8. Sengupta, U., Waxenegger-Wilfing, G., Martin, J., Hardi, J., Juniper, M. P. (accepted for publication in the *International Journal of Spray and Combustion Dynamics*, pending revisions) Forecasting Thermoacoustic Instabilities in Liquid Propellant Rocket Engines Using Multimodal Bayesian Deep Learning. arXiv preprint: 2107.06396.
9. Sun, Y., Sengupta, U., Juniper, M.P. (submitted to the *Journal of Computational Physics*) Physics-informed Deep Learning for Simultaneous Surrogate Modelling and PDE-constrained Optimization. preprint DOI: 10.13140/RG.2.2.23162.57280

References

- Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for regression*. MIT, 1996.
- Matthew P Juniper. Sensitivity analysis of thermoacoustic instability with adjoint helmholtz solvers. *Physical Review Fluids*, 3(11):110509, 2018.
- Jose Guillermo Aguilar Perez. *Sensitivity analysis and optimization in low order thermoacoustic models*. PhD thesis, University of Cambridge, 2019.
- S. Gröning, J. S. Hardi, D. Suslov, and M. Oswald. Injector-driven combustion instabilities in a Hydrogen/Oxygen rocket combustor. *Journal of Propulsion and Power*, 32(3):560–573, 2016. doi: 10.2514/1.B35768.
- Rayleigh. The explanation of certain acoustical phenomena. *Nature*, 1878. ISSN 00280836. doi: 10.1038/018319a0.
- Boa Teh Chu. On the energy transfer to small disturbances in fluid flow (Part I). *Acta Mechanica*, 1965. ISSN 00015970. doi: 10.1007/BF01387235.
- Joseph C Oefelein and Vigor Yang. Comprehensive review of liquid-propellant combustion instabilities in f-1 engines. *Journal of Propulsion and Power*, 9(5):657–677, 1993.
- B Zinn, Y Neumeier, B Zinn, and Y Neumeier. An overview of active control of combustion instabilities. In *35th Aerospace Sciences Meeting and Exhibit*, page 461, 1997.
- Jakob J Keller. Thermoacoustic oscillations in combustion chambers of gas turbines. *AIAA journal*, 33(12):2280–2287, 1995.
- José G Aguilar and Matthew P Juniper. Thermoacoustic stabilization of a longitudinal combustor using adjoint methods. *Physical Review Fluids*, 5(8):083902, 2020.
- Tsubasa Kobayashi, Shogo Murayama, Takayoshi Hachijo, and Hiroshi Gotoda. Early detection of thermoacoustic combustion instability using a methodology combining complex networks and machine learning. *Physical Review Applied*, 11(6):064034, 2019a.
- M. P. Juniper and R. I. Sujith. Sensitivity and nonlinearity of thermoacoustic oscillations. *Annual Review of Fluid Mechanics*, 2018. doi: 10.1146/annurev-fluid-122316-045125.
- Tim Lieuwen. Online Combustor Stability Margin Assessment Using Dynamic Pressure Data. *Journal of Engineering for Gas Turbines and Power*, 2005. ISSN 07424795. doi: 10.1115/1.1850493.

- Vineeth Nair, Gireehkumaran Thampi, Sulochana Karuppusamy, Saravanan Gopalan, and R. I. Sujith. Loss of chaos in combustion noise as a precursor of impending combustion instability. *International Journal of Spray and Combustion Dynamics*, 5(4):273–290, 2013. doi: 10.1260/1756-8277.5.4.273. URL <https://doi.org/10.1260/1756-8277.5.4.273>.
- Hiroshi Gotoda, Hiroyuki Nikimoto, Takaya Miyano, and Shigeru Tachibana. Dynamic properties of combustion instability in a lean premixed gas-turbine combustor. *Chaos*, 2011. ISSN 10541500. doi: 10.1063/1.3563577.
- Vineeth Nair, Gireeshkumaran Thampi, and RI Sujith. Intermittency route to thermoacoustic instability in turbulent combustors. *Journal of Fluid Mechanics*, 756:470, 2014.
- Vineeth Nair and R. I. Sujith. Multifractality in combustion noise: predicting an impending combustion instability. *Journal of Fluid Mechanics*, 747:635–655, may 2014. ISSN 0022-1120. doi: 10.1017/jfm.2014.171. URL <https://www.cambridge.org/core/product/identifier/S0022112014001712/type/journal{ }article>.
- Soumalya Sarkar, Satyanarayanan R Chakravarthy, Vikram Ramanan, and Asok Ray. Dynamic data-driven prediction of instability in a swirl-stabilized combustor. *International Journal of Spray and Combustion Dynamics*, 8(4):235–253, 2016.
- Meenatchidevi Murugesan and RI Sujith. Detecting the onset of an impending thermoacoustic instability using complex networks. *Journal of Propulsion and Power*, 32(3):707–712, 2016.
- Devesh K Jha, Nurali Virani, and Asok Ray. Markov modeling of time series via spectral analysis for detection of combustion instabilities. In *Handbook of Dynamic Data Driven Applications Systems*, pages 123–138. Springer, 2018.
- Sudepta Mondal, Najah F Ghalyan, Asok Ray, and Achintya Mukhopadhyay. Early detection of thermoacoustic instabilities using hidden markov models. *Combustion Science and Technology*, 2018.
- Takayoshi Hachijo, Shinga Masuda, Takuya Kurosaka, and Hiroshi Gotoda. Early detection of thermoacoustic combustion oscillations using a methodology combining statistical complexity and machine learning. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(10):103123, 2019. doi: 10.1063/1.5120815. URL <https://doi.org/10.1063/1.5120815>.
- Tsubasa Kobayashi, Shogo Murayama, Takayoshi Hachijo, and Hiroshi Gotoda. Early detection of thermoacoustic combustion instability using a methodology combining complex networks and machine learning. *Phys. Rev. Applied*, 11:064034, Jun 2019b. doi: 10.1103/PhysRevApplied.11.064034. URL <https://link.aps.org/doi/10.1103/PhysRevApplied.11.064034>.
- Michael McCartney, Thomas Indlekofer, and Wolfgang Polifke. Online prediction of combustion instabilities using machine learning. *Journal of Engineering for Gas Turbines and Power*, 2020.

- Tryambak Gangopadhyay, Vikram Ramanan, Adedotun Akintayo, Paige K Boor, Soumalya Sarkar, Satyanarayanan R Chakravarthy, and Soumik Sarkar. 3d convolutional selective autoencoder for instability detection in combustion systems. *Energy and AI*, 4:100067, 2021. ISSN 2666-5468. doi: <https://doi.org/10.1016/j.egyai.2021.100067>. URL <https://www.sciencedirect.com/science/article/pii/S2666546821000215>.
- Sudepta Mondal, Ashesh Chattopadhyay, Achintya Mukhopadhyay, and Asok Ray. Transfer learning of deep neural networks for predicting thermoacoustic instabilities in combustion systems. *Energy and AI*, page 100085, 2021. ISSN 2666-5468. doi: <https://doi.org/10.1016/j.egyai.2021.100085>. URL <https://www.sciencedirect.com/science/article/pii/S2666546821000392>.
- Ann P Dowling and Aimee S Morgans. Feedback control of combustion oscillations. *Annu. Rev. Fluid Mech.*, 37:151–182, 2005.
- Iain Murray Zoubin Ghahramani. A note on the evidence and Bayesian Occam’s razor. *Gatsby Unit Technical Report*, 2005. doi: <http://eprints.pascal-network.org/archive/00001165/01/occam.pdf>.
- HongLin Li, Payam Barnaghi, Shirin Enshaeifar, and Frieder Ganz. Continual learning using bayesian neural networks, 2019.
- F. A. Williams. Turbulent combustion. *The mathematics of combustion*, pages 97–131, 1985.
- H. Yu, T. Jaravel, M. Ihme, M. P. Juniper, and L. Magri. Data assimilation and optimal calibration in nonlinear models of flame dynamics. *Journal of Engineering for Gas Turbines and Power*, 141, 2019a.
- Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.
- David Eriksson, Kun Dong, Eric Hans Lee, David Bindel, and Andrew Gordon Wilson. Scaling gaussian process regression with derivatives. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 6868–6878, 2018.
- Robert Dale. Gpt-3: What’s it good for? *Natural Language Engineering*, 27(1):113–118, 2021.
- Deng Huang, Theodore T Allen, William I Notz, and Ning Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of global optimization*, 34(3):441–466, 2006.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight Uncertainty in Neural Networks. may 2015. URL <http://arxiv.org/abs/1505.05424>.
- Hippolyt Ritter, Aleksandar Botev, and David Barber. A scalable laplace approximation for neural networks. In *6th International Conference on Learning Representations, ICLR 2018- Conference Track Proceedings*, volume 6. International Conference on Representation Learning, 2018.

- Tim Pearce, Mohamed Zaki, Alexandra Brintrup, Nicolas Anastassacos, and Andy Neely. Uncertainty in Neural Networks: Bayesian Ensembling. oct 2018. URL <http://arxiv.org/abs/1810.05546>.
- Diederik P Kingma and Jimmy Lei Ba. Adam: A method for stochastic gradient descent. *ICLR: International Conference on Learning Representations*, 2015.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR, 2017a.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 14. MIT Press, 2006. ISBN 978-0-262-18253-9. URL <http://www.gaussianprocess.org/gpml/>.
- Radford M. Neal. *Bayesian Learning for Neural Networks*, volume 118 of *Lecture Notes in Statistics*. Springer New York, New York, NY, 1996. ISBN 978-0-387-94724-2. doi: 10.1007/978-1-4612-0745-0. URL <http://link.springer.com/10.1007/978-1-4612-0745-0>.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Warren Scott, Peter Frazier, and Warren Powell. The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression. *SIAM Journal on Optimization*, 21(3):996–1026, 2011.
- Jacob Gardner, Geoff Pleiss, Ruihan Wu, Kilian Weinberger, and Andrew Wilson. Product kernel interpolation for scalable gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 1407–1416. PMLR, 2018.
- Daniel James Lizotte. *Practical bayesian optimization*. University of Alberta, 2008.
- Paul G Constantine, Eric Dow, and Qiqi Wang. Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM Journal on Scientific Computing*, 36(4): A1500–A1524, 2014.
- Ziyu Wang, Masrour Zoghi, Frank Hutter, David Matheson, Nando De Freitas, et al. Bayesian optimization in high dimensions via random embeddings. In *IJCAI*, pages 1778–1784, 2013.
- A. Fuller, Z. Fan, C. Day, and C. Barlow. Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8:108952–108971, 2020.
- Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 2020. ISSN 0027-8424. doi: 10.1073/pnas.1912789117. URL <https://www.pnas.org/content/early/2020/05/28/1912789117>.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538, 2015.

- Stefan T Radev, Ulf K Mertens, Andreass Voss, Lynton Ardizzone, and Ullrich Köthe. Bayesflow: Learning complex stochastic models with invertible neural networks. *arXiv preprint arXiv:2003.06281*, 2020.
- J. W. Strutt. *The Theory of Sound (Vol II)*. Macmillan and Co., 1878.
- S. Hemchandra. *Dynamics of Turbulent Premixed Flames in Acoustic Fields*. PhD Thesis, 2009.
- H. Yu, M. P. Juniper, and L. Magri. Combined state and parameter estimation in level-set methods. *J. Comp. Phys.*, 399, 2019b.
- U. Sengupta, M. Amos, J. Scott Hosking, C. E. Rasmussen, M. P. Juniper, and P. J. Young. Ensembling geophysical models with bayesian neural networks. *34th Conference on Neural Information Processing Systems (NeurIPS), Vancouver, Canada, 2020a*.
- U. Sengupta, M. L. Croci, and M. P. Juniper. Real-time parameter inference in reduced-order flame models with heteroscedastic bayesian neural network ensembles. *Machine Learning and the Physical Sciences Workshop at the 34th Conference on Neural Information Processing Systems (NeurIPS), Vancouver, Canada, 2020b*.
- K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *IEEE International Conference on Computer Vision (ICCV)*, pages 1026–1034, 2015.
- T. Pearce, M. Zaki, A. Brintrup, N. Anastassacos, and A. Neely. Uncertainty in neural networks: Bayesian ensembling. *International Conference on Artificial Intelligence and Statistics*, 2020.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *31st Conference on Neural Information Processing Systems (NIPS), Long Beach, CA, USA, 2017*.
- Matthias Katzfuss, Jonathan R Stroud, and Christopher K Wikle. Understanding the ensemble kalman filter. *The American Statistician*, 70(4):350–357, 2016.
- S. Hochreiter and J Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997.
- Franz Joos. Thermoakustik von brennkammern. *Technische Verbrennung: Verbrennungstechnik, Verbrennungsmodellierung, Emissionen*, pages 737–774, 2006.
- L Magri and MP Juniper. Sensitivity analysis of a time-delayed thermo-acoustic system via an adjoint-based approach. *Journal of Fluid Mechanics*, 719:183–202, 2013.
- Georg A Mensah and Jonas P Moeck. Acoustic damper placement and tuning for annular combustors: an adjoint-based optimization study. *Journal of Engineering for Gas Turbines and Power*, 139(6), 2017.
- Camilo F Silva, Luca Magri, Thomas Runte, and Wolfgang Polifke. Uncertainty quantification of growth rates of thermoacoustic instability by an adjoint helmholtz solver. *Journal of Engineering for Gas Turbines and Power*, 139(1), 2017.

- Stefano Falco and Matthew P Juniper. Shape optimization of thermoacoustic systems using a two-dimensional adjoint helmholtz solver. *Journal of Engineering for Gas Turbines and Power*, 143(7):071025, 2021.
- Georg A Mensah, Alessandro Orchini, and Jonas P Moeck. *Adjoint-based computation of shape sensitivity in a Rijke-tube*. Universitätsbibliothek der RWTH Aachen, 2019.
- Luca Magri. Adjoint methods as design tools in thermoacoustics. *Applied Mechanics Reviews*, 71(2), 2019.
- Henry Jones, Renaud Gaudron, and Aimee S Morgans. Shape optimisation of an unstable combustor using a genetic algorithm. In *Symposium on Thermoacoustics in Combustion: Academia meets Industry*, 2021.
- R Balachandran, BO Ayoola, CF Kaminski, AP Dowling, and E Mastorakos. Experimental investigation of the nonlinear response of turbulent premixed flames to imposed inlet velocity oscillations. *Combustion and Flame*, 143(1-2):37–55, 2005.
- Ann P. Dowling and Yasser Mahmoudi. Combustion noise. *Proceedings of the Combustion Institute*, 2015. ISSN 15407489. doi: 10.1016/j.proci.2014.08.016.
- S. Zelditch. Spectral determination of analytic bi-axisymmetric plane domains. *Geometric and Functional Analysis*, 10(3):628–677, sep 2000. ISSN 1016-443X. doi: 10.1007/PL00001633. URL <http://link.springer.com/10.1007/PL00001633>.
- Mark Kac. Can One Hear the Shape of a Drum? *The American Mathematical Monthly*, 1966. ISSN 00029890. doi: 10.2307/2313748.
- I. Duran, S. Moreau, F. Nicoud, T. Livebardon, E. Bouty, and T. Poinso. Combustion noise in modern aero-engines. *AerospaceLab Journal*, 2014.
- A. A. Putnam. Combustion Roar as Observed in Industrial Furnaces. *Journal of Engineering for Power*, 2009. ISSN 00220825. doi: 10.1115/1.3227357.
- M J Lighthill. Sound generated aerodynamically. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 1961.
- Warren C. Strahle. On combustion generated noise. *Journal of Fluid Mechanics*, 1971. ISSN 14697645. doi: 10.1017/S0022112071002167.
- F. E. Marble and S. M. Candel. Acoustic disturbance from gas non-uniformities convected through a nozzle. *Journal of Sound and Vibration*, 1977. ISSN 10958568. doi: 10.1016/0022-460X(77)90596-X.
- Warren C. Strahle and B. N. Shivashankara. A rational correlation of combustion noise results from open turbulent premixed flames. *Symposium (International) on Combustion*, 1975. ISSN 00820784. doi: 10.1016/S0082-0784(75)80397-3.
- D. I. Abugov and O. I. Obrezkov. Acoustic noise in turbulent flames. *Combustion, Explosion, and Shock Waves*, 1978. ISSN 15738345. doi: 10.1007/BF00789719.

- Matthias Ihme and Heinz Pitsch. On the Generation of Direct Combustion Noise in Turbulent Non-Premixed Flames. *International Journal of Aeroacoustics*, 11(1):25–78, mar 2012. ISSN 1475-472X. doi: 10.1260/1475-472X.11.1.25. URL <http://journals.sagepub.com/doi/10.1260/1475-472X.11.1.25>.
- V Godavarthi, VR Unni, EA Gopalakrishnan, and RI Sujith. Recurrence networks to study dynamical transitions in a turbulent combustor. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(6):063113, 2017.
- Suraj Nair and Tim Lieuwen. Acoustic Detection of Blowout in Premixed Flames. *Journal of Propulsion and Power*, 2008. ISSN 0748-4658. doi: 10.2514/1.5658.
- Shogo Murayama, Kentaro Kaku, Motoi Funatsu, and Hiroshi Gotoda. Characterization of dynamic behavior of combustion noise and detection of blowout in a laboratory-scale gas-turbine model combustor. *Proceedings of the Combustion Institute*, 2019. ISSN 15407489. doi: 10.1016/j.proci.2018.07.034.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic Attribution for Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*, mar 2017b. URL <http://arxiv.org/abs/1703.01365>.
- Michael Schumm, Eberhard Berger, and Peter A. Monkewitz. Self-excited oscillations in the wake of two-dimensional bluff bodies and their control. *Journal of Fluid Mechanics*, 1994. ISSN 14697645. doi: 10.1017/S0022112094001679.
- C-K Peng, Shlomo Havlin, H Eugene Stanley, and Ary L Goldberger. Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series. *Chaos: an interdisciplinary journal of nonlinear science*, 5(1):82–87, 1995.
- Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. Towards better understanding of gradient-based attribution methods for deep neural networks. 01 2018.
- K Fröhlke, A Haberzettl, O Haidn, S Heinrich, M Sion, P Vuillermoz, K Froehlke, A Haberzettl, O Haidn, S Heinrich, et al. First hot fire test campaign at the french/german research facility p8. In *33rd Joint Propulsion Conference and Exhibit*, page 2929, 1997.
- W Koschel, J Haidn, and G Krülle. P8-the new french/german test facility for h2/o2-high pressure rocket engine combustion research. In *11th World Hydrogen Energy Conference, Stuttgart, June 23-28, 1996*, pages 1815–1825. DECHEMA, Frankfurt, Germany, 1996.
- W. Armbruster, J. S. Hardi, D. Suslov, and M. Oswald. Injector-driven flame dynamics in a high-pressure multi-element oxygen-hydrogen rocket thrust chamber. *Journal of Propulsion and Power*, 35:632–644, 2019. doi: 10.2514/1.B37406.
- Tsungnan Lin, B.G. Horne, P. Tino, and C.L. Giles. Learning long-term dependencies in narx recurrent neural networks. *IEEE Transactions on Neural Networks*, 7(6):1329–1338, 1996. doi: 10.1109/72.548162.
- Ushnish Sengupta, Carl Rasmussen, and Matthew Juniper. Bayesian machine learning for the prognosis of combustion instabilities from noise. *Journal of Engineering for Gas Turbines and Power*, 2020c.

Michael McCartney, Ushnish Sengupta, and Matthew Juniper. Reducing uncertainty in the onset of combustion instabilities using dynamic pressure information and bayesian neural networks. In *Proceedings of ASME Turbo Expo 2021*. ASME, 2021.

Peter Welch. 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):70–73, 1967.

Lukas Fröhlich, Edgar Klenske, Julia Vinogradska, Christian Daniel, and Melanie Zeilinger. Noisy-input entropy search for efficient robust bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 2262–2272. PMLR, 2020.