

Adjoint-accelerated Bayesian inference in thermoacoustics



Matthew Yoko

Supervisor: Prof. M.P. Juniper

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Trinity College

January 2024

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this report 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 dissertation contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Matthew Yoko
January 2024

Abstract

We demonstrate a new approach to thermoacoustic modelling, in which we use an efficient Bayesian inference framework to assimilate experimental data into thermoacoustic models. The framework provides three main tools: (i) parameter inference, (ii) model comparison, and (iii) optimal experiment design. With Bayesian parameter inference, we use experimental observations to infer the most probable parameters of a thermoacoustic model. At the same time, we quantify the uncertainty in the parameters and the model. If there are several plausible models, we use Bayesian model comparison to quantitatively rank the models to select the best one. Parameter inference and model comparison can require a lot of data if the experiments are not designed well. With Bayesian optimal experiment design, we use the model to identify the most informative experiments, reducing the experimental cost and effort.

Bayesian inference is often considered too computationally expensive to be applied to problems in fluid dynamics. This is because many Bayesian inference frameworks use sampling techniques to construct the posterior, which require thousands of model evaluations. This is not practical if the model evaluations are expensive, as is usually the case in fluid dynamics. We reduce this cost significantly in two ways. Firstly, we use an approximate Bayesian inference framework called Laplace's method, which reduces the task of computing the posterior to a quadratic optimization problem. We then solve this optimization problem with the fewest model evaluations by using gradient-based optimization, with the gradients calculated using adjoint methods.

We apply this framework to three simple thermoacoustic systems: an electrically heated Rijke tube, a ducted laminar conical flame, and a ducted turbulent conical flame. With the electrically heated Rijke tube, we demonstrate the full set of tools in the Bayesian framework. Using these tools, we create a thermoacoustic model that (i) is quantitatively accurate over the entire operating range, (ii) has quantified uncertainty bounds, (iii) can extrapolate beyond the observed data, and (iv) can be trained on a few optimal experiments. This is a significant improvement on previous attempts to model the electrically heated Rijke tube in the literature.

With both sets of ducted flames, we demonstrate a second use for Bayesian inference in thermoacoustics: obtaining flame transfer functions, and their uncertainties, from pressure observations. We do this with the flame *in-situ*, without a need for optical access. If the flame's response is sensitive to its environment, which is usually the case, then this is preferable to measuring flame transfer functions *ex-situ* using optical methods.

This thesis provides a proof of concept by applying adjoint-accelerated Bayesian inference to a set of canonical problems in thermoacoustics. In doing so, we show that the Bayesian framework provides a powerful set of tools for combining the work of the experimentalist and the modeller in a mutually beneficial way. The resulting models are accurate, interpretable, and capable of extrapolation. The models can be trained with only a few experimental observations, making Bayesian inference feasible for applications where experiments are expensive.

This framework is currently being implemented on the Rolls–Royce SCARLET industrial test rig, which will investigate how the tools handle more complex thermoacoustic systems. Work is also being done on applying the framework to assimilate video footage of flames into models of the flame dynamics. This will allow us to predict flame transfer functions at operating conditions and frequencies that have not been observed experimentally. This, combined with the acoustic network models presented in this thesis, would form a powerful design tool that could be used to inform more elegant and robust interventions to thermoacoustic instability with fewer prototyping iterations.

Table of contents

List of figures	xi
List of tables	xv
Nomenclature	xvii
Thesis structure	1
I Introduction	3
1 Thermoacoustics	5
1.1 Historical background	5
1.2 Modelling thermoacoustic oscillations	7
2 Bayesian Inference	11
2.1 Adjoint-accelerated Bayesian inference	12
2.1.1 Parameter inference	12
2.1.2 Uncertainty quantification	14
2.1.3 Model selection	19
2.2 Bayesian optimal experiment design	20
2.2.1 Optimal design for parameter inference	22
2.2.2 Optimal sensor placement for parameter inference	24
2.2.3 Optimal design for model comparison	25
2.3 Summary	26

II	Applications in thermoacoustics	29
3	Hot wire Rijke tube	31
3.1	Introduction	32
3.1.1	Experimental configuration	34
3.1.2	Experimental procedure & data processing	36
3.1.3	Experiment design	38
3.1.4	Physics-based model of the hot wire Rijke tube	39
3.2	Generating a quantitatively accurate model of a hot wire Rijke tube	41
3.2.1	Calibrating the base flow model	42
3.2.2	Characterizing the cold acoustics	44
3.2.3	Inferring the fluctuating heat release rate at the heater	52
3.2.4	Extrapolation	56
3.2.5	Summary	58
3.3	Minimizing the data required to train a model of a hot wire Rijke tube	59
3.3.1	Optimal design for parameter inference	59
3.3.2	Optimal sensor placement for parameter inference	67
3.3.3	Optimal design for model comparison	71
3.3.4	Note on application	74
3.4	Conclusion	75
4	Ducted laminar conical flames	79
4.1	Introduction	80
4.1.1	Experimental configuration	82
4.1.2	Experimental procedure & data processing	84
4.1.3	Experiment design	85
4.1.4	Thermoacoustic model of the ducted laminar conical flame	86
4.2	Assimilating data into the thermoacoustic model	87
4.2.1	Characterization of the cold rig	88
4.2.2	Assimilating heat release rate models	91
4.2.3	Conclusion	104
5	Ducted turbulent conical flames	107
5.1	Introduction	108
5.1.1	Experimental configuration	108

5.1.2	Direct measurements of the fluctuating heat release rate	109
5.1.3	Flame properties	110
5.1.4	Thermoacoustic model of the ducted flame rig	112
5.2	Inferring flame transfer functions of turbulent conical flames	112
5.2.1	Characterizing the cold rig	113
5.2.2	Assimilating heat release rate models from pressure data	114
5.2.3	Validation of inferred fluctuating heat release rate	119
5.3	Conclusion	121
 III Conclusion		123
 6 Adjoint-accelerated Bayesian inference in thermoacoustics		125
6.1	Summary of main results	125
6.2	Outlook	127
 References		131
 Appendix A Properties of the flames studied in Chapters 4 and 5		145
 Appendix B Photographs of the experimental rigs		147

List of figures

2.1	Parameter inference using gradient-based optimization.	14
2.2	Uncertainty quantification using Laplace’s method.	16
2.3	Systematic uncertainty estimation on a toy problem.	18
2.4	Model comparison on a toy problem: predictions of three candidate models.	20
2.5	Model comparison on a toy problem: model ranking metrics.	21
3.1	Diagram of the electrically heated Rijke tube rig.	35
3.2	Diagram of the electric heater.	35
3.3	Raw and processed signals from a forced response test.	38
3.4	Model diagram for the hot wire Rijke tube.	42
3.5	Predictions of the calibrated base flow model.	43
3.6	Posterior probability distributions of R and η for the hot wire Rijke tube.	46
3.7	Performance of seven models for the drag and blockage of the heater support prongs.	47
3.8	Performance of the posterior model of the cold Rijke tube.	51
3.9	Local feedback coefficients for the heat release rate fluctuations of the hot wire, inferred using two different models.	53
3.10	Performance of a model of the hot wire Rijke tube based on other studies in the literature.	55
3.11	Performance of the model of the hot wire Rijke tube constructed from the component models with the highest marginal likelihoods.	57
3.12	As for Fig. 3.11 but assimilating data from just six of the 120 hot operating points.	57
3.13	Three steps of optimal experiment design for inferring the visco-thermal dissipation of the cold heater.	60

3.14	Learning rate when applying optimal experiment design for inferring the visco-thermal dissipation of the cold heater.	62
3.15	Four steps of optimal experiment design for inferring the fluctuating heat release rate.	64
3.16	Learning rate when applying optimal experiment design for inferring the fluctuating heat release rate.	67
3.17	Model predictions after three stages of optimal sensor placement.	68
3.18	Posterior joint distributions after three stages of optimal sensor placement. . .	69
3.19	Comparison of uncertainty collapse when using the best and worst sensor placements.	70
3.20	Learning rate when applying optimal sensor placement.	71
3.21	Predictions of three candidate models for the fluctuating heat release rate. . .	72
3.22	Sequential assimilation of best and worst experiments for model comparison. . .	73
4.1	Diagram of the ducted laminar conical flame rig.	82
4.2	Extracting the linear growth rate from non-linear growth.	86
4.3	Processed steady flame images from the 24 laminar conical flames.	87
4.4	Model diagram for the ducted laminar conical flame.	88
4.5	Characterization of the cold ducted laminar flame rig.	90
4.6	Prior and posterior joint parameter probability distributions for the cold ducted laminar flame rig.	91
4.7	Diagrams of models for the fluctuating heat release rate of a ducted laminar conical flame.	93
4.8	Predictions of four flame transfer function models for the ducted laminar conical flame.	94
4.9	Model ranking metrics for four flame transfer function models for the ducted laminar conical flame.	96
4.10	Inferred uncertainties for four flame transfer function models for the ducted conical flame.	97
4.11	Experimental data collected from 24 laminar conical flames.	98
4.12	Posterior model predictions of the thermoacoustics of 24 laminar conical flames - view 1.	99
4.13	Posterior model predictions of the thermoacoustics of 24 laminar conical flames - view 2.	100

4.14	Polar plot of the inferred flame transfer functions for all 24 laminar conical flames.	101
4.15	Comparison of inferred flame transfer functions with those from the literature.	103
5.1	Diagram of turbulent ducted flame rig.	109
5.2	Images of a perturbed turbulent flame.	110
5.3	Properties of the 15 turbulent conical flames.	111
5.4	Model diagram for the ducted turbulent conical flame.	112
5.5	Characterization of the cold ducted turbulent conical flame rig.	114
5.6	Prior and posterior joint parameter probability distributions for the cold ducted turbulent flame rig.	115
5.7	Experimental <i>vs</i> predicted eigenvalues for 15 turbulent flames.	116
5.8	Inferred flame transfer functions for 15 turbulent conical flames (Bode plot). .	117
5.9	Inferred flame transfer functions for 15 turbulent conical flames (Polar plot). .	118
5.10	Direct measurements <i>vs</i> inferred fluctuating heat release rates.	119
B.1	Photograph of the hot wire Rijke tube.	147
B.2	Photograph of the electric heater.	148
B.3	Photograph of the ducted flame rig.	149

List of tables

3.1	Summary of the experiments conducted on the hot wire Rijke tube	39
3.2	Comparison of seven models for the heater support prongs.	49
3.3	Comparison of three models for the drag and blockage of the heater.	51
3.4	Comparison of 9 models for the hot heater.	54
5.1	Summary of the flow properties of the 15 turbulent conical flames.	111
A.1	Properties of the 24 laminar conical flames.	145
A.2	Properties of the 15 turbulent conical flames.	146

Nomenclature

Roman Symbols

a	Parameter vector
A	Area
<i>c</i>	Speed of sound
C_{aa}	Covariance matrix describing parameter variance
C_{ee}	Covariance matrix describing experimental variance
C_{ss}	Covariance matrix describing model prediction variance
C_{tt}	Covariance matrix describing total variance
D	Diameter
D_{KL}	Kullback-Leibler divergence
<i>\mathcal{F}</i>	Flame/heater transfer function
F	Forward travelling wave (frequency domain)
<i>f</i>	Forward travelling wave (time domain)
G	Backward travelling wave (frequency domain)
<i>g</i>	Backward travelling wave (time domain)
H	Hessian tensor
<i>\mathcal{H}</i>	Candidate model
I	Identity matrix
I	Spatially integrated image intensity
<i>i</i>	Imaginary unit
Im(•)	Imaginary part of a complex number
J	Jacobian matrix
<i>\mathcal{J}</i>	Cost function
J	Bessel function of the first kind
<i>k_{ep}</i>	Linear feedback from pressure into the energy equation
<i>k_{eu}</i>	Linear feedback from velocity into the energy equation

k_{mp}	Linear feedback from pressure into the momentum equation
k_{mu}	Linear feedback from velocity into the momentum equation
L	Length
\mathcal{N}	Normal (Gaussian) distribution
n	Interaction index of the n - τ model
P	Fourier-decomposed pressure
p	Pressure
$p(\bullet)$	Probability
Pr	Prandtl number
Q	Heat release rate
R	Reflection coefficient
Re(\bullet)	Real part of a complex number
Re	Reynolds number
\mathbf{s}	Model prediction vector
S	Shannon entropy
s_i	Angular frequency
s_r	Growth/decay rate
St	Strouhal number
t	Time
tr	Trace of a matrix
U	Expected utility
u	Velocity
$u(\bullet)$	Utility function
V	Variance of a univariate Gaussian distribution
\mathbf{x}	Experiment design variable vector
X	Axial position
\mathbf{z}	Data vector
S_L	Unstretched laminar flame speed

Greek Symbols

η	Visco-thermal dissipation strength
γ	Ratio of specific heats
μ	Mean of a univariate Gaussian distribution or viscosity
ρ	Density

σ	Shear number, $\sigma = 1/2D\sqrt{\rho s_i/\mu}$
τ	Time delay of the n - τ model
τ_w	Shear stress at a wall

Superscripts

MP	Most probable
----	---------------

Other Symbols

$ \bullet $	Magnitude of a complex number
$\angle(\bullet)$	Phase angle of a complex number
$\bar{\bullet}$	Mean quantity
\bullet'	Fluctuating quantity
$\hat{\bullet}$	Relative fluctuating quantity, $\hat{\bullet} = \bullet'/\bar{\bullet}$

Acronyms / Abbreviations

BFL	Best fit likelihood
CFD	Computational fluid dynamics
DAQ	Data acquisition system
LES	Large eddy simulation
MCMC	Markov chain Monte Carlo
ML	Marginal likelihood
MLR	Marginal likelihood ratio
MRV	Magnetic resonance velocimetry
OF	Occam factor
PID	Proportional, integral, derivative control
RANS	Reynolds averaged Navier–Stokes

Thesis structure

This thesis is divided into three parts. Part I contains two introductory chapters, which provide context and motivation for the thesis. In chapter 1, we provide historical context of the problem of thermoacoustic oscillations, with a focus on aerospace propulsion. We then discuss the difficulty of producing quantitatively accurate models of thermoacoustic oscillations, and the need for such models. In chapter 2 we introduce the adjoint-accelerated Bayesian inference framework, which we use in this thesis to improve the accuracy of thermoacoustic models. The framework provides three main tools: (i) parameter inference, (ii) model selection, and (iii) optimal experiment design. We develop the key equations for each of the tools, and provide an intuitive explanation of each one.

Part II contains three results chapters, which are based on a set of papers that were written during the course of this research. In each of the chapters, we apply adjoint-accelerated Bayesian inference to a different canonical problem in thermoacoustics. In chapter 3 we study a hot wire Rijke tube, on which we demonstrate all three of the tools in our framework. In chapter 4 we study a ducted laminar conical flame, on which we demonstrate parameter inference and model selection. In chapter 5 we study a ducted turbulent conical flame, on which we demonstrate parameter inference.

Part III contains one conclusion chapter. We begin with a summary of the main conclusions drawn in each of the results chapters (detailed conclusions can be found in each of the results chapters). We then provide an outlook on the current state of this research, and its potential in the future.

PART I:

INTRODUCTION

Chapter 1

Thermoacoustics

This chapter provides a historical context of the problem of thermoacoustic oscillations in aerospace propulsion, and discusses the challenge of generating quantitatively accurate thermoacoustic models. Further literature review on the specific systems studied in this thesis is provided in Chapters 3 and 4.

1.1 Historical background

Thermoacoustic oscillations were first reported in the 1800s by Higgins [1], Sondhauss [2] and Rijke [3]. Nearly eighty years later, the first physically correct description of the underlying mechanism was provided by Rayleigh [4]:

“If heat be periodically communicated to, and abstracted from, a mass of air vibrating (for example) in a cylinder bounded by a piston, the effect produced will depend upon the phase of the vibration at which the transfer of heat takes place. If heat be given to the air at the moment of greatest condensation, or taken from it at the moment of greatest rarefaction, the vibration is encouraged. On the other hand, if heat be given at the moment of greatest rarefaction, or abstracted at the moment of greatest condensation, the vibration is discouraged.”

Rayleigh’s insight reveals that if heat release rate fluctuations are sufficiently in phase with pressure oscillations, the flame drives the acoustic oscillations. If the driving exceeds the dissipation of acoustic energy through damping or acoustic radiation, the amplitude of the acoustic oscillations grow, and the system is thermoacoustically unstable.

In the 1800s, thermoacoustic oscillations were only a curious phenomenon experienced by glass blowers and demonstrated in laboratories. With the advent of rocket engines in the 1930s, however, they became a significant engineering challenge [5]. Rocket engines are essentially hard-walled pressure vessels with little acoustic damping, and the mean heat release rates are exceptionally high. If even a small percentage of the mean heat release rate is converted to work, the system becomes unstable (about 0.1%, according to Huang and Yang [6]).

Historically, gas turbine engines were relatively impervious to thermoacoustic instability, predominantly because the highly perforated liners of early designs provided ample acoustic damping. In the drive for higher efficiencies and lower emissions, however, designers have been pushed towards making gas turbine combustion chambers look more like their rocket counterparts. As a result, thermoacoustic instability has become increasingly problematic for gas turbine manufacturers, often preventing them from exploiting higher performance designs and operating regimes [7].

Despite over a century of research, thermoacoustic instability remains a significant challenge today [8]. Designers are typically unable to predict thermoacoustic instability early in the design cycle. Instead, it is often only discovered during the first full-scale prototype testing [9, 10]. This requires design changes late in the development cycle, increasing cost and causing delays. The most famous example of this is the F1 rocket engine, which was developed by Rocketdyne for the Apollo program. The development program was plagued by combustion instabilities, which required over 2000 full scale tests to remove. This process took 4 years and cost around 2 billion US dollars in the 1960s [11].

In the hope of achieving stability at lower cost and with less delay, active feedback control was proposed [12, 13]. This involves, in its simplest form, introducing intentional perturbations to the system with an actuator. A controller selects the phase of these perturbations such that they cancel the thermoacoustic oscillations, and stabilize the system. This has been demonstrated on full scale engines through modulation of the fuel flow [14, 15]. It has not been widely adopted, however, because the actuators present a single point of failure with potentially catastrophic consequences. More recently, active control of thermoacoustics has been revisited, but with the control achieved through plasma discharge near the base of the flame [16–18]. It remains to be seen whether this will face the same trepidation as its predecessor.

The reluctance to adopt active feedback control motivates a continual search for passive control mechanisms. Acoustic dampers, such as liners and Helmholtz resonators, are widely used in practice as a mitigation strategy for thermoacoustic instability. If an engine is found to

be thermoacoustically unstable during full scale testing, it can be retro-fitted with a Helmholtz resonator, or a series of Helmholtz resonators, tuned to damp the unstable mode(s). This is, however, dependent on the resonator(s) fitting within the packaging constraints and mass budget of the engine.

The goal of engine manufacturers, therefore, is to be able to design systems that are linearly stable throughout the operating envelope, without having to retro-fit dampers or avoid certain operating conditions. We now explore why this is so difficult to achieve.

1.2 Modelling thermoacoustic oscillations

Thermoacoustic instabilities are extremely sensitive to small changes in the system's design or operating conditions [19]. From the designer's perspective, this extreme sensitivity can be beneficial, because only small design changes are required to stabilize the system [9, 11]. The challenge, however, lies in devising the appropriate design change to make. This can be done elegantly by using adjoint shape optimization [20], if a quantitatively accurate model is available.

The modeller, however, is typically unable to achieve quantitative accuracy. This is largely a consequence of the extreme sensitivity mentioned above, which makes model predictions sensitive to the model parameters. The values of the model parameters are not known a-priori, and often depend on detailed dynamics. For example, it is clear from Rayleigh's description that we need to know the phase difference between velocity perturbations and heat release rate fluctuations. In a typical flame, this phase difference depends on phenomena such as spray dynamics, droplet evaporation, detailed flame kinematics and flame-flow interactions. These detailed dynamics occur over a wide range of spatio-temporal scales and involve many physical processes, presenting a significant modelling challenge.

Recent studies have used large eddy simulations (LES) to resolve these effects in laboratory burners [21–26], and even industrially relevant geometries [8, 27–29]. This research has provided valuable insight into the detailed interactions described above. The computational cost of LES, however, continues to prevent its widespread adoption by industry [8, 30]. While it is currently used for spot-checks at the detailed design phase, it cannot feasibly be used as the design tool that industry needs.

The acoustic network model is a versatile reduced-order modelling framework, which is suitable for use as a design exploration tool. In this framework, the system is modelled as a network of 1D or 2D axisymmetric acoustic elements [31, 32]. The acoustic waves are modelled

as forward and backward travelling waves of constant strength within each acoustic element. Adjacent elements are connected through jump conditions, enforcing the conservation of momentum and energy¹. The complex amplitudes of the waves in each element are determined by solution of the linear system comprising the jump and boundary conditions. This framework has been widely applied in both industry and academia, but requires a reduced-order model for the detailed flame dynamics described above.

The approach proposed by Crocco and Cheng [33], variations of which are still largely in use today, is to lump together all the detailed dynamics, and model the flame response as a “black-box” input-output system. In Crocco’s framework, the relationship between velocity perturbations and heat release rate fluctuations is modelled as a time lag: $Q' = nu'(t - \tau)$, where Q' is the fluctuating heat release rate, u' is the acoustic velocity, n is the interaction index and τ is the time delay. This has since been extended to account for dependence on forcing frequency (the flame transfer function) [34], dependence on forcing amplitude (the flame describing function) [35–37], and dependence on pressure fluctuations (the flame transfer matrix) [38, 39]. More recently, Polifke [40] proposed a distributed time delay approach, where several discrete time delays replace the single time delay in Crocco’s model. Polifke describes this as a “grey-box” model, because the discrete time delays can often be linked to physical length and velocity scales in the system.

While all these approaches simplify the analysis, they introduce new parameters, to which model predictions are sensitive. These parameters are typically obtained experimentally [37, 41–43] or using computational fluid dynamics (CFD) [41, 44–46]. The response of a flame to acoustic forcing has, however, been shown to change sensitively (i) with changes in the operating condition [47, 48], (ii) with changes to the confinement of the flame [49, 50], and (iii) when the flame is combined with other flames [37, 43]. Therefore, unless experiments or simulations are representative of the full system, the parameters will be inaccurate, and the overall thermoacoustic model will be qualitatively accurate at best.

The modeller can, however, exploit the extreme sensitivity of thermoacoustics if a data-driven modelling approach is taken. This is because the sensitivity makes the uncertain model parameters easy to observe from the experimental data. While there has been a lot of interest in data-driven methods in the thermoacoustics community, this has mostly focussed on early prediction [51–55], or applying physics-agnostic frameworks with millions of parameters [56–59]. In this thesis we adopt a probabilistic data-driven framework that allows us to assimilate

¹It is only necessary to enforce the conservation of mass when the Mach number of the mean flow is significant. This is beyond the scope of this thesis, where we study systems with negligible mean flow.

experimental data into physics-based models with few parameters, making them quantitatively accurate.

Chapter 2

Bayesian Inference

Bayesian probability theory has been applied to scientific problems since the late 1700s. Laplace wrote extensively on its use in astronomy, and developed an elegant set of tools for scientific inference [60, 61]. These tools, however, were largely abandoned in favour of the frequentist framework shortly after Laplace's death [62]. Beginning in the early 20th century, Laplace's approach was revived, updated and expanded upon by several authors [63–67]. The latest of these, the work of MacKay [67], is the approach to Bayesian inference that we use in this thesis. In its broader forms, Bayesian inference is still actively applied in astronomy and astrophysics [68–71], but has also been adopted in biology [72–74], economics [75–77], geophysics and meteorology [78–81], and engineering, where it has predominantly been applied in structural mechanics [82–84].

By comparison, engineers working in fluid dynamics have made very little use of the Bayesian framework. In their review of machine learning in fluid dynamics, Brunton et al. [85] argue that, for fluid mechanics problems, Bayesian inference may be superior to other machine learning techniques because of its robustness, but that it is hampered by the cost of the thousands of model evaluations required to compute the posterior distribution. While this is true of typical sampling methods, such as Markov Chain Monte Carlo, it is not true in the approximate inference framework of Laplace. This framework reduces the required model evaluations, making Bayesian inference feasible for computationally expensive models [81, 86].

In this chapter we provide an overview of the Bayesian inference framework that we use in this thesis. In § 2.1 we present three tools that allow us to use noisy, sparse experimental observations to improve the accuracy of physics-based models. With these tools, we can (i) infer

the most probable parameters of a model, (ii) quantify the uncertainty in the parameters and model predictions, (iii) estimate the systematic and structural uncertainty in the experiments and model, and (iv) compare several candidate models and select the best one. In § 2.2 we extend the framework to Bayesian optimal experiment design, which allows us to minimize the experimental effort required to apply Bayesian inference.

2.1 Adjoint-accelerated Bayesian inference

We consider the case where we have a set of candidate models, \mathcal{H}_j , each with a set of unknown parameters, \mathbf{a} . Each model predicts the system state, \mathbf{s} , which we test against noisy, sparse experimental observations, \mathbf{z} . We apply Bayesian inference at two levels. At the first level, we infer the most probable parameters of each model, \mathbf{a}_{MP} , and quantify the uncertainty in these parameters. At the second level, we compare the candidate models to select the best model, given the data.

2.1.1 Parameter inference

At the first level of inference we assume that the candidate model, \mathcal{H}_j , is structurally correct, and we attribute all model inaccuracy to errors in the unknown parameters. This assumption will rarely be correct, so we will revisit it later. We encode our level of uncertainty in the parameter values through a probability distribution, which we denote $p(\bullet)$. Using any prior knowledge we have about the unknown parameters (which may be none at all), we propose a prior probability distribution over the parameter values, $p(\mathbf{a}|\mathcal{H}_j)$. We then assimilate the data, \mathbf{z} , by performing a Bayesian update on the parameter values:

$$p(\mathbf{a}|\mathbf{z}, \mathcal{H}_j) = \frac{p(\mathbf{z}|\mathbf{a}, \mathcal{H}_j)p(\mathbf{a}|\mathcal{H}_j)}{p(\mathbf{z}|\mathcal{H}_j)} \quad (2.1)$$

The quantity on the left-hand side of Eq. (2.1) is the posterior probability of the parameters, given the data. It is generally computationally intractable to calculate the full posterior, because it requires integration over parameter space. The integral typically cannot be evaluated analytically, and requires thousands of model evaluations to compute numerically. At the parameter inference stage, however, we are only interested in finding the most probable parameters, which are those that maximize the posterior. We therefore use an optimization algorithm to find the peak of the posterior without evaluating the full distribution. This process is made computationally efficient by (i) assuming that the experimental uncertainty is Gaussian

distributed, and (ii) choosing the prior parameter distribution to be Gaussian. Assumption (i) is reasonable for well-designed experiments in which the uncertainty is dominated by random error, which is typically Gaussian distributed. For assumption (ii) we note that the choice of prior is often the prerogative of the researcher, and we are free to exploit the mathematical convenience offered by the Gaussian distribution.

When finding the most probable parameters, we neglect the denominator of the right-hand side of Eq. (2.1), because it does not depend on the parameters. It is then convenient to define a cost function, \mathcal{J} , as the negative log of the numerator of Eq. (2.1), which we minimize:

$$\begin{aligned} \mathcal{J} = & \frac{1}{2}(\mathbf{s}(\mathbf{a}) - \mathbf{z})^T \mathbf{C}_{ee}^{-1}(\mathbf{s}(\mathbf{a}) - \mathbf{z}) \\ & + \frac{1}{2}(\mathbf{a} - \mathbf{a}_p)^T \mathbf{C}_{aa}^{-1}(\mathbf{a} - \mathbf{a}_p) + K \end{aligned} \quad (2.2)$$

where \mathbf{s} and \mathbf{z} are column vectors of the model predictions and experimental observations respectively, \mathbf{C}_{ee} is the covariance matrix describing the experimental uncertainty, \mathbf{a} and \mathbf{a}_p are column vectors of the current and prior parameter values respectively, \mathbf{C}_{aa} is the covariance matrix describing the uncertainty in the prior, and K is a constant from the Gaussian pre-exponential factors, which has no impact on the most probable parameters.

To find the minimum of \mathcal{J} with the fewest model evaluations, we use gradient-based optimization. We see from Eq. (2.2) that the cost function gradient, $\partial \mathcal{J} / \partial \mathbf{a}$, depends on the model's parameter sensitivities, $\partial \mathbf{s} / \partial \mathbf{a}$. We obtain these cheaply using first order adjoint methods, which allow us to evaluate the gradient of a function with respect to many parameters, with a computational cost that is independent of the number of parameters [87, 88]. In fluid mechanics, the adjoint method been widely applied to shape optimization [20, 89, 90], so adjoint codes are readily available. We note that automatic differentiation presents similar opportunities, but care must be taken that the code is structured such that it is differentiable. Careful application of the adjoint approach always produces a well-posed differentiable code.

The parameter inference process is illustrated in Fig. 2.1 for a simple system with a single unknown parameter, a , and a single observable variable, z . In (a) we show the marginal probability distributions of the prior, $p(a)$ and the data, $p(z)$. The prior and data are independent, so we construct the joint distribution, $p(a, z)$ by multiplying the two marginals. In (b), we overlay the model predictions, s , for various values of a . Marginalizing along the model predictions yields the true posterior, $p(a|z)$. This is possible for a cheap model with a single

parameter, but exact marginalization quickly becomes intractable as the number of parameters increases. In (c) we plot the cost function, \mathcal{J} , which is the negative log of the unnormalized posterior. We show the three steps of gradient-based optimization that were required to find the local minimum, which corresponds to the most probable parameters, a_{MP} .

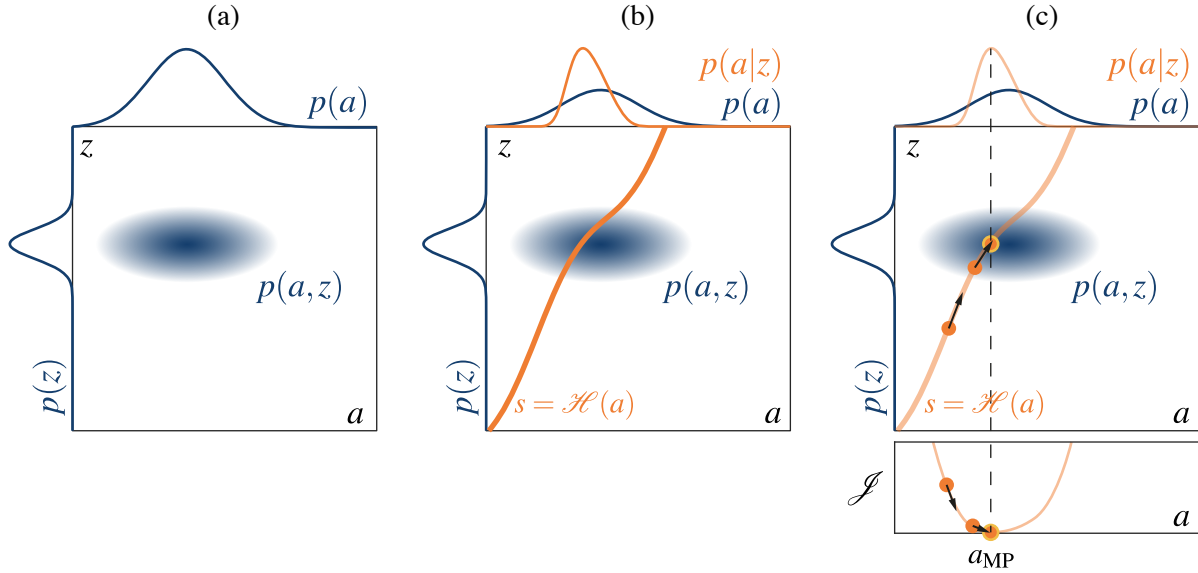


Fig. 2.1 Illustration of parameter inference on a simple univariate system. (a) the marginal probability distributions of the prior and data, $p(a)$ and $p(z)$, as well as their joint distribution, $p(a, z)$ are plotted on axes of parameter value, a , vs observation outcome, z . (b) the model, \mathcal{H} , imposes a functional relationship between the parameters, a , and the predictions, s . Marginalizing along the model predictions yields the true posterior, $p(a|z)$. This cannot be done for computationally expensive models with even moderately large parameter spaces. (c) instead of evaluating the full posterior, we use gradient-based optimization to find its peak. This yields the most probable parameters, a_{MP} .

2.1.2 Uncertainty quantification

Uncertainty quantification can be split into three steps: (i) quantifying the uncertainty in the parameters, (ii) propagating the parametric uncertainty to the model predictions, and (iii) estimating the systematic and structural uncertainty in the experiments and model predictions. We will deal with these separately.

Parametric uncertainty

Once we have found the most probable parameter values by minimizing \mathcal{J} in Eq. (2.2), we estimate the uncertainty in these parameter values using Laplace's method [63, 67]. This

method approximates the posterior probability distribution as a Gaussian, centred around the most probable parameters, \mathbf{a}_{MP} . Laplace showed that the inverse covariance of the Gaussian that best approximates the true posterior is given by the Hessian of \mathcal{J} :

$$\begin{aligned} \mathbf{C}_{aa}^{\text{MP}^{-1}} &\approx \frac{\partial^2 \mathcal{J}}{\partial a_l \partial a_m} \\ &= \mathbf{C}_{aa}^{-1} + \mathbf{J}^T \mathbf{C}_{ee}^{-1} \mathbf{J} + (\mathbf{s}(\mathbf{a}_{\text{MP}}) - \mathbf{z})^T \mathbf{C}_{ee}^{-1} \mathbf{H} \end{aligned} \quad (2.3)$$

where \mathbf{J} is the Jacobian matrix containing the parameter sensitivities of the model predictions, $\partial s_l / \partial a_m$, and \mathbf{H} is the rank three tensor containing the second order sensitivities, $\partial^2 s_l / \partial a_m \partial a_n$. Both \mathbf{J} and \mathbf{H} are evaluated at \mathbf{a}_{MP} .

Equation (2.3) provides the first evidence of a claim that was made at the end of Chapter 1: the extreme sensitivity of thermoacoustics can be beneficial to the modeller, because it makes the parameters easy to observe from the data. The inverse covariance matrix quantifies the precision to which the parameters are known. It is then clear from Eq. (2.3) that we can infer the parameters more precisely when the model is sensitive to the parameter values.

The accuracy of Laplace's method depends on the functional dependence of the model on the parameters. This is shown graphically in Fig. 2.2, where we compare the uncertainty quantification process for three univariate systems. In (a), the model is linear in the parameters. Marginalizing a Gaussian joint distribution along any intersecting line produces a Gaussian posterior distribution, so Laplace's method is exact. In (b), the model is weakly nonlinear in the parameters. The true posterior is skewed, but the Gaussian approximation is still reasonable. This panel also shows a geometric interpretation of Laplace's method: the approximate posterior is given by linearizing the model around \mathbf{a}_{MP} , and marginalizing the joint distribution along the linearized model. In (c), the model is strongly nonlinear in the parameters, so the true posterior is multi-modal and the main peak is highly skewed. Laplace's method underestimates the uncertainty in this case. Furthermore, the cost function has two local minima, but the parameter inference step will only find one peak, which will depend on the choice of initial condition for the optimization.

This simple example seems to imply that Laplace's method is only suitable for weakly nonlinear models. It has, however, only considered the case where a single data point is assimilated. In many cases, the true posterior tends to a Gaussian distribution as the number of observations increases, even for models that are strongly nonlinear in the parameters [91, § 10.2]. For a given model, the accuracy of Laplace's method can be checked *a-posteriori*

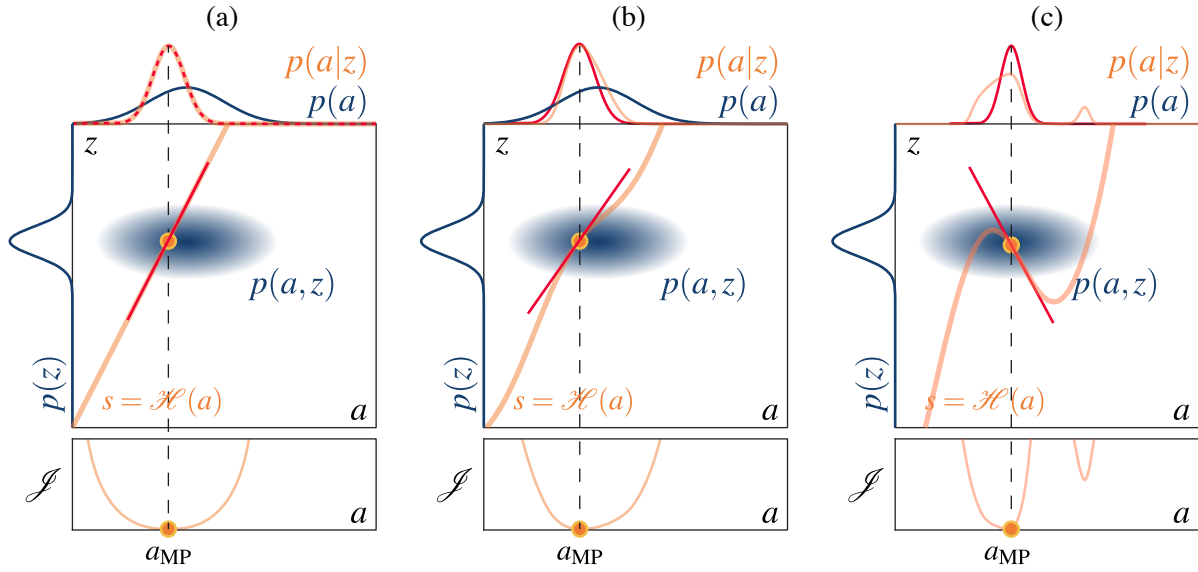


Fig. 2.2 Illustration of uncertainty quantification for three univariate systems. (a) the model is linear in the parameters, so the true posterior is Gaussian and Laplace’s method is exact. (b) the model is weakly nonlinear in the parameters, the true posterior is slightly skewed, but Laplace’s method yields a reasonable approximation. (c) the model is strongly nonlinear in the parameters, the posterior is multi-modal and Laplace’s method underestimates the uncertainty.

using a sampling method such as Markov Chain Monte Carlo (MCMC). Previous work has applied MCMC to thermoacoustic network models [92] and more complex models in fluid mechanics [93], both of which showed the posteriors to be approximately Gaussian. If the true posterior is found to be poorly approximated by a Gaussian, the researcher can attempt to reduce the extent of the nonlinearity captured by the joint distribution by (i) shrinking the joint distribution by providing more precise prior information or more precise experimental data, or (ii) re-parameterizing the model to reduce the strength of the nonlinearity [67, Chapter 27].

Uncertainty propagation

To quantify the parametric uncertainty in the model predictions, we propagate the parameter uncertainties through the model. This is done cheaply by linearizing the model around \mathbf{a}_{MP} and propagating the uncertainties through the linear model. The uncertainty in the model predictions is given by:

$$\mathbf{C}_{ss} = \mathbf{J}^T \mathbf{C}_{aa} \mathbf{J} \quad (2.4)$$

where \mathbf{C}_{ss} is the covariance matrix describing the variance in the model predictions due to the variance in the parameters. The marginal variance in each predicted variable, s_l , is given by the diagonal elements of \mathbf{C}_{ss} , because the prediction uncertainty is Gaussian distributed [94, Theorem 1.2.6 & 1.2.7].

This allows us to quantify the uncertainty in the model predictions due to the uncertainty in the parameters, but we have still been working under the assumption that the model is structurally correct. We now relax this assumption, and introduce a method for estimating the systematic and structural uncertainty in the experiments and model predictions.

Systematic uncertainty

In most cases, experimental data will contain some systematic uncertainty, and models will contain some structural uncertainty. These uncertainty sources cannot be quantified *a-priori*, and are often referred to as “unknown unknowns”. We can, however, construct a total covariance matrix, \mathbf{C}_{tt} , which encodes the total uncertainty due to (i) the known experimental uncertainty, (ii) the unknown systematic experimental uncertainty, and (iii) the unknown structural model uncertainty. We can then estimate this total covariance from the posterior discrepancy between the model and the data. This must be done simultaneously with parameter inference, because the posterior parameter distribution depends on the total uncertainty in the model and data. We therefore replace \mathbf{C}_{ee} with \mathbf{C}_{tt} in Eq. (2.2), and estimate the total uncertainty by simultaneously minimizing \mathcal{J} with respect to \mathbf{a} and \mathbf{C}_{tt}^{-1} .

We begin by calculating the derivative of \mathcal{J} with respect to \mathbf{C}_{tt}^{-1} , assuming that the observed variables are uncorrelated, and keeping in mind that the normalizing constant, K , contains the term $\log(|\mathbf{C}_{tt}|^{1/2})$. This gives:

$$\frac{\partial \mathcal{J}}{\partial \mathbf{C}_{tt}^{-1}} = \frac{1}{2}(\mathbf{s}(\mathbf{a}) - \mathbf{z})(\mathbf{s}(\mathbf{a}) - \mathbf{z})^T \circ \mathbf{I} - \frac{1}{2}\mathbf{C}_{tt} \quad (2.5)$$

where \mathbf{I} is the identity matrix, and \circ denotes the Hadamard product. For a given set of parameters, the most probable \mathbf{C}_{tt} sets Eq. (2.5) to zero. This gives the estimate:

$$\mathbf{C}_{tt} = (\mathbf{s}(\mathbf{a}) - \mathbf{z})(\mathbf{s}(\mathbf{a}) - \mathbf{z})^T \circ \mathbf{I} \quad (2.6)$$

which is the expected result that the total variance in the model and data is the square of the discrepancy between the model predictions and the data. Although we cannot directly identify the source of the unknown uncertainty because the experimental and model uncertainties cannot

be disentangled, the inferred total uncertainty can assist the researcher with identifying potential error sources. For example, if the unknown error in a single sensor is unexpectedly large, this could indicate a faulty sensor or bad installation. If the unknown error at a certain experimental operating condition is large, this could prompt the researcher to repeat that experiment. If the unknown error grows with one of the input variables, the researcher might investigate the model to see if any important physical phenomena may have been neglected.

In Fig. 2.3, we demonstrate this with a toy problem in which we infer the parameters and total uncertainty of two polynomial models. We generate data by sampling a ground truth polynomial, which is of degree seven, and corrupting the data with random error. In (a) the model is structurally correct, so when we infer the parameters and total uncertainty we find that the model can fit the data, and the inferred total uncertainty is close to the known error. In (b) the model is a fifth degree polynomial, so it is structurally incorrect. We find that the total uncertainty is much larger than the known random error, indicating systematic or structural error.

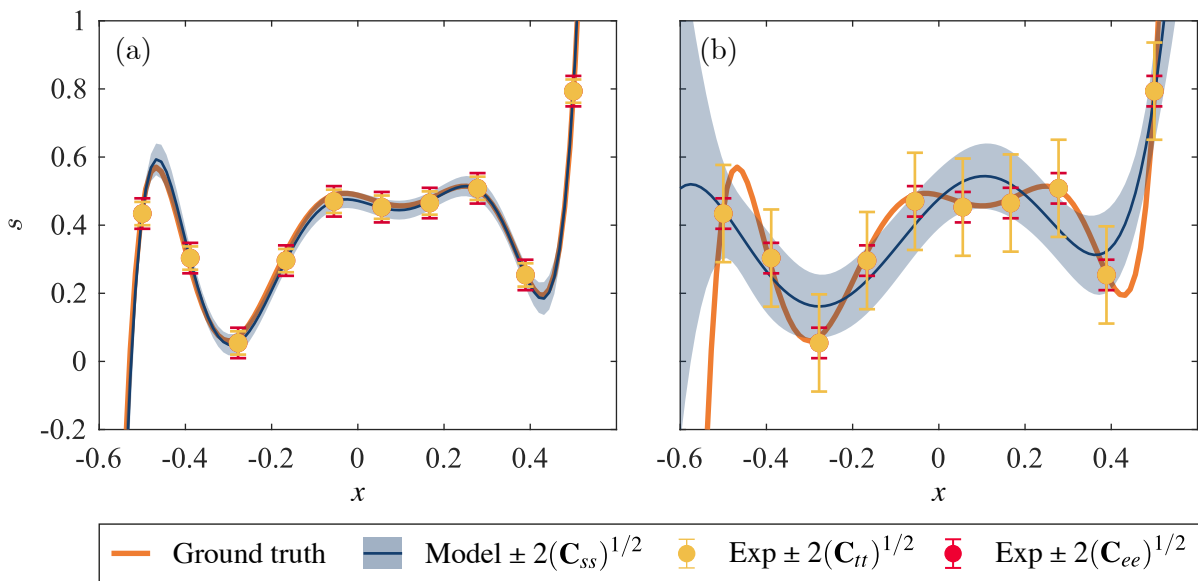


Fig. 2.3 Demonstration of systematic uncertainty estimation on a toy problem of inferring the coefficients of a polynomial. In each case we collect data by sampling the ground truth polynomial (orange) at 10 points, and corrupting the data with Gaussian random noise (red dots and error bars). The model predictions (blue lines) are shown with two standard deviations of parametric uncertainty (blue patch). The inferred total uncertainty is plotted as a second set of data points (yellow dots and error bars). In (a) we infer the parameters and systematic error of a polynomial of the correct degree, and we find that the model can fit the data, and the inferred total uncertainty is close to the known error. In (b) we infer the parameters and systematic error of a polynomial of the wrong degree, and we find that the total uncertainty is much larger than the known random error, indicating systematic or structural error.

2.1.3 Model selection

At the second level of inference, we calculate the posterior probability of each model, given the data. This allows us to compare several candidate models quantitatively. We use Bayes' theorem applied to the models, \mathcal{H}_j , and data, \mathbf{z} :

$$p(\mathcal{H}_j|\mathbf{z}) \propto p(\mathbf{z}|\mathcal{H}_j)p(\mathcal{H}_j) \quad (2.7)$$

The first factor on the right-hand side of Eq. (2.7) is the denominator of Eq. (2.1), which is referred to as the marginal likelihood, or evidence. The second factor is the prior probability that we assign to each model. If we have no reason to prefer one model over another, we assign equal probabilities to all models and rank them according to their evidence. The evidence is calculated by integrating the numerator of Eq. (2.1) over parameter space. Recalling that the posterior distribution has been approximated as a Gaussian, we can cheaply calculate the evidence using Laplace's approximation for the integral of a multivariate Gaussian:

$$p(\mathbf{z}|\mathcal{H}_j) \approx p(\mathbf{z}|\mathbf{a}_{\text{MP}}, \mathcal{H}_j) \times p(\mathbf{a}_{\text{MP}}|\mathcal{H}_j) |C_{aa}^{\text{MP}}|^{1/2} \quad (2.8)$$

The first factor on the right-hand side of Eq. (2.8), called the best fit likelihood, is a measure of how well the model fits the data. The second factor, called the Occam factor, penalizes the model based on its parametric complexity, where the complexity is measured by how precisely the parameter values must be tuned for the model to fit the data to within the experimental uncertainty. The model with the largest evidence is the simplest model that is capable of describing the data, for given measurement error and given priors. This process therefore naturally enforces Occam's razor to select the best model.

We once again demonstrate this with a toy problem in which we infer the parameters and total uncertainty of several polynomial models. The ground truth polynomial has degree seven, and we generate synthetic data by sampling the ground truth and corrupting the samples with random noise. We then assimilate the data into models of degree two to degree twelve, and calculate the model comparison metrics for each model. Fig. 2.4 shows the posterior predictions of three of the candidate models: (a) degree five, (b) degree seven, and (c) degree twelve.

The model comparison metrics are plotted for all eleven models in Fig. 2.5. We see that models of degree lower than ground truth (1-6) are penalized by a small best fit likelihood (orange), but their simplicity is rewarded by a small negative Occam factor. Their simplicity does not, however, make up for the poor fit, and their log-evidence (marginal likelihood) is

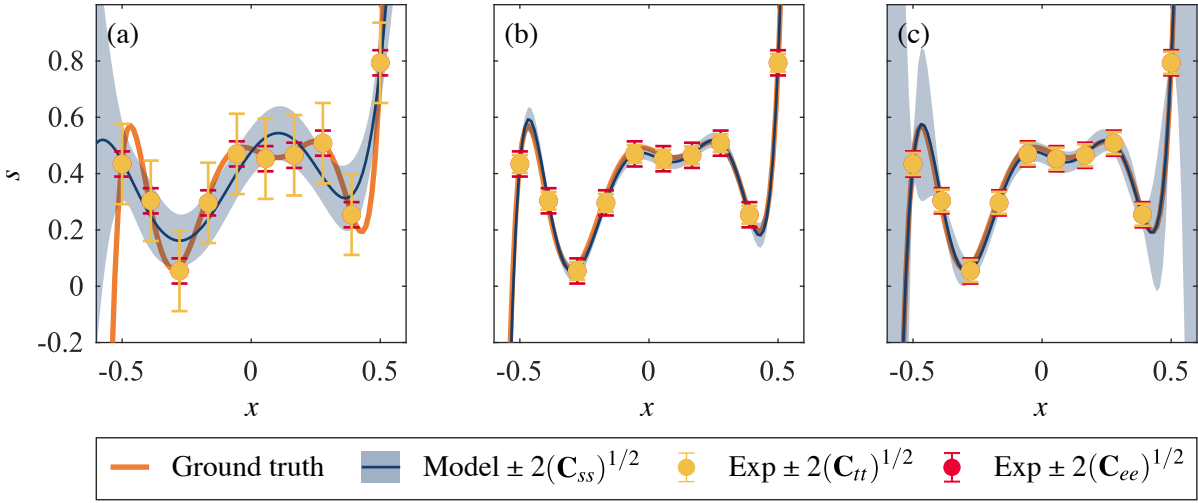


Fig. 2.4 As for Fig. 2.3, but showing three candidate models of degree (a) five, (b) seven, and (c) twelve. Ground truth is a polynomial of degree seven.

small. Models of degree higher than ground truth fit the data well, but their complexity is penalized by a large negative Occam factor. We see that the model that is structurally correct, degree seven, has the largest evidence and would be selected as the best model. Of course, the outcome of this toy problem depends on the random noise that is added. If a more flexible model happens to fit the noise well, it can become more likely than the structurally correct model. We repeat this process with 100 realizations of random noise, and find that the structurally correct model is chosen 88% of the time. Of the remaining cases, seven are marginal (less than 5% difference in evidence), while in five cases the more flexible model has significantly higher evidence.

2.2 Bayesian optimal experiment design

The framework described in § 2.1 allows the modeller to benefit from the work of the experimentalist by using experimental observations to improve the accuracy of physics-based models. We now extend this framework to enable the modeller to inform the experimental design process in order to minimize the experimental cost. For this section, we add to our notation the vector of experimental design variables, \mathbf{x} , which are the variables that the experimentalist can control.

There are a number of ways in which the modeller could inform the experiment design. In this thesis, we focus on three specific questions an experimentalist may face: (i) which experimental design would provide the maximum information about the unknown parameters, (ii) where should the sensors be placed to provide the maximum information about the unknown




Fig. 2.5 Model ranking metrics for eleven candidate models. The best fit likelihood, $\log(\text{BFL})$ in orange, is a measure of how well the model fits the data. The Occam factor, $\log(\text{OF})$ in yellow, penalizes the model based on its parametric complexity. The evidence or marginal likelihood, $\log(\text{ML})$ in blue, is the sum of the best fit likelihood and the Occam factor. The model with the largest evidence is the simplest model that is capable of describing the data, for given measurement error and given priors.

parameters, and (iii) which experimental design would maximize the discrimination between candidate models?

Each of these questions can be answered by using metrics from information theory to quantify the information content of a candidate experiment. We follow the general approach proposed by Lindley [95], which has been applied to many other optimal experiment design studies [96–99]. For each of the experimental questions listed above, we define a suitable utility function, $u(\mathbf{x}, \mathbf{z}, \mathbf{a})$. We then calculate the *expected* utility by integrating over all possible realizations of the parameters and data:

$$U(\mathbf{x}) = \int_{\mathbf{Z}} \int_{\mathbf{A}} u(\mathbf{x}, \mathbf{z}, \mathbf{a}) p(\mathbf{a}|\mathbf{x}, \mathbf{z}) p(\mathbf{z}|\mathbf{x}) d\mathbf{a} d\mathbf{z} \quad (2.9)$$

where $U(\mathbf{x})$ is the expected utility, $u(\mathbf{x}, \mathbf{z}, \mathbf{a})$ is the utility function, and \mathbf{Z} and \mathbf{A} are the support of $p(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{a}|\mathbf{x}, \mathbf{z})$ respectively. The optimal experiment to perform is the one that maximizes the expected utility.

Previous studies have faced difficulties with the cost of computing the expected utility within their inference frameworks, which typically requires Monte Carlo integration over high dimensional spaces [96–99]. We will show that in the adjoint-accelerated Bayesian inference framework, the expected utility can be computed cheaply.

2.2.1 Optimal design for parameter inference

We consider the situation where we have assimilated data from i experiments, where i could be zero. We want to know which experiment to perform next in order to gain maximal information about the unknown parameters. The information content can be quantified through the Shannon entropy of the parameter probability distribution:

$$S_i = - \int_{\mathbf{A}} p(\mathbf{a}|\mathbf{z}_i) \log_2(p(\mathbf{a}|\mathbf{z}_i)) \, d\mathbf{a} \quad (2.10)$$

where S_i is the Shannon entropy of the parameter probability distribution after the i^{th} experiment has been assimilated¹. The information we gain by assimilating the data point \mathbf{z}_{i+1} is given by the change in Shannon entropy: $\Delta S_{i+1} = S_i - S_{i+1}$ [95], which we choose as our utility function for this case:

$$\begin{aligned} u(\mathbf{x}, \mathbf{z}, \mathbf{a}) &= \Delta S_{i+1} = S_i - S_{i+1} \\ &= \int_{\mathbf{A}} p(\mathbf{a}|\mathbf{z}_i) \log_2(p(\mathbf{a}|\mathbf{z}_i)) \, d\mathbf{a} - \int_{\mathbf{A}} p(\mathbf{a}|\mathbf{z}_{i+1}) \log_2(p(\mathbf{a}|\mathbf{z}_{i+1})) \, d\mathbf{a} \end{aligned} \quad (2.11)$$

We see that the utility function involves integration over \mathbf{a} , so $u(\mathbf{a}, \mathbf{x}, \mathbf{z}) = u(\mathbf{x}, \mathbf{z})$. The expected utility is therefore:

$$\begin{aligned} U &= \int_{\mathbf{Z}} \int_{\mathbf{A}} u(\mathbf{x}, \mathbf{z}) p(\mathbf{a}|\mathbf{x}, \mathbf{z}) p(\mathbf{z}|\mathbf{x}) \, d\mathbf{a} d\mathbf{z} \\ &= \int_{\mathbf{Z}} u(\mathbf{x}, \mathbf{z}) p(\mathbf{z}|\mathbf{x}) \, d\mathbf{z} \end{aligned} \quad (2.12)$$

To evaluate the remaining integral, we once again exploit Laplace's method, recalling that the parameter probability distribution is always Gaussian in this framework. The difference in Shannon entropy between two Gaussians is given by:

$$\Delta S_{i+1} = \frac{1}{2} \log_2 \left(\frac{|\mathbf{C}_{aa}^{-1}|_{i+1}}{|\mathbf{C}_{aa}^{-1}|_i} \right) \quad (2.13)$$

from which we see that maximising the information gained about the parameters is equivalent to finding the maximum reduction in parameter uncertainty. Recall that, from Laplace's approximation in Eq. (2.3), assimilating the measurement \mathbf{z}_{i+1} updates the parameter uncertainty

¹For the purpose of optimal experiment design, the choice of the logarithm base is arbitrary. We have chosen to work with base 2 logarithms, which provide information in units of bits.

according to:

$$(\mathbf{C}_{aa}^{-1})_{i+1} \approx (\mathbf{C}_{aa}^{-1})_i + \mathbf{J}_{i+1}^T \mathbf{C}_{ee}^{-1} \mathbf{J}_{i+1} + (\mathbf{s}(\mathbf{a}_i) - \mathbf{z}_{i+1})^T \mathbf{C}_{ee}^{-1} \mathbf{H}_{i+1} \quad (2.14)$$

where the Jacobian, \mathbf{J}_{i+1} , and Hessian, \mathbf{H}_{i+1} , are evaluated with the previous most probable parameters, \mathbf{a}_i , and with the candidate experiment design parameters, \mathbf{x}_{i+1} .

From Eqs. (2.13) and (2.14) we can see that the choice of which experiment to perform next depends on the outcome of that experiment, \mathbf{z}_{i+1} , as one might expect. However, the data-dependence only occurs in the second order term in Eq. (2.14), which is exactly zero for models that are linear in the parameters, and is often small compared to the first order term for nonlinear models². We therefore proceed with a first order approximation, in which the posterior covariance is independent of the experimental outcome:

$$(\mathbf{C}_{aa}^{-1})_{i+1} \approx (\mathbf{C}_{aa}^{-1})_i + \mathbf{J}_{i+1}^T \mathbf{C}_{ee}^{-1} \mathbf{J}_{i+1} \quad (2.15)$$

If the second order sensitivities are found to be non-negligible for a model of interest, we may still be able to neglect the data-dependent term on the grounds that the discrepancy, $(\mathbf{s}(\mathbf{a}_i) - \mathbf{z}_{i+1})$, will become small as the model is updated with more data. In this case we expect that the initial experiments may not be optimal, but that the chosen experiments will become optimal as more experiments are assimilated and the discrepancy decreases. In many cases this may be an acceptable sacrifice for the computational cost savings of the proposed approach.

By neglecting the data-dependent term, we are able to plan the subsequent experiment using only the model and its adjoint:

$$U = \Delta S_{i+1} = \frac{1}{2} \log_2 \left(\frac{\left| (\mathbf{C}_{aa}^{-1})_i + \mathbf{J}_{i+1}^T \mathbf{C}_{ee}^{-1} \mathbf{J}_{i+1} \right|}{\left| (\mathbf{C}_{aa}^{-1})_i \right|} \right) \quad (2.16)$$

We follow a greedy sequential experiment design process, where we iteratively (i) use Eq. (2.16) to identify the most informative experiment, (ii) perform that experiment to collect the data, \mathbf{z}_{i+1} , and (iii) assimilate the data to obtain \mathbf{a}_{i+1} and \mathbf{C}_{aa}^{i+1} .

To assist with interpretation of Eq. (2.16), we consider the special case of a system with univariate design and observable variables (multiple unknown parameters are still permitted).

²We have used second order adjoint methods to confirm that this assumption is valid for the thermoacoustic models of interest in this study.

Under these conditions, the Jacobian, \mathbf{J} , reduces to a column vector, \mathbf{j} , and the experimental covariance, \mathbf{C}_{ee} , reduces to the variance, V_e . The quantity $\mathbf{J}^T \mathbf{C}_{ee}^{-1} \mathbf{J}$ in the numerator of Eq. (2.16) reduces to $V_e^{-1} \mathbf{j}^T \mathbf{j}$, which is a rank-one perturbation of the inverse covariance matrix \mathbf{C}_{aa}^{-1} . Using the identities for rank-one perturbations, Eq. (2.16) reduces to:

$$\Delta S_{i+1} = \frac{1}{2} \log_2 (1 + V_e^{-1} \mathbf{j}_{i+1}^T \mathbf{C}_{aa}^i \mathbf{j}_{i+1}) \quad (2.17)$$

This recovers the result that MacKay [100] arrived at for scalar interpolation problems: to maximize information gain we must (i) maximize V_e^{-1} , which is to say that we learn the most when we make precise observations, or (ii) maximize $\mathbf{j}^T \mathbf{C}_{aa} \mathbf{j}$, which is the posterior variance of the model predictions, as shown in Eq. (2.4). This produces the intuitive result that, for univariate systems, we learn the most about the unknown parameters when we perform experiments where the model is most uncertain.

While this aids with interpretation, MacKay's result isn't directly applicable to multivariate systems, unless an ad-hoc decision is made about which variable's uncertainty should drive the decision of which experiment to perform next. This is avoided in our framework by maximizing the change in Shannon entropy, which automatically balances the information gained from each observed variable based on how sensitive it is to the unknown parameters.

2.2.2 Optimal sensor placement for parameter inference

We now consider the case where one or more of the observed variables, \mathbf{z} , are spatially varying and are observed with point measurement sensors. We would like to know where to place the sensors in order to gain as much information as possible about the unknown parameters. We may have an existing rig with a fixed number of sensors, and we would like to know where to place the sensors to gain maximal information. Alternatively, we may be designing a new rig, for which we already have a qualitative model, and we would like to know (i) how many sensors we need to buy, and (ii) where we should make provision for instrument access.

To answer these questions, we add the sensor locations to the vector of design parameters, \mathbf{x} , and find the design parameters that maximize the change in Shannon entropy following the steps in § 2.2.1. We place the sensors sequentially, with each sensor location selected to provide maximum information based on (i) the model and (ii) the information from the existing sensor layout.

This process naturally accounts for the local reduction in information in the vicinity of existing sensors, with the correlation length determined by the model sensitivity. This removes the need to define ad-hoc methods to avoid sensor clustering, as done in previous studies on optimal sensor placement [101, 102]. We note that [102] additionally uses the variance field to identify optimal sensor placement, requiring a second ad-hoc decision about which variable should guide sensor placement, as discussed in the previous section.

2.2.3 Optimal design for model comparison

Finally, we consider the case where we are trying to identify the best model from a set of candidate models $\mathcal{H}_j, j = \{1, 2, \dots, M\}$. We assume that we have already performed the optimal experiments to learn the unknown parameters of each model. We now want to identify the optimal experiment design, \mathbf{x}_{i+1} , which maximizes the discrimination between the candidate models.

At the (as yet undetermined) experiment design \mathbf{x}_{i+1} , each candidate model will make a slightly different prediction, $\mathbf{s}_{i+1,j}$, with slightly different uncertainty, $\mathbf{C}_{ss}^{i+1,j}$. Before we make the next observation, \mathbf{z}_{i+1} , each model encodes a belief that the data will occur with a probability distribution given by a Gaussian centred around the model prediction:

$$\begin{aligned} P_j &= p(\mathbf{z}_{i+1} | \mathcal{H}_j) = \mathcal{N}(\boldsymbol{\mu}_j, \mathbf{C}_j), \\ \boldsymbol{\mu}_j &= \mathbf{s}_{i+1,j} = \mathcal{H}_j(\mathbf{x}_{i+1}, \mathbf{a}_{i,j}), \\ \mathbf{C}_j &= \mathbf{C}_{ss}^j + \mathbf{C}_{ee} \end{aligned} \quad (2.18)$$

where \mathbf{C}_{ss} is the prediction uncertainty, which is calculated using Eq. (2.4), and we have introduced the shorthand $P_j = p(\mathbf{z}_{i+1} | \mathcal{H}_j)$ to simplify the subsequent notation. In order to maximally discriminate between the models, we choose \mathbf{x}_{i+1} such that the distributions P_j are as dissimilar as possible. We therefore define our utility function to be a measure of distance between the distributions. In this case we have chosen the average divergence [103, 104]:

$$u = \frac{1}{M(M-1)} \sum_{j=1}^M \sum_{k=1}^M D_{KL}(P_j || P_k) \quad (2.19)$$

where D_{KL} is the directed Kullback-Leibler divergence, defined as:

$$D_{KL}(P_j || P_k) = \int_{\mathbf{z}} P_j \log_2 \frac{P_j}{P_k} d\mathbf{z} \quad (2.20)$$

For a multivariate Gaussian distribution with d observed variables, the directed Kullback-Leibler divergence is:

$$D_{KL}(P_j||P_k) = \frac{1}{2} \left(\log_2 \frac{|\mathbf{C}_j|}{|\mathbf{C}_k|} - d + \text{tr}(\mathbf{C}_k^{-1} \mathbf{C}_j) + (\boldsymbol{\mu}_j - \boldsymbol{\mu}_k)^T \mathbf{C}_k^{-1} (\boldsymbol{\mu}_j - \boldsymbol{\mu}_k) \right) \quad (2.21)$$

where the number of observed variables, d , does not impact the choice of experiment.

As before, we find that after approximating the parameter probability distribution with a Gaussian, the utility function depends only on the candidate experiment design, \mathbf{x}_{i+1} . We are therefore able to select the subsequent experiments using only the model and its adjoint.

We once again seek an intuitive understanding of this process by considering the case of comparing two candidate models with univariate design and observable variables. Under these conditions, Eq. (2.19) reduces to the symmetric Kullback-Leibler divergence between a pair of univariate Gaussians:

$$u = \frac{1}{2} \left\{ \left(\frac{1}{V_1^2} + \frac{1}{V_2^2} \right) (\mu_1 - \mu_2)^2 + \left(\frac{V_1 - V_2}{V_1 V_2} \right)^2 \right\} \quad (2.22)$$

where μ_j and V_j are the expected value and variance of the distributions $p(\mathbf{z}_{i+1}|\mathcal{H}_j)$. We again recover the result that MacKay [100] arrived at for scalar interpolation problems: to maximize the discrimination between models we must (i) gather data where the model predictions maximally disagree, measured relative to the confidence in their predictions, and (ii) gather data where the confidence in the models are maximally different. The first result serves to maximize the best-fit-likelihood reward on the model which fits the new data point better. The second result, which is perhaps less intuitive than the first, serves to maximize the Occam penalty on the model with more parametric flexibility.

2.3 Summary

In this chapter we have introduced the Bayesian inference framework that will be demonstrated in this thesis. In § 2.1.1 we turn the task of inferring unknown model parameters into an optimization problem, which we solve cheaply using adjoint methods. In § 2.1.2 we introduce an approximation to quantify the uncertainty in the unknown parameters. We then propagate this uncertainty to the model predictions, and estimate the systematic uncertainty in the data and the structural uncertainty in the model. In § 2.1.3 we introduce a method for quantitatively comparing candidate models in order to select the best one. Finally, in § 2.2 we extend the

Bayesian inference framework to include optimal experiment design. This allows us to use the model to forecast which experiments would be the most informative, thereby reducing the effort and cost of collecting the data required for Bayesian inference.

PART II:

APPLICATIONS IN THERMOACOUSTICS

Chapter 3

Hot wire Rijke tube

This chapter contains the results of two studies, each of which has been published as a conference paper and subsequently extended and published as a journal paper. Section 3.2 contains excerpts from two papers on generating a quantitatively-accurate model of a hot wire Rijke tube using Bayesian parameter inference and model comparison [105, 106]. This work was done in the first year of the PhD in close collaboration with Matthew Juniper. Matthew Yoko performed the experiments, processed the data, and assimilated the data into the base flow model. Matthew Juniper developed the concept and the adjoint thermoacoustics code, assimilated the data into the thermoacoustics code, and wrote the text of the original paper. Section 3.3 contains excerpts from two papers on Bayesian optimal experiment design for this rig [107, 108] ([108] is currently under review). This work was performed independently by Matthew Yoko. The experiments presented in this chapter were carried out on an existing rig, which was modified for the purposes of this study.

3.1 Introduction

The electrically heated Rijke tube is a canonical problem in thermoacoustics that is a commonly used as an initial test case for new modelling and analysis methods. The system consists of a duct and a heating element, which is typically a wire mesh that is supplied with an electrical current. While the system seems simple, it has proven challenging to model accurately.

There are several hundred papers that describe or model the Rijke tube. These are listed in the handful of reviews of this subject over the last 50 years [109–112]. The physical mechanism that causes thermoacoustic instability was correctly described by Rayleigh [4]. Later, Chu [113] re-expressed this in an elegant and complete mathematical framework based on small perturbations to the governing equations.

The thermoacoustic mechanism in the hot wire Rijke tube is described in [111, §3.3.1], which highlights the crucial role played by the time delay between velocity perturbations and heat release rate perturbations at the hot wire. If the hot wire is in the upstream (downstream) half of the tube, then this time delay causes the heat release rate to be slightly in phase (out of phase) with the acoustic pressure of the first acoustic mode, causing acoustic oscillations to be thermoacoustically driven (damped).

Using analytical methods, Carrier [114] estimated the gain and phase of the heat release rate of a hot flat ribbon in a fluctuating air stream. Lighthill [115] performed a similar analysis for a hot circular cylinder and Merk [116] for a hot wire gauze. These analytical methods show that the phase lag is caused by heat conduction through the finite thickness boundary layers around the hot element and explain why the Rijke tube requires a through-flow. This is described by Bayly [117]:

“If the blowing is too weak, the wires are surrounded by very thick jackets of stagnant air, and the fluctuations in the external flow have a comparatively small effect on the heat transfer. On the other hand, if the blowing is strong, the phase lag between the velocity fluctuations and the induced heat transfer fluctuations is small ... Although the heat transfer is more efficient with strong blowing, the absence of sufficient phase matching makes the amplification ineffective.”

The above analytical methods on simplified models, although qualitatively correct, are not quantitatively correct. Subsequent numerical simulations of heat transfer in an oscillating flow around a cylinder have shown that the time delay is two to three times greater than that predicted by Lighthill [118, 119], and that the gain and phase have more intricate dependence

on Reynolds and Strouhal numbers than can be derived analytically [120]. Instead, since the 1970's [117], researchers have tended to use simple linear relations such as the $n - \tau$ law, empirical alterations to King's law [121], or CFD simulations. This is shown in a detailed survey of heat transfer models from 1985 to 2017 [122, Table 3.1]. CFD simulations are accurate but are too expensive for simple thermoacoustic models. All simple thermoacoustic models of the Rijke tube therefore rely on heat release rate models that are, at best, only qualitatively accurate.

Most studies that aim to create quantitatively accurate models of the Rijke tube include viscous dissipation in the acoustic boundary layer along the inside wall of the tube [123, 114, 111, 118]. Viscous dissipation is greatest where the acoustic velocity fluctuation is greatest, which is at the ends of the Rijke tube. Some studies also include thermal dissipation in the acoustic boundary layer, whose effect has a similar magnitude. Thermal dissipation is greatest where the acoustic temperature (equiv. pressure) fluctuation is greatest, which is at the centre of the Rijke tube. In Rayleigh [124, §348–350] and Kinsler et al. [125, §9.5] thermal dissipation is included through an added viscosity. By modelling thermal dissipation as increased viscosity, one loses the important detail that the location of thermal dissipation differs from the location of viscous dissipation. In this chapter we model both separately so that the two mechanisms can be disentangled.

All models of the Rijke tube must include heat transfer between the hot wire and the surrounding air. The mechanism that causes heat transfer (molecular diffusion) is the same as that which causes momentum transfer, so there cannot be one without the other. The question is whether the momentum transfer is so small that it can be neglected. This momentum transfer manifests itself as an acoustic pressure drop across the heater that is, in the linear regime, proportional to the velocity at the heater. This pressure drop is considered by [123, p84], [114, Eq.4.4], [126, Eq.34] [127, §3.3] but is always set to zero on the grounds that it is small compared with the acoustic pressure amplitude. This simplifies the analysis but asserts that the drag from an object blocking the duct is negligible compared with the drag from the acoustic boundary layers of the duct, which is questionable. In this chapter we show that neither can be neglected.

All studies of the Rijke tube agree on the strong influence of the time or phase delay between velocity fluctuations and subsequent heat release rate fluctuations at the heater. This time delay can influence the frequency and the growth rate by similar absolute values, although the relative influence on the growth rate is larger because the growth rate is close to zero, while

the frequency is close to the resonant frequency of the tube. Indeed, after his famous quote about vibrations being encouraged, Rayleigh [4] writes:

“If the air be at its normal density at the moment when the transfer of heat takes place, the vibration is neither encouraged nor discouraged, but the pitch is altered.”

An efficient data-driven approach will therefore use the frequency drift as well as the growth rate drift in order to infer the time delay. This requires an accurate measurement of the speed of sound in the tube.

In this chapter we firstly construct a quantitatively accurate model of an electrically heated Rijke tube by (i) devising qualitatively-accurate physics-based candidate models of each component of the rig; (ii) collecting experimental data using experiments carefully designed to identify the unknown model parameters; (iii) assimilating the experimental data into the candidate models to infer their unknown parameters; (iv) quantifying the evidence (marginal likelihood) for each candidate model and selecting the best candidate model, and (vi) repeating for the next component until the model of the rig is complete. We then revisit the dataset collected for this work, and use it to demonstrate Bayesian optimal experiment design and assess how much data is strictly required in the Bayesian framework.

3.1.1 Experimental configuration

The Rijke tube, shown in Fig. 3.1, is constructed from a 1 m length of stainless steel tube with an inner diameter of 47.4 mm and a wall thickness of 1.7 mm. The tube is mounted vertically and the ends are open to ambient conditions. A photograph of the rig is provided in Fig. B.1 in Appendix B.

The heater is shown in detail in Fig. 3.2, and a photograph is provided in Fig. B.2 in Appendix B. It is constructed from ceramic rings wrapped with 0.6 mm nichrome wire. It is mounted on two support rods and inserted into the upstream (bottom) end of the tube. The support rods are attached to an electrically driven traverse so that the heater position can be controlled through the data acquisition system (DAQ). The heater is connected to a programmable DC power supply (Elektro-Automatix EA-PSI 5080-20), which is controlled using a PID controller implemented in LabVIEW. This allows a desired power to be supplied to the heater, with compensation for the changing resistance of the wire as it heats up.

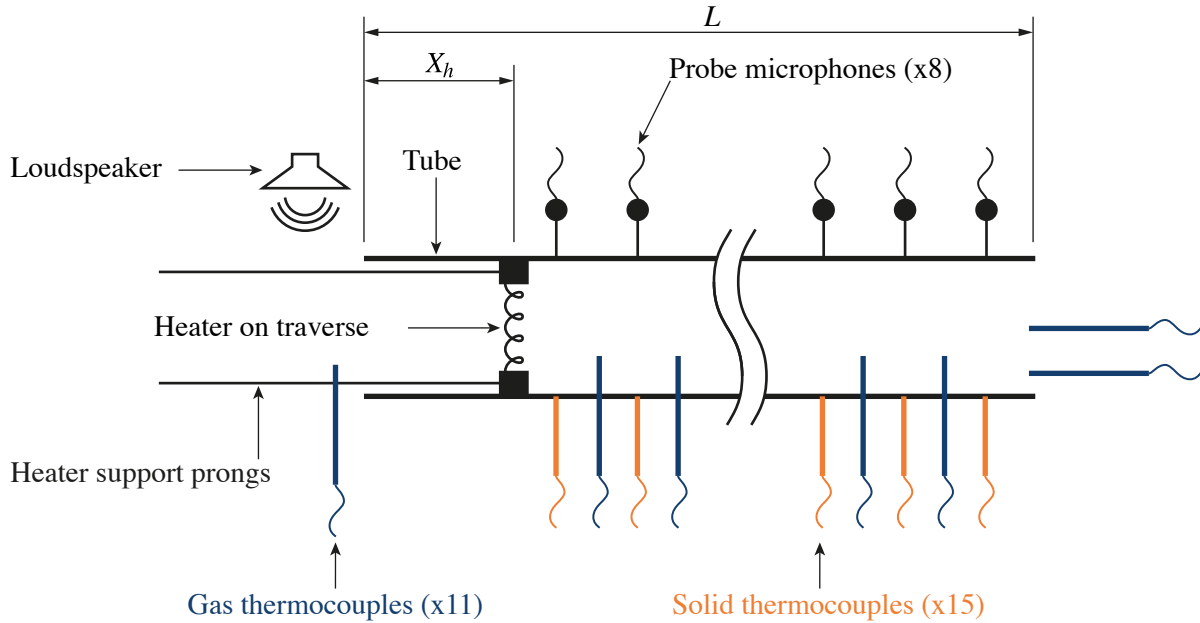


Fig. 3.1 Diagram of the Rijke tube, rotated for convenience. The tube length, $L = 1$ m. The heater is mounted on a traverse and inserted into the bottom of the tube to a distance of X_h . Eight probe microphones are placed along the length of the tube. Eleven thermocouples monitor the gas temperature (blue) and another fifteen monitor the tube temperature (orange). A loudspeaker mounted at the tube inlet provides acoustic forcing.




Fig. 3.2 (a) Top view, (b) side view, and (c) isometric view of the heater, which consists of two identical concentric annular ceramic plates, each wound with nichrome wire. It is held in place by two threaded support prongs (shown). The dimensions are $d = 47$ mm, $d_i = 31.6$ mm, $d_w = 0.6$ mm, $t = 5$ mm, $h = 5$ mm, $d_p = 3$ mm. The power is supplied to the nichrome wire by two fabric-insulated copper wires (not shown), which each have diameter mm.

A loudspeaker (Visaton FRS 8) is mounted at the upstream end of the tube and connected to an audio amplifier (Stage Line STA-500). We use the loudspeaker to force the system so that its response can be measured. The forcing signal is a sine wave generated by LabVIEW.

Instrumentation

We record the acoustic pressure using eight probe microphones (G.R.A.S. 40SA), with the probes inserted through ports in the tube wall. The microphones are distributed with 100 mm spacing, starting at 250 mm from the upstream end of the tube. Each of the microphones is connected to an amplifier (G.R.A.S 12AA), which is then connected to the acquisition system.

The rig is instrumented with K-type thermocouples (0.7 mm diameter) to measure (i) the temperature of the air inside the tube, and (ii) the temperature of the tube walls. Eight thermocouples are inserted through ports in the tube wall to measure the internal air temperature. They are mounted with a spacing of 100 mm, starting at 200 mm from the bottom of the tube. In addition to these eight thermocouples, two thermocouples are mounted vertically with the beads placed on the outlet plane of the tube. Fifteen thermocouples are bonded to the exterior of the tube with a spacing of 50 mm, starting at 150 mm from the bottom of the tube. An additional two thermocouples are mounted 5 mm from either end of the tube. Finally, one thermocouple is mounted 100 mm below the inlet to record the ambient temperature.

Automation & Acquisition

The experiment automation and data acquisition is handled by a PC running LabVIEW. The PC has two National Instruments data acquisition cards (PCI-6035E & PCIe-6343), each of which is linked to an eight channel terminal block (BNC-2110) for input/output. We record the microphone data at a sampling rate of 10 kHz. This is well above the Nyquist frequency for the fundamental mode of the rig, and is used because the higher sampling rate improves the accuracy of the data processing. Temperature data is logged through a bank of four thermocouple DAQ devices (2 x Omega TC-08 & 2 x Pico TC-08).

The LabVIEW code automatically runs through a pre-determined matrix of heater positions and heater powers. For each heater position and power setting, the code repeats the experimental procedure 40 times so that the random error can be quantified. This allows a large, densely sampled dataset to be collected in a relatively short time.

3.1.2 Experimental procedure & data processing

At each heater position and heater power we conduct two tests. The first test measures the system's growth/decay rate and natural frequency following a brief period of harmonic forcing.

The second test uses an acoustic impulse to measure the sound speed. We now describe these tests in detail.

Forced response test

The forced response tests begin by triggering the acquisition system to start recording data from the microphones. A short delay (± 0.2 s) is then enforced to ensure that the acquisition system is running correctly. Following this, the loudspeaker generates a tone to force the tube near its natural frequency, which we refer to as a ‘ping’. The forcing is faded in over a period of 1 s to minimize transients, which occur when a full amplitude signal is commanded instantly. The forcing is sustained for 6 s and then abruptly stopped. The acquisition records for a further 4 s to capture the decay, following which the acquisition terminates.

A sample of the raw data obtained from a forced response test is shown in Fig. 3.3(a). This shows the pressures measured by each of the eight microphones plotted against time. The raw signal is processed as follows:

1. A Fourier transform is performed with four overlapping Hann windows to obtain the complex pressure as a function of time.
2. The decaying portion of the signal is isolated. This is defined as the period between the termination of the forcing command and the point where the pressure magnitude reaches the noise floor.
3. The growth rate is determined by performing linear regression on the logarithm of the decaying signal. This can be seen in Fig. 3.3(b.i).
4. The natural frequency is determined by calculating the mean angular frequency during the decaying period. This can be seen in Fig. 3.3(b.ii).

Sound speed test

A previous study on the same rig inferred the local sound speed from the temperature measurements [92]. We subsequently found that this method was unreliable due to excessive conduction of heat along the stem of the thermocouple, arising from the use of relatively thick thermocouples to measure the temperature of a low velocity flow. Additionally, the thermocouples near the heater are heated through radiation, biasing the sound speed estimation. To overcome this, we devise a new method to measure the local sound speed using the microphone array.

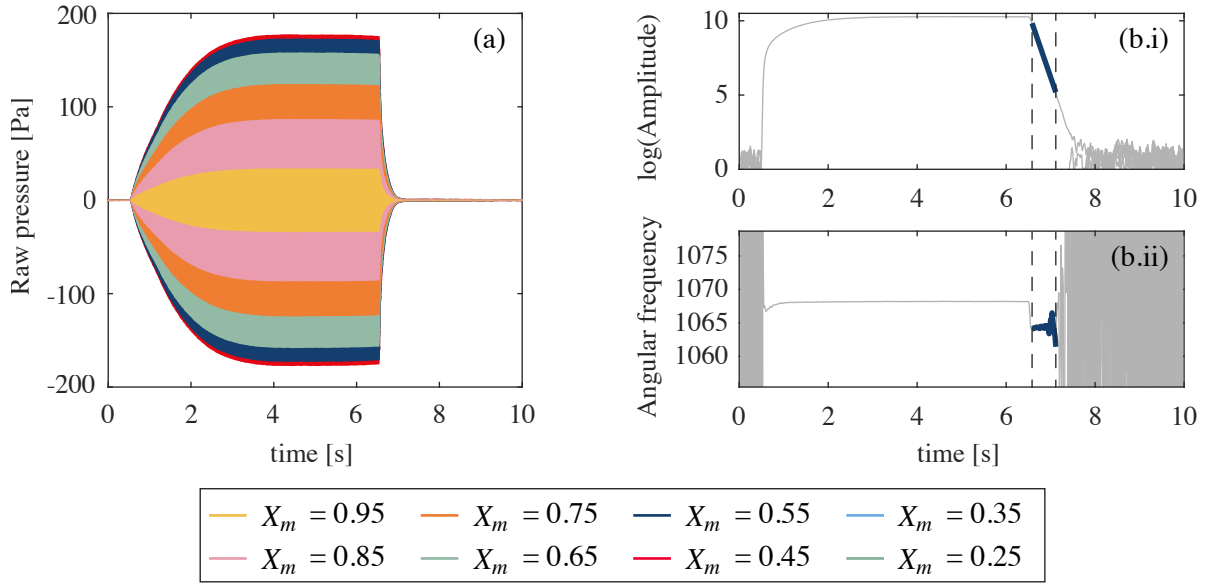


Fig. 3.3 Raw and processed signals from a forced response test. (a) The raw signals from each microphone are plotted against time. (b.i) The log of the absolute value of the Fourier decomposed pressure is plotted against time with the decaying portion highlighted. (b.ii) The angular frequency is plotted against time with the region during decay highlighted.

After each ping, the loudspeaker sends an impulse down the tube, which we refer to as a ‘click’. The 8 probe microphones measure the response along the tube at 62.5 kHz, which is the maximum sampling rate of the acquisition system. We find the phase shift that maximizes the cross-correlation function between the microphones (i.e. brings the measured impulses optimally in phase). This process yields the time at which the impulse arrives at each microphone, and therefore an estimate of the local sound speed.

The acquisition system used for these tests scans the channels sequentially, introducing a different acquisition delay to each channel. We measure these delays by performing several clicks at ambient conditions, where the true sound speed is known. The acquisition delay is then subtracted from the measured transit times in the hot conditions, removing a significant source of systematic error.

3.1.3 Experiment design

The experiments were designed to allow the unknown model parameters to be inferred sequentially, because inferring all unknown parameters at once typically leads to an ill-posed problem. Firstly we collect data from the empty tube, which allows the reflection coefficients and visco-thermal dissipation to be inferred without the influence of the heater. Secondly we

traverse the heater support prongs through the tube with the heater removed, which allows the visco-thermal dissipation and blockage of the support prongs to be inferred. Thirdly we traverse the cold heater through the tube, which allows the visco-thermal dissipation and blockage of the heater to be inferred. Finally, we traverse the heater through the tube, and sweep through several heater powers at each heater position. This allows the parameters of the fluctuating heat release rate models to be inferred. A summary of the tests conducted is provided in table 3.1.

Table 3.1 The number of experimental measurements (40) assimilated at each operating point. Column 1 describes the type of experiment: C1 and C2 denote experiments with the prongs and heater attached but switched off; C3 denote experiments with the prongs only; H denotes experiments with the prongs and heater attached and switched on. Column 2 contains a check mark if the thermocouples were present. Column 3 contains the heater power in Watts. Columns 4 to 23 contain the number of experimental measurements taken at the heater position shown in the second row, where E denotes the empty tube, in which the heater and prongs were removed.

	T	Q	Heater position (cm)																				
			E	5	10	15	20	25	30	35	40	45	50	55	60	65	70	75	80	85	90	95	
C1	-	0	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-	
C2	-	0	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	
C3	-	0	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	
H	✓	0	40	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	7	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	15	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	30	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	50	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	80	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	130	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-
H	✓	180	-	-	-	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	-	-

3.1.4 Physics-based model of the hot wire Rijke tube

As was indicated in the introduction, the hot wire Rijke tube exhibits phenomena which occur on two distinct time scales. The heating of the Rijke tube, which drives the convective flow, exhibits a long timescale, $\mathcal{O}(1000s)$. The decay of the acoustic oscillations occur on a short timescale, $\mathcal{O}(1s)$. At the long timescale we construct a model of the base flow, into which we assimilate measurements of the local sound speed, the local temperature of the gas, and the local temperature of the tube walls. At the short timescale we construct a model of the thermoacoustic oscillations, into which we assimilate measurements of the decay rate, the natural frequency, and the Fourier-decomposed pressure at the microphone locations.

Base flow model

The model of the long timescale flow is an unsteady 1D model of the base flow, including conjugate heat transfer between the gas and the tube walls. This model is an extension of a model from a previous study [92], into which we have included conductive cooling of the tube through contact with the support structure, conductive heating of the tube through thermal contact with the heater, heat loss from the wires upstream of the heating element, and variation of Nusselt number along the inner wall. Four parameters characterize (i) the inviscid drag coefficient of the heater, (ii) the Nusselt number on the outer surface of the tube, (iii) the thermal resistance at the tube mounts, and (iv) the proportion of supplied power that conducts through the tube wall. Two further parameters characterize the Nusselt number distribution inside the tube. For a given configuration, the model outputs the velocity of the convective flow, which cannot be measured directly.

Thermoacoustic model

Acoustic waves are modelled as forward-travelling waves, $f(t - x/c)$, and backward-travelling waves, $g(t + x/c)$, in N acoustic elements within the tube [109]. In element i , the pressure is $p_i = f_i + g_i$ and the velocity is $u_i = (f_i - g_i)/(\rho_i c_i)$. At the interfaces between the acoustic elements, the complex wave amplitudes are related through jump conditions for the momentum and energy equations. The linear influence of all components of the network model can be expressed in terms of local linear feedback from velocity or pressure into the momentum or energy equations [109, §VI(A)]. Labelling these feedback coefficients k_{**} , the jump conditions are:

$$p_{i+1} - p_i = -k_{mu} u_i - k_{mp} p_i \quad (3.1)$$

$$u_{i+1} - u_i = -k_{eu} u_i - k_{ep} p_i \quad (3.2)$$

For example, viscous dissipation is modelled as local feedback from the velocity into the momentum equation (k_{mu}), thermal dissipation is modelled as local feedback from the temperature (equiv. pressure) into the energy equation (k_{ep}), and the heat release rate from the wire is modelled as local feedback from the velocity into the energy equation (k_{eu}). Local feedback from the pressure into the momentum equation, k_{mp} , is not required for this study. In § 3.2.2 these local feedback coefficients are derived from candidate physics-based models

and are expressed in terms of those models' parameters. Wave reflection at the upstream and downstream ends of the tube are modelled by complex reflection coefficients, R_u and R_d .

The equations are converted to the frequency domain through modal decompositions such as $f(t - x/c) = F e^{st} e^{-sx/c}$, where $-is$ is the complex angular frequency of oscillations. There are $2N$ unknown amplitudes (N each for F and G) and $2N$ constraints ($2(N - 1)$ jump conditions and 2 reflection conditions). This creates a nonlinear eigenvalue problem for the eigenvalue s , which is solved with Newton iteration. We calculate the corresponding eigenfunction $P(x, s)$ from F and G . We then use first order [128] and second order [129] adjoint methods to obtain the first and second derivatives of s and P with respect to all the local feedback coefficients k_{**} and reflection coefficients R_* . We then combine these derivatives to obtain the first and second derivatives of s and P with respect to the model parameters of the components of the network model.

The model is shown schematically in Fig. 3.4. The top illustration shows a simplified diagram of the rig. The middle illustration shows a schematic diagram of the network model, with the upstream and downstream reflection coefficients labelled R_u and R_d respectively, and the model for the heater labelled \mathcal{F} . The boundaries and heater model are connected by a straight duct with visco-thermal dissipation with a strength of η . The bottom illustration shows the structure of the network model for this case. The boundary conditions are highlighted in blue, the heater jump condition is highlighted in orange, and the visco-thermal dissipation jump conditions are highlighted in teal. The forward and backward travelling waves, F_i and G_i are shown within one of the acoustic elements.

3.2 Generating a quantitatively accurate model of a hot wire Rijke tube

To generate a quantitatively accurate model of the rig, we begin by calibrating the base flow model by inferring its unknown parameters from data. We then construct a quantitatively accurate model of the thermoacoustics of the cold rig by assembling several component models, each of which is the candidate model with the highest marginal likelihood given the experimental data. Finally, we consider several candidate models for the fluctuating heat release rate, and select the best one. The network model (§ 3.1.4) of the empty tube contains $N = 40$ elements with equal length. When the heater is added to the model, the acoustic element at the heater location is split into two elements, with jump conditions (Eqs. (3.1)-(3.2)) at the heater.

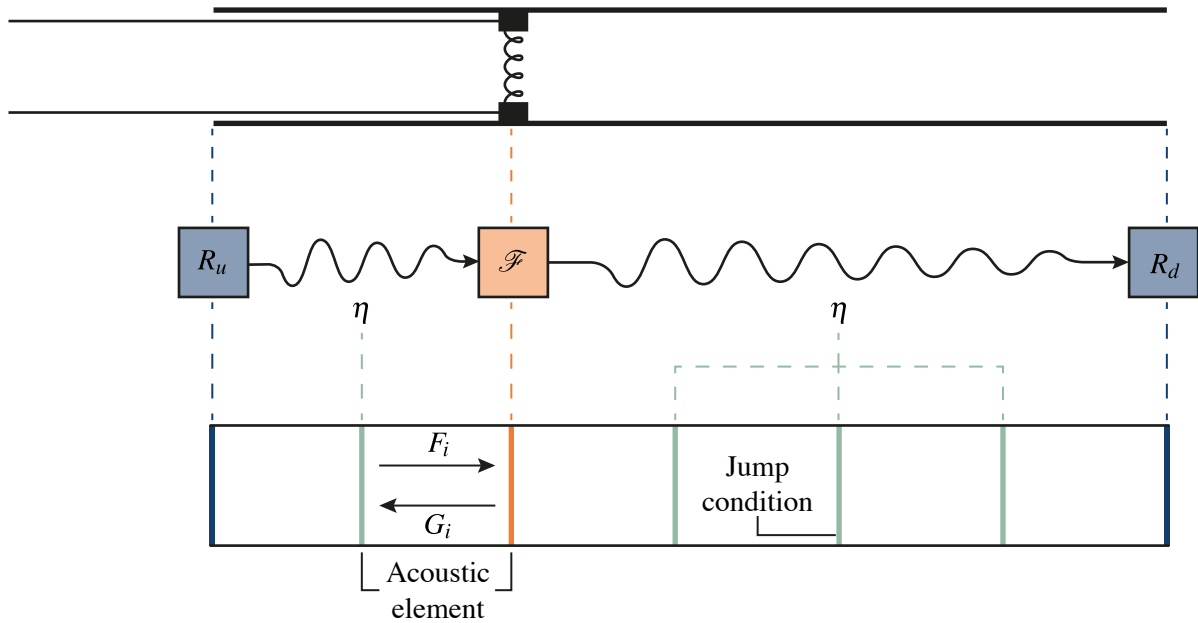


Fig. 3.4 Schematic diagram of the network model for the hot wire Rijke tube. The top illustration shows a simplified diagram of the rig. The middle illustration shows a schematic diagram of the network model. The bottom diagram shows the structure of the network model for this case. R_u and R_d are the upstream and downstream reflection coefficients respectively, \mathcal{F} is the model for the heater, η is the strength of the visco-thermal dissipation, and F_i and G_i are the forward and backward travelling waves within one of the acoustic elements.

The same procedure is followed when the 8 thermocouples are added. The network model therefore contains between 40 and 50 acoustic elements, N , depending on the configuration. We have checked that the model predictions do not change significantly with N for $N > 40$.

3.2.1 Calibrating the base flow model

The four unknown parameters of the base flow model (§ 3.1.4) are inferred from the sound speed and temperature measurements collected from the hot experiments (rows labelled H in Table 3.1). Fig. 3.5 shows (a) the velocity at the heater and (b) the local sound speed, as functions of the heater position, X_h , and heater power, Q_h . As X_h increases, the length of the column of hot air above the heater decreases, so the buoyancy force driving the flow decreases, causing the velocity at the heater, \bar{u}_h , to decrease (Fig. 3.5(a)) and, at a given heater power, the temperature jump across the heater to increase (Fig. 3.5(b)).

There is clear evidence of measurement and model error in Fig. 3.5: (i) \bar{u}_h is expected to decrease smoothly with X_h but does not; (ii) the maximum of $c(x)$ is expected to increase smoothly with X_h but does not at $X_h = 0.75$ m. The main source of error is the fact that the

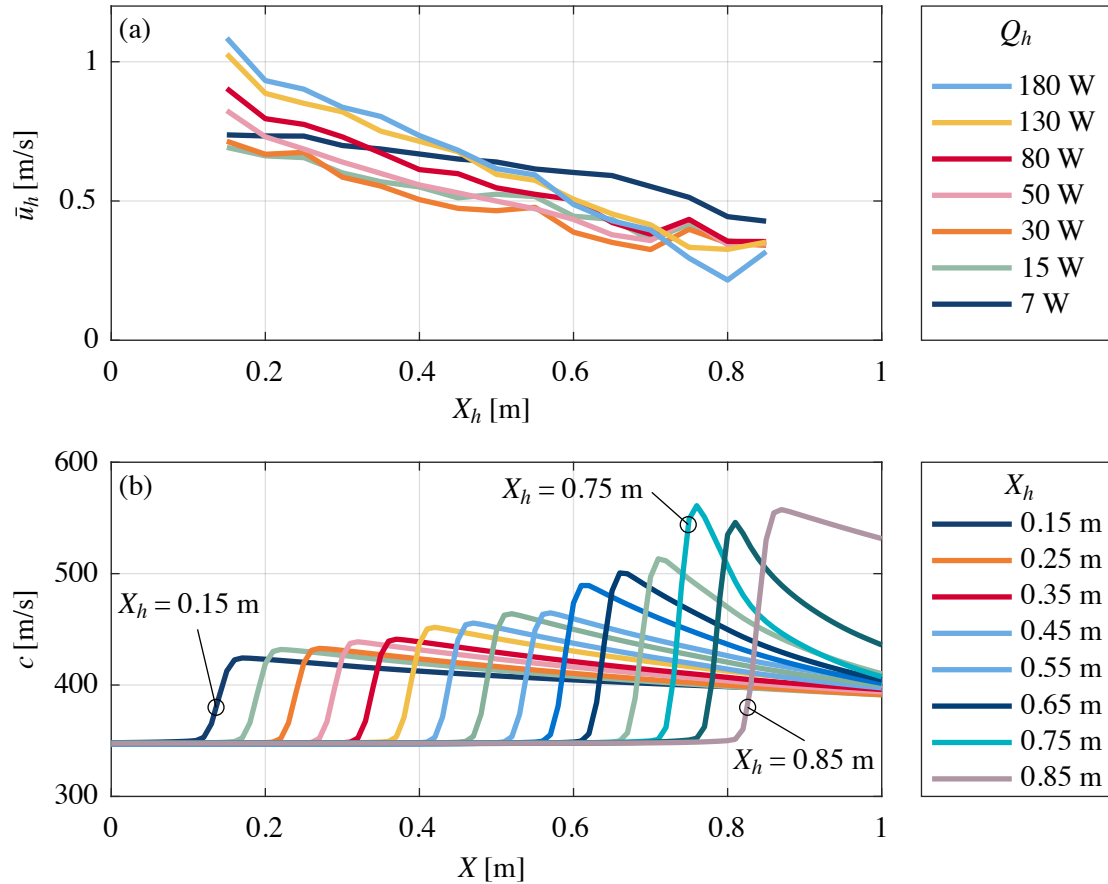


Fig. 3.5 (a) Velocity at the heater, \bar{u}_h , as a function of heater position, X_h , and heater power, Q_h . (b) Local speed of sound, $c(x)$, as a function of heater position, X_h , at $Q_h = 180$ W. These results are calculated from the base flow model (§ 3.1.4) for the long timescale flow, after assimilating the local sound speed and temperature measurements.

number of microphones downstream of the heater decreases as X_h increases, so there are fewer measurements of the downstream sound speed during the click test. It is impractical to remove all this error. Instead, we group the measurement and model errors and check a-posteriori that the combined error is small compared with the trends in the model predictions.

We note that \bar{u}_h is derived from several measured quantities, whose errors accumulate into the error in \bar{u}_h , and that the model assumes uniform flow at the heater location, which is a simplification. We must therefore be aware that \bar{u}_h contains more systematic uncertainty and model error than the other variables.

3.2.2 Characterizing the cold acoustics

Visco-thermal boundary layer drag and reflection coefficients

To model the viscous dissipation in the boundary layers inside the tube, we derive an expression for the shear stress in the boundary layer. Tijdeman [130] provides an analytical expression for the acoustic velocity in a cylindrical duct as a function of axial and radial position [130, Eq. (B19)]. Using this expression, we obtain:

$$\tau'_w = -\mu \left. \frac{\partial u'}{\partial y} \right|_w = \frac{2\mu}{D} \left[\frac{i^{3/2} \sigma J_1(i^{3/2} \sigma)}{J_0(i^{3/2} \sigma) - 1} \right] u' \quad (3.3)$$

where τ'_w is the fluctuating shear stress at the wall, μ is the viscosity of the air, D is the tube diameter, u' is the fluctuating velocity in the bulk. The shear number, σ , is defined as $\sigma = 1/2D\sqrt{\rho s_i/\mu}$, where ρ is the mean density of the air, and s_i is the angular frequency. J is the Bessel function of the first kind.

This stress acts on the perimeter area $\pi D \delta x$, where δx is the length of an element in the network model. The pressure on either side of this element acts on a cross-sectional area $\pi D^2/4$. The viscous boundary layer drag is therefore modelled through a jump condition in the pressure (Eq. (3.1)) via a local feedback coefficient:

$$k_{mu} = \frac{8\mu \delta x}{D^2} \left[\frac{i^{3/2} \sigma J_1(i^{3/2} \sigma)}{J_0(i^{3/2} \sigma) - 1} \right] \quad (3.4)$$

Similarly, we use Tijdeman's expression for the acoustic temperature in a cylindrical duct [130, Eq. (B22)] to obtain the local feedback coefficient for thermal dissipation:

$$k_{ep} = \frac{8\mu \delta x}{\rho^2 D^2 c^2} \frac{\gamma - 1}{\text{Pr}} \left[\frac{i^{3/2} \sigma \sqrt{\text{Pr}} J_1(i^{3/2} \sigma \sqrt{\text{Pr}})}{J_0(i^{3/2} \sigma \sqrt{\text{Pr}})} \right] \quad (3.5)$$

where c is the speed of sound, γ is the ratio of specific heats, and Pr is the Prandtl number. For the assimilation process, we multiply k_{mu} and k_{ep} by a real constant, η , which we infer from the data. The value of η would be unity if this model were perfect.

The ends of the tube are identical so, when the tube is empty, we assume that their reflection coefficients are the same: $R_u = R_d = R$. For both empty tube experiments (column E in Table 3.1), Levine-Schwinger's (LS) calculations [131] give $|R|_{\text{LS}} = 0.9975$ and $\angle R_{\text{LS}} = 3.0550$. If we set $\eta = 1.0$ and assimilate R directly from the experimentally-measured growth rates and frequencies of the two empty tube experiments, we obtain $|R| = 0.9953$ and 0.9952 ,

and $\angle R = 3.0597$ and 3.0641 . Alternatively if we set $R = R_{LS}$ and assimilate η we obtain $\eta = 1.1235$ and $\eta = 1.1198$. It might be tempting to accept a 0.23% discrepancy in $|R|$ or a 12% discrepancy between the boundary layer model and experiment, but careful use of Bayesian inference enables us to improve on both values. For demonstration, we will now infer R and η simultaneously, first when there is insufficient information in the data, and second when there is sufficient information. We set one standard deviation of the measurement uncertainties to be 0.1 rad s^{-1} for the decay rate and 1 rad s^{-1} for the frequency. For the prior expected values we set $R = R_{LS}$ and $\eta = 1.0$. We set large prior variances.

Fig. 3.6(a) shows the first inference problem, in which we attempt to infer $|R|$, $\angle R$ and η from the decay rates and frequencies only. Each set of axes shows the posterior joint distribution between each pair of parameters in the form of ellipses denoting 1, 2, and 3 standard deviations from the expected values. We see that rather than learning unique values for each parameter, we have learned a set of correlations between the parameter pairs, indicated by diagonally stretched ellipses. This is because we have attempted to infer three parameters from two independent observations. As one would expect, the data can be explained by a wide range of $|R|$, $\angle R$ and η , where the main source of uncertainty in one is the value of the other.

To disentangle the parameters, we need to (i) devise additional experiments to provide more information about one of the parameters, or (ii) provide stronger prior information about one of the parameters, if it is available. We demonstrate (i) in Fig. 3.6(b), which shows the posterior expected values and covariances when we attempt to infer R and η from the decay rates, frequencies, and the relative amplitudes and phases of all 8 microphones. Comparing Figs. 3.6(a) and (b), we see that the information provided by the relative amplitudes and phases of the microphones has allowed us to decorrelate the values of $|R|$ and $\angle R$. This is because the observations of $|P|$ provide information about the distance between the pressure node and the end of the tube, which uniquely identifies $\angle R$. The values of $|R|$ and η , however, remain correlated, although the correlation is weaker than in Fig. 3.6(a). In theory, a unique value for $|R|$ can be obtained from the observations of $\angle P$ near the ends of the tube, which indicates the work done at the ends of the tube. However, it was not possible to phase-calibrate the microphone array sufficiently accurately to resolve this effect. The microphones could only be repeatably phase-calibrated to within $\mathcal{O}(10^{-2})$ radians, but the model suggests that the phase shift at the ends of the tube is $\mathcal{O}(10^{-3})$ radians. We can therefore not decorrelate the parameters any further without enforcing a stronger prior for one of the parameters.

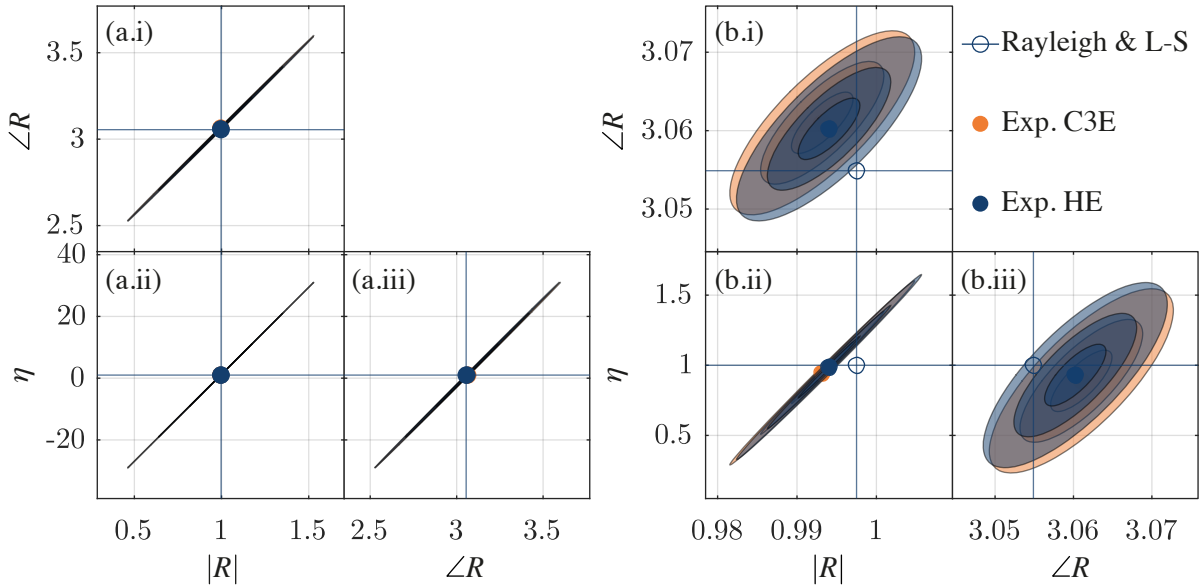


Fig. 3.6 Posterior joint probability distributions after assimilating data from the two empty tube experiments labelled E in Table 3.1. The contours show 1, 2, and 3 standard deviations from the assimilated expected values for (i) $\mathbf{a} = (|R|, \angle R)$, (ii) $\mathbf{a} = (|R|, \eta)$, and (iii) $\mathbf{a} = (\angle R, \eta)$. (a) The data contains decay rate and frequency only, measured by the microphone at $x/L = 0.75$. (b) The data contains the decay rate, frequency, and relative pressures of all 8 microphones. This shows that the posterior probability distributions collapse only if the data contains sufficient information for that model. The values predicted by Tijdeman [130] and Levine-Schwinger [131] are shown for comparison.

We see from the posterior expected values in Fig. 3.6(b) that η is remarkably close to 1.0, indicating that Tijdeman's model performs well for this experiment, but that the inferred R deviates slightly from that predicted by the Levine-Schwinger model. This slight deviation is to be expected because our tube is not infinitely-thin and the exterior volume is not infinitely large, which are both assumptions of the Levine-Schwinger model. We calculate a correction factor to the Levine-Schwinger model, k_{LS} , such that $R = k_{LS} \times R_{LS}$ and hardwire this empirical coefficient into the model for R used in the rest of this section. This allows us to use the corrected Levine-Schwinger model to predict the downstream reflection coefficient for the hot experiments.

Drag and blockage by the heater prongs

The heater is held by two 3 mm diameter threaded prongs and fed by two 4 mm diameter fabric-coated copper wires. These are inserted from the upstream end ($x = 0$). The black dots in Fig. 3.7 show the measured growth rates and frequencies as the prongs are inserted without the heater attached. On physical grounds, we may suspect that these variations are caused by

the extra visco-thermal dissipation and blockage caused by the prongs and wires. We may not be sure, however, which physical effects need to be included in a low order model. On the one hand, we need to include all physical effects having a significant influence. On the other hand, if we include too many effects, it will be impossible to distinguish between them with the available data, as was shown in Fig. 3.6(a). In this section we propose seven reasonable physics-based models, assimilate the data into those models, and use the marginal likelihood to identify the most likely model given the data.

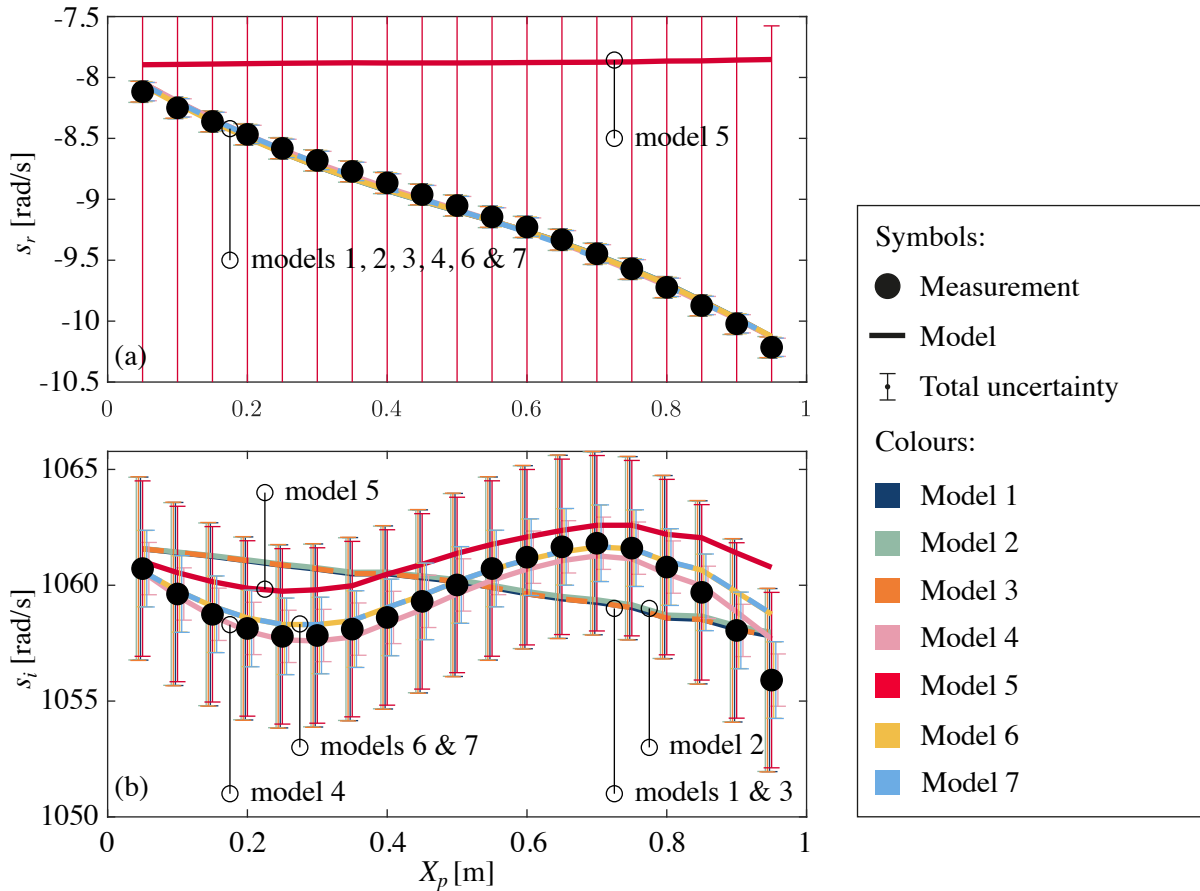


Fig. 3.7 (a) Growth rate, s_r , and (b) angular frequency, s_i , as the heater prongs are inserted through the tube from $x = 0$ to $x = X_p$. The black dots show the experimental measurements. The coloured lines show the predictions from the models in Table 3.2 after assimilating the experimental data. The coloured error bars show the inferred total uncertainty, which includes random and systematic experimental error, and structural model error. For each measurement (black dot) the error bars are ordered from model 1 on the left to model 7 on the right.

We denote P_p as the sum of the perimeters of the prongs and heater wires, and P_t as the perimeter of the tube. At each axial location containing the prongs, their viscous dissipation and

thermal dissipation are modelled by multiplying the viscous dissipation and thermal dissipation of the tube's boundary layer at that location by P_p/P_t , and then following the procedures outlined in the next paragraph. The effect of their blockage is modelled as an area change at the heater position x_h . We follow the analysis in [132, §4.4.2], which allows the effect of the area change to be expressed as local feedback from u into the energy equation: $k_{eu_h} = 1 - (\rho_1 A_1)/(\rho_2 A_2)$.

Model 1 multiplies the local viscous dissipation, k_{mu} , and local thermal dissipation, k_{ep} , by the same real parameter to account for a possible change in the magnitude of drag between the tube walls and the prong walls due to the different surface roughness and materials. Model 2 multiplies the viscous dissipation and thermal dissipation by the same complex parameter to account for possible changes in the magnitude and phase of each type of dissipation. Model 3 multiplies the viscous dissipation by one real parameter and the thermal dissipation by another real parameter to account for possible independent changes in the magnitudes of the viscous and thermal dissipation. Model 4 multiplies the viscous dissipation by one complex parameter and the thermal dissipation by another complex parameter to account for possible independent changes in the magnitude and phases of the viscous and thermal dissipation. Model 5 multiplies k_{eu_h} by one real parameter to account for the blockage but contains no visco-thermal dissipation from the prongs. Model 6 combines models 1 and 5. Model 7 combines models 3 and 5.

We simultaneously infer the parameters and the total uncertainty in the model and data. This gives the fairest comparison of the marginal likelihoods because, if the measurement uncertainties are fixed to the random experimental error, then simple models (which tend to have larger model error) become overwhelmingly unlikely even if they miss just a few of the datapoints.

Fig. 3.7 shows the model predictions *vs* experimental measurements, while Table 3.2 summarizes the details of each model and shows their best fit likelihood (BFL), Occam Factor (OF), and marginal likelihood (ML). Models 1, 2, 3, and 5 have large negative values of $\log(\text{BFL})$ and $\log(\text{ML})$, meaning that they fit the data badly and are not supported by the data. Models 1, 2, and 3 (no blockage) cannot match the experimental frequency, which implies that a successful model must include the prong blockage. Model 5 (no drag) cannot match the experimental growth rate, which implies that a successful model must include the prong drag. Models 4 (complex visco-thermal dissipation but no blockage) and model 6 and 7 (real drag and blockage) have large positive values of $\log(\text{BFL})$, meaning that they fit the data well. Of these, model 6 has the largest value of ML because it achieves this good fit with just two parameters, while model 7 contains three parameters and model 4 contains four parameters.

It is reassuring that the model with the most evidence (the highest marginal likelihood) is the simplest model that contains the expected physics. While a human may have been able to identify this model by hand, this Bayesian framework provides a rigorous and quantifiable measure of its suitability compared with other plausible candidate models.

Table 3.2 Description and performance of seven models for the drag and blockage caused by the heater prongs. Column 2 contains the number of parameters, \mathbf{a} , in each model. Columns 3 and 4 describe how the viscous prong drag, k_{mu_p} , and thermal prong drag, k_{ep_p} , are modelled: as the tube perimeter boundary layer drag multiplied by a real parameter, a complex parameter, or zero. Column 5 describes how the blockage of the prongs, k_{eu_h} , is modelled: as a real parameter or zero. Column 6 contains the log best fit likelihood (BFL) per datapoint, where $\text{BFL} = p(D|\mathbf{a}_{\text{MP}}, \mathcal{H}_i)$; a more positive $\log(\text{BFL})$ implies a better fit to the data. Column 7 contains the log Occam Factor (OF) per datapoint, where $\text{OF} = p(\mathbf{a}_{\text{MP}}|\mathcal{H}_i) (\det(\mathbf{A}/2\pi))^{-1/2}$; a more negative $\log(\text{OF})$ means that the parameter space has collapsed more when the data arrive. This indicates that the model requires excessive precision in the parameters in order to fit the model, which is a sign of an overly-complex model. Column 8 contains the log Marginal Likelihood (ML) per datapoint, where $\text{ML} = \text{BFL} \times \text{OF}$; a more positive $\log(\text{ML})$ implies higher evidence for the model, given the data. The model with the largest ML is the simplest model that fits the data well, and is taken as the most likely model

Model	Params	k_{mu_p}	k_{ep_p}	k_{eu_h}	$\log(\text{BFL})$	$\log(\text{OF})$	$\log(\text{ML})$
1	1	real		zero	-0.3635	-0.3561	-0.7197
2	2	complex		zero	-0.3575	-0.5567	-0.9141
3	2	real	real	zero	-0.3552	-0.4151	-0.7702
4	4	complex	complex	zero	+0.9452	-1.2319	-0.2867
5	1	zero	zero	real	-3.7064	-0.1736	-3.8799
6	2	real		real	+0.6705	-0.5767	+0.0938
7	3	real	real	real	+0.7010	-0.8191	-0.1181

Drag and blockage by the heater itself

The two heater traverses (rows C1 and C2 of Table 3.1) were performed on different days with the heater attached but switched off and with no thermocouples present. The teal circles in Fig. 3.8 show the growth rates and frequencies as the heater is traversed through the tube. By comparing them with the blue circles, which are for the prong only experiments, it can be seen that the heater significantly changes the growth rate and frequency when at the ends of the tube and slightly changes the growth rate when at the centre of the tube. In this section, we propose physical reasons for this and compare the evidence for three corresponding physical models.

The heater (Fig. 3.2) consists of two orifice plates, one behind the other. From [132, §4.4.3], the pressure drop across a single orifice plate with length L and cross-sectional area A_d in a pipe of radius A is given by $\Delta p = -\rho L_{\text{eff}}(A/A_d)(du'/dt)$, due to the inertia of the air in the

orifice plate. The effective length, L_{eff} , equals $L + 2\delta$ where δ is typically $(8/3\pi)(A_d/\pi)^{1/2}$. The local feedback mechanism is therefore from the velocity into the momentum equation. When converted to the frequency domain, the local feedback coefficient is $k_{mu_h} = \rho L_{\text{eff}}(A/A_d)s$ where s is the complex angular frequency, whose imaginary part is much greater than its real part.

Nichrome wire is wound around each orifice plate. It is reasonable to suppose that the viscous dissipation from the wire is non-zero because momentum transfer from the wire relies on the same physical mechanism as heat transfer, and heat transfer will be crucial in later sections. The question is whether the data shows that the contributions to the viscous dissipation coefficient (k_{mu_h}) and thermal dissipation coefficient (k_{ep_h}) are negligible compared with the contribution to k_{mu_h} from the orifice. (Note that the orifice and the viscous dissipation both contribute to the same local feedback coefficient, although the orifice affects only the imaginary component.)

We evaluate three models. Model 1 models the orifice alone, by assimilating an imaginary value of k_{mu_h} and setting the other local feedback coefficients ($k_{mp_h}, k_{ep_h}, k_{eu_h}$) to zero. Model 2 is the same as model 1 but allows k_{mu_h} to be complex, which models the orifice and the viscous dissipation, but not the thermal dissipation. Model 3 models the orifice and the wire's visco-thermal dissipation by allowing complex non-zero values of k_{mu_h} and k_{ep_h} . Table 3.3 shows the BFL, OF, and ML of all three models. Of these, the data supports model 3. Although not shown here, model 1 fails because it cannot model the influence of the heater on the growth rate at the ends of the tube, where the acoustic velocity is high, and the viscous dissipation is important. Model 2 fails because it cannot model the influence of the heater on the growth rate at the centre of the tube, where the acoustic temperature fluctuation is high, and the thermal dissipation is important. Model 3 models the orifice blockage, the viscous dissipation, and the thermal dissipation. It is worth noting that the imaginary part of k_{mu_h} is six times greater than its real part, indicating that the orifice blockage is significantly more influential than the viscous dissipation.

The experiments in row H-0 of Table 3.1 were performed on different days with the heater attached but switched off and with thermocouples in place. The pink circles in Fig. 3.8 show the growth rates and frequencies of these experiments. By comparing them with the teal circles, which are for the same experiments with the thermocouples removed, it can be seen that the thermocouples significantly affect the growth rate and slightly affect the frequency. We model the thermocouples as identical local visco-thermal dissipation elements and assimilate their

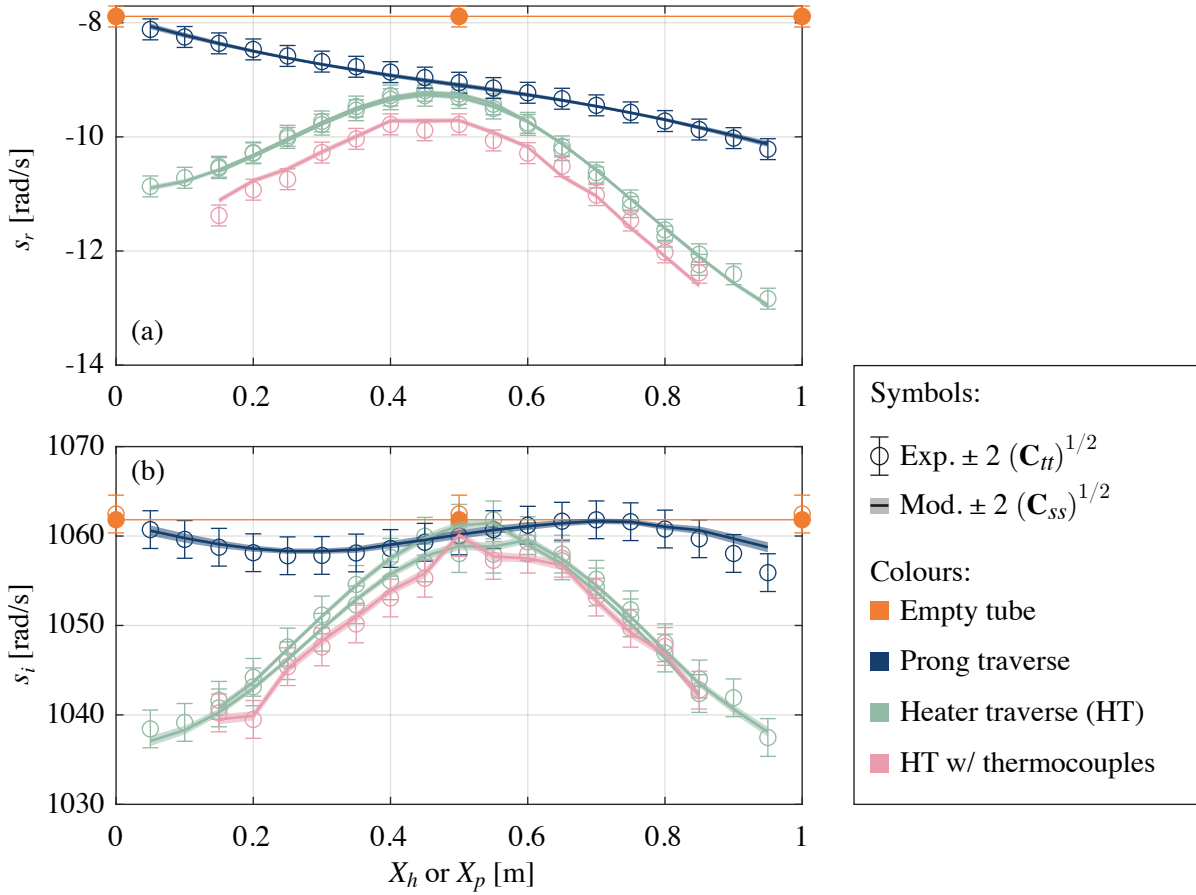


Fig. 3.8 Experimental measurements and model predictions of (a) the growth rate, s_r , and (b) the angular frequency, s_i , for all the cold experiments (rows C1, C2, C3, and H at $Q = 0$ in Table 3.1). The model combines R and η from Fig. 3.6, model 6 from Table 3.2 for the prongs, model 3 from Table 3.3 for the heater, and visco-thermal dissipation for the thermocouples. The experiments with the heater in place were performed on different days at different ambient temperatures, which accounts for the jagged lines that are particularly visible in the growth rate.

Table 3.3 As for Table 3.2 but assimilating models for the drag and blockage caused by the heater (Fig. 3.2). Column 3 describes whether k_{mu_h} is imaginary or complex: if imaginary, the heater is modelled as an orifice plate without viscous dissipation; if complex, the model also includes viscous dissipation. Column 4 describes whether k_{ep_h} is zero or complex: if complex, the model includes thermal dissipation.

Model	Params	k_{mu_h}	k_{ep_h}	$\log(\text{BFL})$	$\log(\text{OF})$	$\log(\text{ML})$
1	1	imag	none	-3.7472	-0.0949	-3.8420
2	2	complex	none	-1.6648	-0.1947	-1.8595
3	4	complex	complex	+0.4209	-0.4047	+0.0162

complex visco-thermal dissipation coefficients k_{mu_t} and k_{ep_t} . Fig. 3.8 shows the experimental measurements and calibrated model predictions from all the cold experiments, using calibrated

Levine-Schwinger for the reflection coefficients, Tijdeman's model for the tube boundary layer, model 6 from Table 3.2 for the prongs, model 3 from Table 3.3 for the heater, and assimilated visco-thermal dissipation coefficients for the thermocouples.

3.2.3 Inferring the fluctuating heat release rate at the heater

Having carefully developed a quantitatively-accurate model of the elements of the Rijke tube from the cold experiments, we now assimilate k_{eu_h} , which is the local feedback from the velocity to the heat release rate at the heater. Fig. 3.9(a) shows the expected values of $k_{eu_h} \pm 2$ standard deviations, calculated independently at each heater power and heater position. The parameter k_{eu_h} is inferred from the thermoacoustic behaviour of the system. The thermoacoustic effect is strongest when the heater is placed around $X_h = 0.25$ and $X_h = 0.75$, and is weakest when placed around $X_h = 0.5$. The uncertainty in k_{eu_h} should therefore be smallest when the heater is placed around $X_h = 0.25$ and $X_h = 0.75$ and greatest when placed around $X_h = 0.5$. This can be observed clearly in Fig. 3.9(a). Away from $X_h = 0.5$ and for heater powers above 15 Watts, the values of $|k_{eu_h}|/Q_h$ and $\angle k_{eu_h}$ are, with little uncertainty, re-assuringly independent of the heater position and heater power. The biggest deviation is seen when the heater is near the top end of the tube ($X_h \sim 0.8$), and, by inspecting the two right-most temperature profiles in Fig. 3.5, we suspect this is due to a systematic error in the base flow calculation.

For comparison, Fig. 3.9(b) shows k_{eu_h} calculated from the same data, but without including models for the visco-thermal dissipation and blockage of the heater and prongs. It is tempting to ignore the visco-thermal dissipation and blockage of the heater and prongs, as in [123, 114, 126, 127], because they are difficult to model and seem *a priori* to have little influence. With these ignored, however, $|k_{eu_h}|/Q_h$ and $\angle k_{eu_h}$ seem to change significantly with heater power and heater position. Most worryingly, $\angle k_{eu_h}$ seems to change discontinuously around $X_h = 0.5$. There is, of course, no physical justification for this apparent dependence. It arises simply because k_{eu_h} is being used to accommodate deficiencies elsewhere in the model. A machine learning algorithm applied to a neural network would learn this dependence but, consequently, would only be able to interpolate between previous observations. The advantage of our physics-based Bayesian approach is that it reveals bad models, such as that used to create Fig. 3.9(b), and forces the researcher to develop good models, which tend to have simple physically-interpretable behaviour, such as that used to create Fig. 3.9(a).

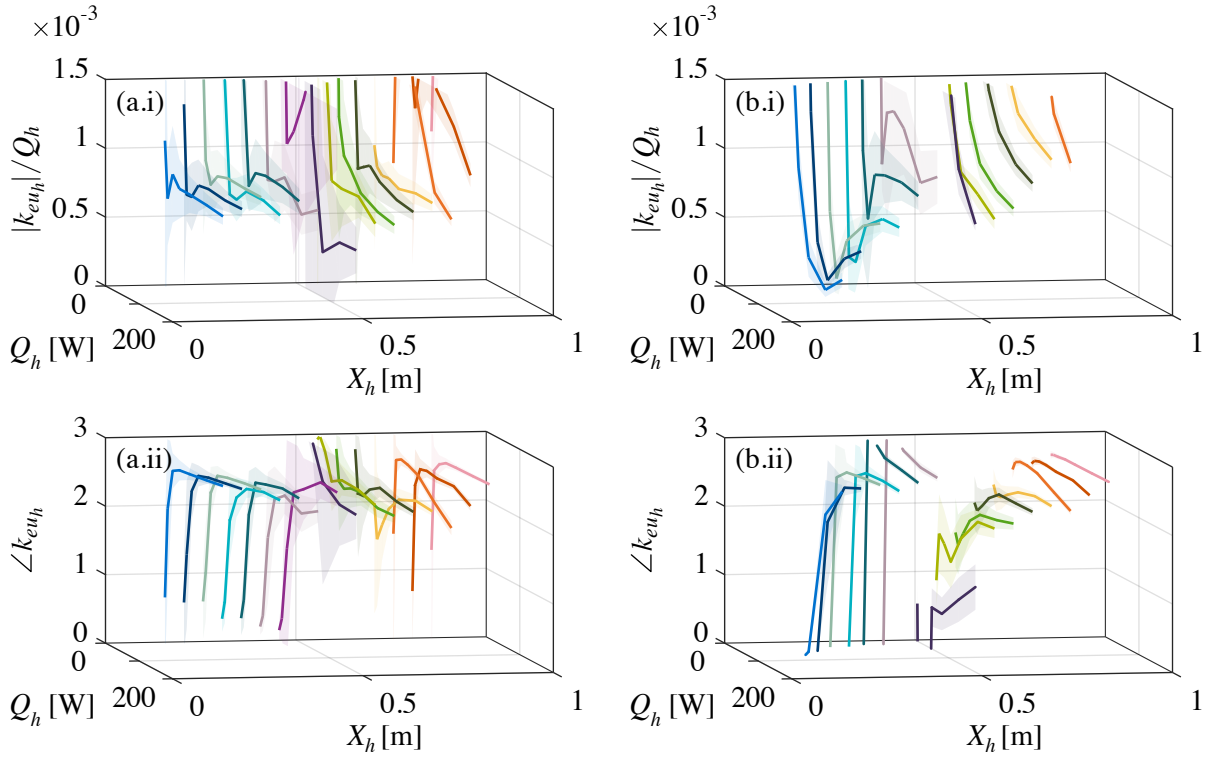


Fig. 3.9 Local feedback coefficient from velocity fluctuations to heat release rate fluctuations at the heater, k_{eu_h} , inferred independently at each heater power and heater position. The solid lines show the expected values. The patches show ± 2 standard deviations. (i) $|k_{eu_h}|/Q_h$; (b) $\angle k_{eu_h}$. These values were inferred by assimilating the data from the experiments in rows H of Table 3.1, into two models. (a) R and η from Fig. 3.6, model 6 from Table 3.2 for the prongs, model 3 from Table 3.3 for the heater, and visco-thermal dissipation for the thermocouples. (b) R and η from Fig. 3.6, but neglecting blockage visco-thermal dissipation from the prongs and heater.

Modelling the fluctuating heat release rate at the heater

We now propose nine candidate models for the complex local feedback coefficient k_{eu_h} . These models are listed in Table 3.4, where k_i with numerical subscripts are the model parameters. Models 1, 2, 7, 8, & 9 model $|k_{eu_h}|$ as $k_1 \times Q_h$, taking inspiration from Fig. 3.9, in which this quantity is seen to be nearly uniform over the operating regime. Model 3 models $|k_{eu_h}|$ as $k_1 \times Q_h^{k_3}$, to see whether the added flexibility of the power law in Q_h increases the marginal likelihood of the model. Models 4, 5, & 6 model $|k_{eu_h}|$ as $k_1 \times Q_{\text{King}}$, where Q_{King} is calculated from King's law [133, Eq(33)]:

$$Q_{\text{King}} = \frac{1}{(\text{Re Pr } \pi/2)^{-0.5} + 2} \frac{Q_h}{\bar{u}_h} \quad (3.6)$$

Models 1, 3, 4, & 9 model $\angle k_{eu_h}$ as constant k_2 , taking inspiration from Fig. 3.9, in which this is nearly uniform over most of the operating regime. Models 2, 5, 7, & 8 model $\angle k_{eu_h}$ as $k_2 \times \text{Im}(s)$, which corresponds to a constant time delay model in which $\tau = k_2$. Model 6 models $\angle k_{eu_h}$ as $k_2 \times s_i \tau_L$ where $\tau_L \equiv 0.2D_{\text{wire}}/\bar{u}_h$ is the time delay calculated by Lighthill [115]. If this is accurate then k_2 should assimilate to 1.

Models 1 to 6 assume that the visco-thermal dissipation at the heater does not change when the heater element becomes hot. Model 7 allows the viscous dissipation to drift from its cold value in proportion to Q_h . Models 8 and 9 allow the viscous and thermal dissipation to drift from their cold values in proportion to Q_h .

Table 3.4 As for Table 3.2 but assimilating models for the hot heater. Column 2 contains the number, N , of parameters, which are labelled $k_1 \dots k_N$. Column 3 contains the model for $|k_{eu_h}|$, where Q_h is the heater power and Q_{King} comes from Eq. (3.6). Column 4 contains the model for $\angle k_{eu_h}$ where s_i is the frequency and $\tau_L = 0.2D_{\text{wire}}/\bar{u}_h$ is Lighthill's time delay. Columns 5 and 6 contain the viscous and thermal dissipation coefficients where models 7-9 are allowed to deviate from their cold values, k_c .

Mod	N	$ k_{eu_h} $	$\angle k_{eu_h}$	k_{mu_h}	k_{ep_h}	log(BFL)	log(OF)	log(ML)
1	2	$k_1 \times Q_h$	k_2	k_c	k_c	-4.7055	-0.1016	-4.8071
2	2	$k_1 \times Q_h$	$k_2 \times s_i$	k_c	k_c	-4.5805	-0.1030	-4.6836
3	3	$k_1 \times Q_h^{k_3}$	k_2	k_c	k_c	-4.6477	-0.1385	-4.7862
4	2	$k_1 \times Q_{\text{King}}$	k_2	k_c	k_c	-4.7218	-0.1058	-4.8275
5	2	$k_1 \times Q_{\text{King}}$	$k_2 \times s_i$	k_c	k_c	-4.6956	-0.1072	-4.8028
6	2	$k_1 \times Q_{\text{King}}$	$k_2 \times s_i \tau_L$	k_c	k_c	-5.7721	-0.1019	-5.8741
7	4	$k_1 \times Q_h$	$k_2 \times s_i$	$k_c + Q_h(k_3 + ik_4)$	k_c	-3.4889	-0.2655	-3.7544
8	6	$k_1 \times Q_h$	$k_2 \times s_i$	$k_c + Q_h(k_3 + ik_4)$	$k_c + Q_h(k_5 + ik_6)$	-3.2512	-0.3964	-3.6476
9	6	$k_1 \times Q_h$	k_2	$k_c + Q_h(k_3 + ik_4)$	$k_c + Q_h(k_5 + ik_6)$	-3.6496	-0.3837	-4.0333

Of the models with constant visco-thermal dissipation at the hot wire (models 1 to 6), model 2 has the highest marginal likelihood. The extra flexibility of model 3 turns out not to be beneficial. Modelling the heat release rate with King's law is marginally less accurate, but this could be due to systematic errors in \bar{u}_h calculated with the long timescale code in § 3.1.4. The marginal likelihood increases significantly when the visco-thermal dissipation at the hot wire is allowed to increase with the heater power in models 7, 8, & 9. This has little effect on the frequency but a significant effect on the growth rate when the heater is placed near the ends of the tube, where the acoustic velocity is greatest. Comparing models 2, 7 & 8 we see that the extra viscous dissipation is more influential than the extra thermal dissipation, but that including both is best. Comparing models 1 & 2 and models 8 & 9 shows that there is more evidence for models with constant time delay τ than constant phase angle. Model 8 is the best model, but models 7 and 9 are also good.

Fig. 3.10(a) compares the experimental measurements with the model predictions of a model similar to those used by Refs. [123, 114, 126, 127]. This model has Tijdeman’s visco-thermal boundary layers, Levine-Schwinger’s reflection coefficients, and model 2 for the heat release rate, which was the best model when the visco-thermal dissipation from the heater is not included. The most likely coefficients of model 2 are assimilated from the data but the agreement is not good. This model is under-damped so, although it matches the frequency reasonably well, it always over-predicts the growth rate. Indeed at high powers and $X_h \approx 0.25$, this model predicts a positive growth rate. It is acceptable for the model to predict a positive growth rate in general, because the model is physics-based and can therefore extrapolate beyond the experimental regime. In this case, however, the model has not extrapolated beyond the experimental regime, so the positive growth rates are revealing model error. The structural error in the model is further revealed by the excessive total uncertainty that was required to maximize the posterior likelihood.

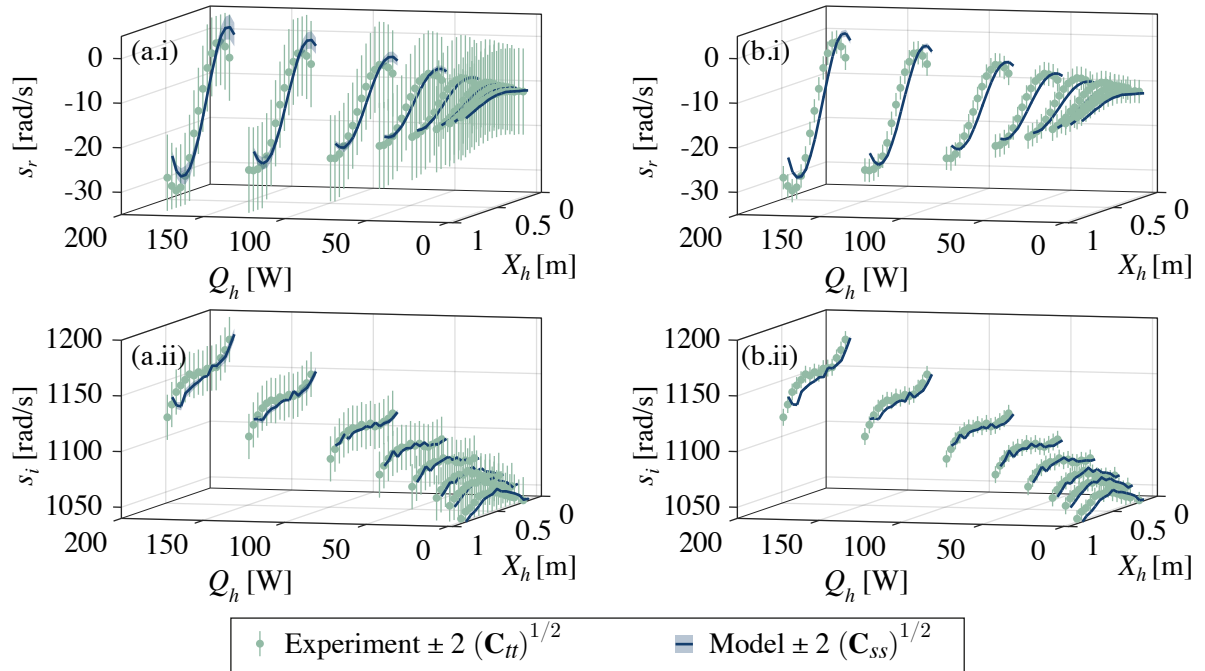


Fig. 3.10 Model predictions of (i) growth rate and (ii) frequency vs. experimental measurements assimilated with (a) model 2 from Table 3.4 for the heat release rate, Tijdeman’s model for the visco-thermal boundary layer, and Levine-Schwinger for the reflection coefficients; and (b) the same model but with the visco-thermal dissipation multiplied by η , where η is inferred from the hot data. The error bars on the experimental measurements show two standard deviations of the inferred total uncertainty, C_{tt} . The errorbars on the model show two standard deviations of the parametric uncertainty, C_{ss} .

Fig. 3.10(b) shows the model predictions when the researcher has noticed that the first model is under-damped and has chosen to model the extra damping as if it were all caused by the visco-thermal boundary layer. This model therefore assimilates the coefficient η as well as k_{eu_h} . The assimilated constant η is 1.67, which corresponds to a 67% increase in the boundary layer drag. This seems to be a reasonable approach and the model is able to approximate the growth rates and frequencies over the operating range. The inferred total uncertainty is smaller than in Fig. 3.10(a), but it is still larger than can be reasonably expected for these experiments. Additionally, there remains a significant discrepancy in the growth rate predictions when the heater is at the extremities of the tube and the heater power is high.

Fig. 3.11 shows the predictions for a model that contains the components with the highest marginal likelihoods: R and η from Fig. 3.6(b), model 6 from Table 3.2 for the prongs, model 3 from Table 3.3 for the heater, visco-thermal dissipation for the thermocouples, and model 8 from Table 3.4 for the fluctuating heat release rate. This model is qualitatively and quantitatively accurate over the entire operating range, as desired. The largest discrepancy occurs in the frequency prediction when the heater is placed in the downstream half of the tube, and set to the highest power. This is likely to be caused by a systematic error in the sound speed, which is difficult to measure accurately when the heater is in the downstream half of the tube, as described in § 3.2.1. With that said, the absolute value of the largest discrepancy is only 2.3 Hz, which is a 1.7% error.

3.2.4 Extrapolation

When constructing a model from a range of candidate models, it is beneficial to have a large amount of data. The densely sampled dataset presented in this section has allowed us to critically evaluate each candidate model over the full operating range of the rig, which has revealed structural errors in many commonly used models. Once a good model has been selected, however, it can be trained on sparse data. Fig. 3.12 compares model predictions with experimental measurements when data from just six arbitrarily selected operating points are assimilated into the model in Fig. 3.11. The results are almost indistinguishable from Fig. 3.11, for which 120 hot operating points were assimilated.

This highlights two major advantages of assimilating data into physics-based models over a physics-agnostic machine learning approach: (i) much less data is required, and (ii) the model can successfully extrapolate outside the training dataset, which is a common shortfall of physics-agnostic machine learning approaches.

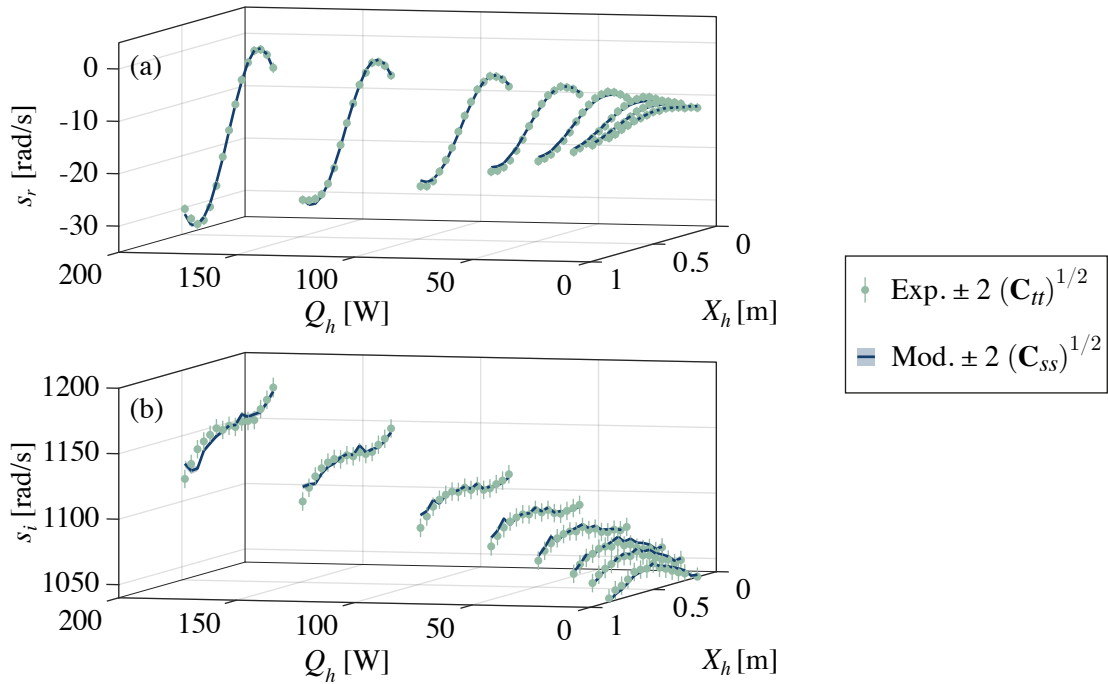


Fig. 3.11 As for Fig. 3.10 but using the component models with the highest marginal likelihoods: R and η from Fig. 3.6(b), model 6 from Table 3.2 for the prongs, model 3 from Table 3.3 for the heater, visco-thermal dissipation for the thermocouples, and model 8 from Table 3.4 for the fluctuating heat release rate.

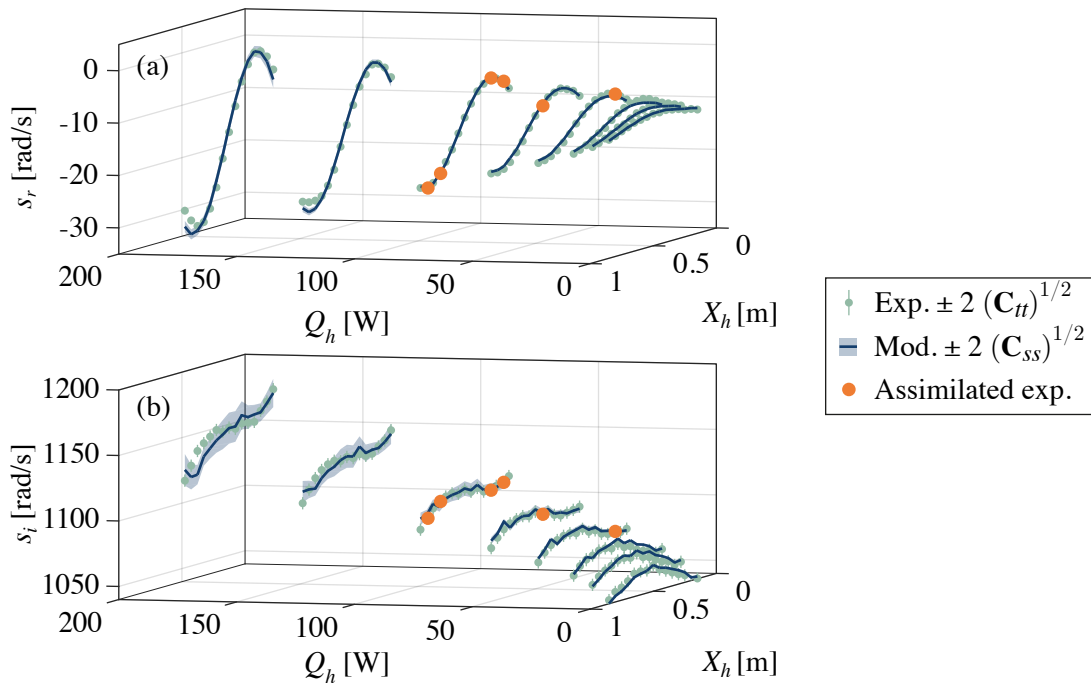


Fig. 3.12 As for Fig. 3.11 but assimilating data from just six of the 120 hot operating points.

3.2.5 Summary

We have assembled a quantitatively-accurate physics-based model of the Rijke tube, component by component. For each component we have proposed several candidate models and have then selected the component model with the highest marginal likelihood, given the data. The selected component model is always the simplest model that contains all the physics necessary to explain the data. The final assembled model is therefore as small as possible, quantitatively accurate, and physically interpretable. The model extrapolates successfully because it is physics-based, and little data is required to train it because it contains few parameters.

3.3 Minimizing the data required to train a model of a hot wire Rijke tube

At the end of the previous section we showed that it requires only a few experiments to infer the unknown parameters of a physics-based model. This was done, however, using arbitrarily chosen experiments. We now approach this more rigorously by applying Bayesian optimal experiment design to assess how much data is strictly necessary. We demonstrate the three applications of Bayesian experiment design described in § 2.2.1-2.2.3. These are (i) optimal design for parameter inference, (ii) optimal sensor placement for parameter inference, and (iii) optimal design for model comparison.

In each case, we demonstrate Bayesian optimal experiment design by recursively selecting the optimal experiment from the densely sampled set. On each recursion, we (i) identify the most informative experiment, (ii) assimilate the data from this experiment, and (iii) compare the model predictions to the data that has not yet been assimilated. The data that has not been assimilated acts as a validation set, which is hidden from the model and the assimilation process. This allows us to assess the performance of Bayesian optimal experiment design using real experimental data.

3.3.1 Optimal design for parameter inference

We first demonstrate the optimal experiment design framework on the simple sub-problem of inferring the visco-thermal dissipation of the cold heating element using as few experimental observations as possible. We then apply the method to the more complex problem of inferring the model parameters for the fluctuating heat release rate when the heater is switched on.

Assimilating heater dissipation from cold experiments

When traversing the cold heater through the tube, the vector of design parameters, \mathbf{x} , contains only the position of the heater within the tube, and the measurement vector, \mathbf{z} , contains the complex eigenvalue whose real part is the growth rate and imaginary part is the natural frequency of oscillations. We assimilate this data into a model with two complex parameters, \mathbf{a} , which are used to model the strength of the viscous and thermal dissipation of the heater.

In Fig. 3.13(a.i)-(a.ii), we compare the predictions of the prior model to the experimental data, which has not yet been assimilated. As expected, the prior model performs poorly, because

the prior parameter values are inaccurate. However, we have specified correspondingly large uncertainties in the prior parameters, so the model predictions have appropriately large error bars, which extend beyond the limits of the y-axes.

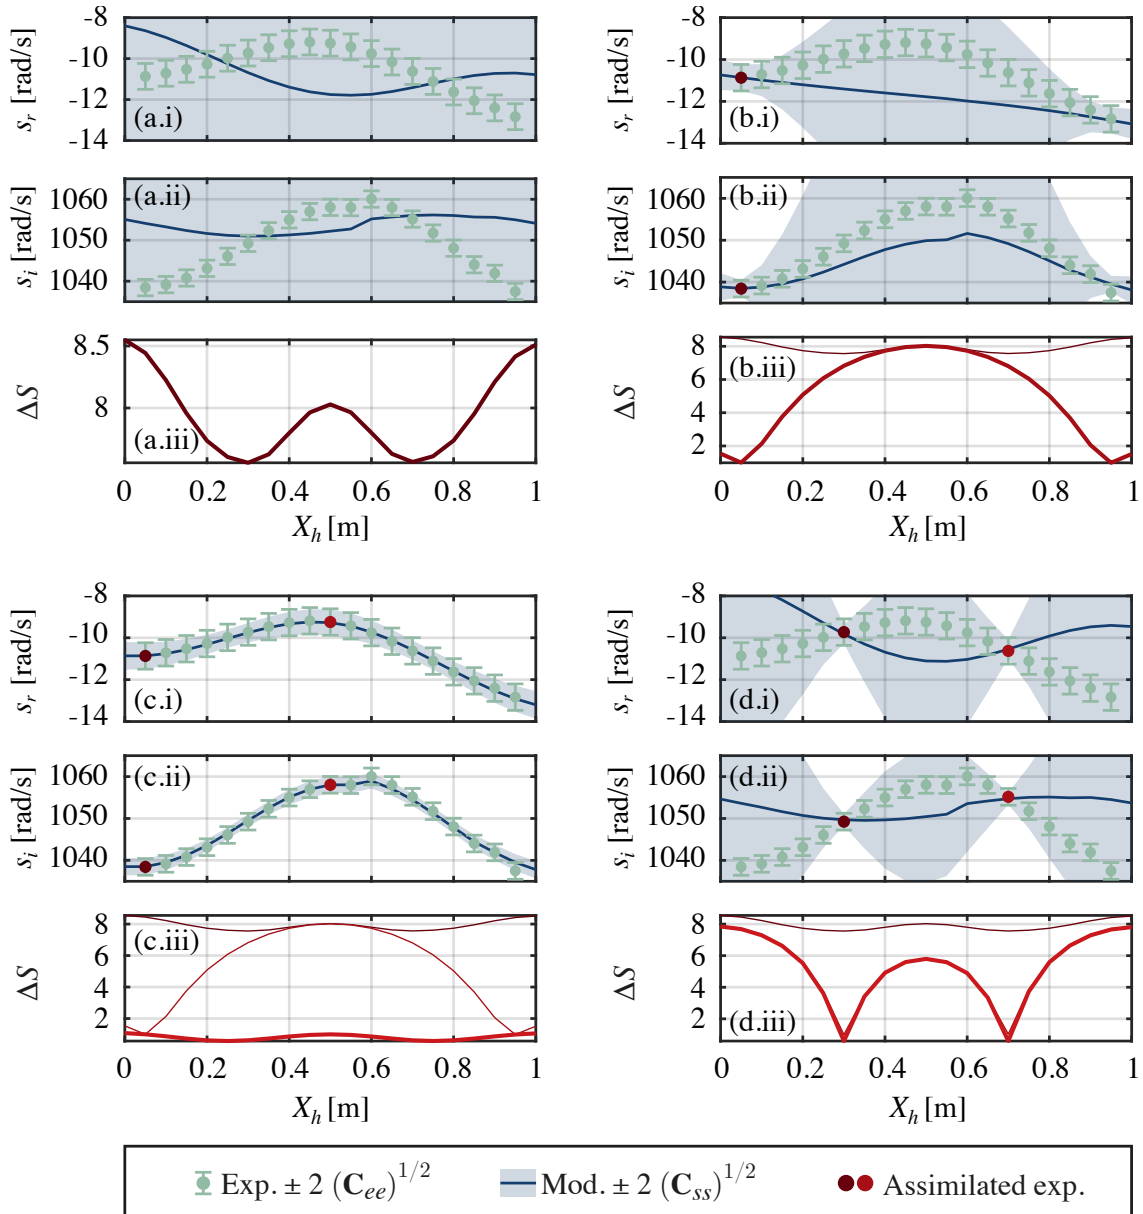


Fig. 3.13 Three steps of Bayesian optimal experiment design comparing experimental data to model predictions after assimilating: (a) no data, (b) the first datapoint with maximum information content (c) the second datapoint with maximum information content. For each step we show (i) growth rate, z_r , (ii) angular frequency, z_i , and (iii) information content, ΔS , plotted against heater position, X_h . For comparison, we also show (d) the result from assimilating the two experiments with minimum information content.

Fig. 3.13(a.iii) shows the potential information gain from an experiment performed at each heater position, as estimated using Eq. (2.16). As a reminder, this is calculated based solely on properties of the model, and knowledge of the data is not required in order to select the optimal experiment. We see that the expected change in Shannon entropy is maximal at the boundaries of the tube, indicating that the most information can be gained by performing an experiment there. This is an intuitive result because it places the heater near the point of maximum acoustic velocity fluctuations, and therefore maximum viscous damping. Furthermore, the acoustic pressure (equiv. temperature) fluctuations approach zero at the boundaries, so we expect the thermal dissipation to approach zero there. With this experiment, we can therefore easily infer the parameters controlling viscous dissipation without being influenced by effects of thermal dissipation.

Due to symmetry, the information gain at $X_h = 0.05$ and $X_h = 0.95$ is equal, so either experiment is a viable candidate experiment. We arbitrarily choose to assimilate the measurement taken at $X_h = 0.05$, which updates the parameter probability distribution. The new model predictions are plotted in Fig. 3.13(b.i)-(b.ii). We see that near the assimilated data point, the model is now accurate and the uncertainty has collapsed to the uncertainty in the data. Away from the assimilated data point, the model predictions deviate from the data and the uncertainty increases. We see, however, that because the symmetry of the problem is encoded in our physics-based model, we have also learned about the system response when the heater is placed at the opposite end of the tube.

Looking at the expected information gain in Fig. 3.13(b.iii), we see that we do not expect to learn much more by conducting another experiment near the ends of the tube. However, there is still information to be gained from an experiment conducted at the centre of the tube. Once again, this is an intuitive result because it now places the heater in the acoustic pressure (equiv. temperature) anti-node and velocity node, making thermal dissipation maximum and viscous dissipation zero. With this experiment, we easily infer the parameters controlling thermal dissipation, in the absence of viscous dissipation.

We assimilate the measurement taken at $X_h = 0.5$ and plot the new model predictions in Fig. 3.13(c.i)-(c.ii). This shows that, after assimilating just these two data points, the model matches the experimental data for all heater positions, and the model uncertainty is now similar in magnitude to the measurement uncertainty at all heater positions.

The expected information gain, shown as the bold red line in Fig. 3.13(c.iii), is now relatively flat and close to zero at all heater position. This shows that no experiment will

be substantially more informative than any other, and there is little information to be gained by assimilating more data. We have therefore gained all meaningful information about the unknown parameters using just two carefully chosen experimental observations. Although a skilled experimentalist may be able to identify the most informative experiments from physical insight alone, this becomes increasingly difficult when assimilating data into more complex models containing more physical mechanisms and more model parameters.

An expected value of two parameters can, of course, be obtained from any two independent observations. Their posterior variances, however, depend strongly on which two observations are selected. To demonstrate this, we assimilate the two data points with the lowest information content, as shown in Fig. 3.13(d.i)–(d.iii). The results clearly show that assimilating these two points only achieves local accuracy and local certainty in the model predictions.

As a final demonstration of the effectiveness of this approach, we sequentially assimilate all the available data and record the information gained as each experiment is assimilated. Fig. 3.14(a) shows the Shannon entropy plotted against the number of experiments assimilated, while Fig. 3.14(b) shows the change in Shannon entropy against the number of experiments assimilated. For comparison, we also plot the information gained by applying an experiment design strategy in which the worst available experiment is assimilated at each step.

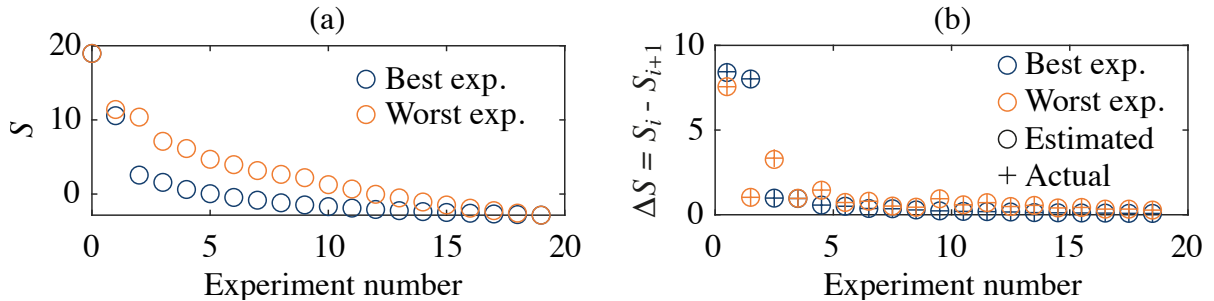


Fig. 3.14 Comparison of learning rate when selecting the best vs worst experiments. Plot (a) shows how the Shannon entropy of the parameter probability distribution decreases as additional experiments are assimilated. Assimilating the best experiment at each step (blue) leads to a rapid decrease in Shannon entropy, until the model uncertainty falls below the experimental uncertainty. Assimilating the worst experiment at each step (orange) leads to a gradual decrease in Shannon entropy throughout the data set. Plot (b) shows the information gained from each experiment, which is given by the change in Shannon entropy. We show the information gain estimated before the data is assimilated using Eq. (2.16) (+), as well as the actual achieved information gain, calculated after the experiment is assimilated (o).

We see from Fig. 3.14(a) that assimilating the best available experiment at each step causes the Shannon entropy to reduce rapidly with the first two observations, i.e. we quickly become confident in the most probable value of the unknown parameters. We then reach a point where

the model uncertainty reduces below the experimental uncertainty, and further experiments are relatively uninformative. By contrast, assimilating the worst experiments results in a gradual reduction in uncertainty.

This is reinforced by Fig. 3.14(b), which shows that when assimilating the best experiments, we gain significant information from the first two observations, and negligible information thereafter. By comparison, when we assimilate the worst experiments, we gain significant information from the first observation, and then a small but meaningful amount of information throughout the data set. We see that the information gained by assimilating the worst experiments tends to oscillate between a small information gain, followed by a negligible information gain. This is because once an experiment is assimilated, its symmetric pair becomes the next least informative experiment, which carries almost no new information.

In Fig. 3.14(b) we also compare the information gain estimated using Eq. (2.16) against the actual information gain, which is calculated from the reduction in Shannon entropy when the experiment is assimilated. We see that the values are almost indistinguishable, so in this case there was no harm in neglecting the data-dependant term in Eq. (2.3), because the second order parameter sensitivities are indeed small.

Assimilating a heat release rate model from hot experiments

We apply the same approach to experiments with the heater turned on, using the best model identified in § 3.2.2. In this case, the design parameter vector, \mathbf{x} , includes the heater position and heater power, while the measurement vector, \mathbf{z} , contains the complex eigenvalue. The model includes three complex parameters, \mathbf{a} , one of which models the heater transfer function and the remaining two of which model how the heater's visco-thermal dissipation changes with heater power.

In Fig. 3.15(a.i)-(a.ii), we compare the predictions of the prior model to the experimental data, which have not been assimilated yet. Once again we see that the prior model performs poorly, because the prior parameter values are inaccurate, but that the model predictions have appropriately large error bars.

Fig. 3.15(a.iii) shows the potential information gain for each experiment, which is estimated using Eq. (2.16). We see that, as expected, we gain more information by performing experiments at higher heater powers, where both the thermoacoustic effect and the variation in visco-thermal dissipation will be strongest. Furthermore, the information gain peaks at heater positions that produce the strongest thermoacoustic effect (where the magnitude of the product of

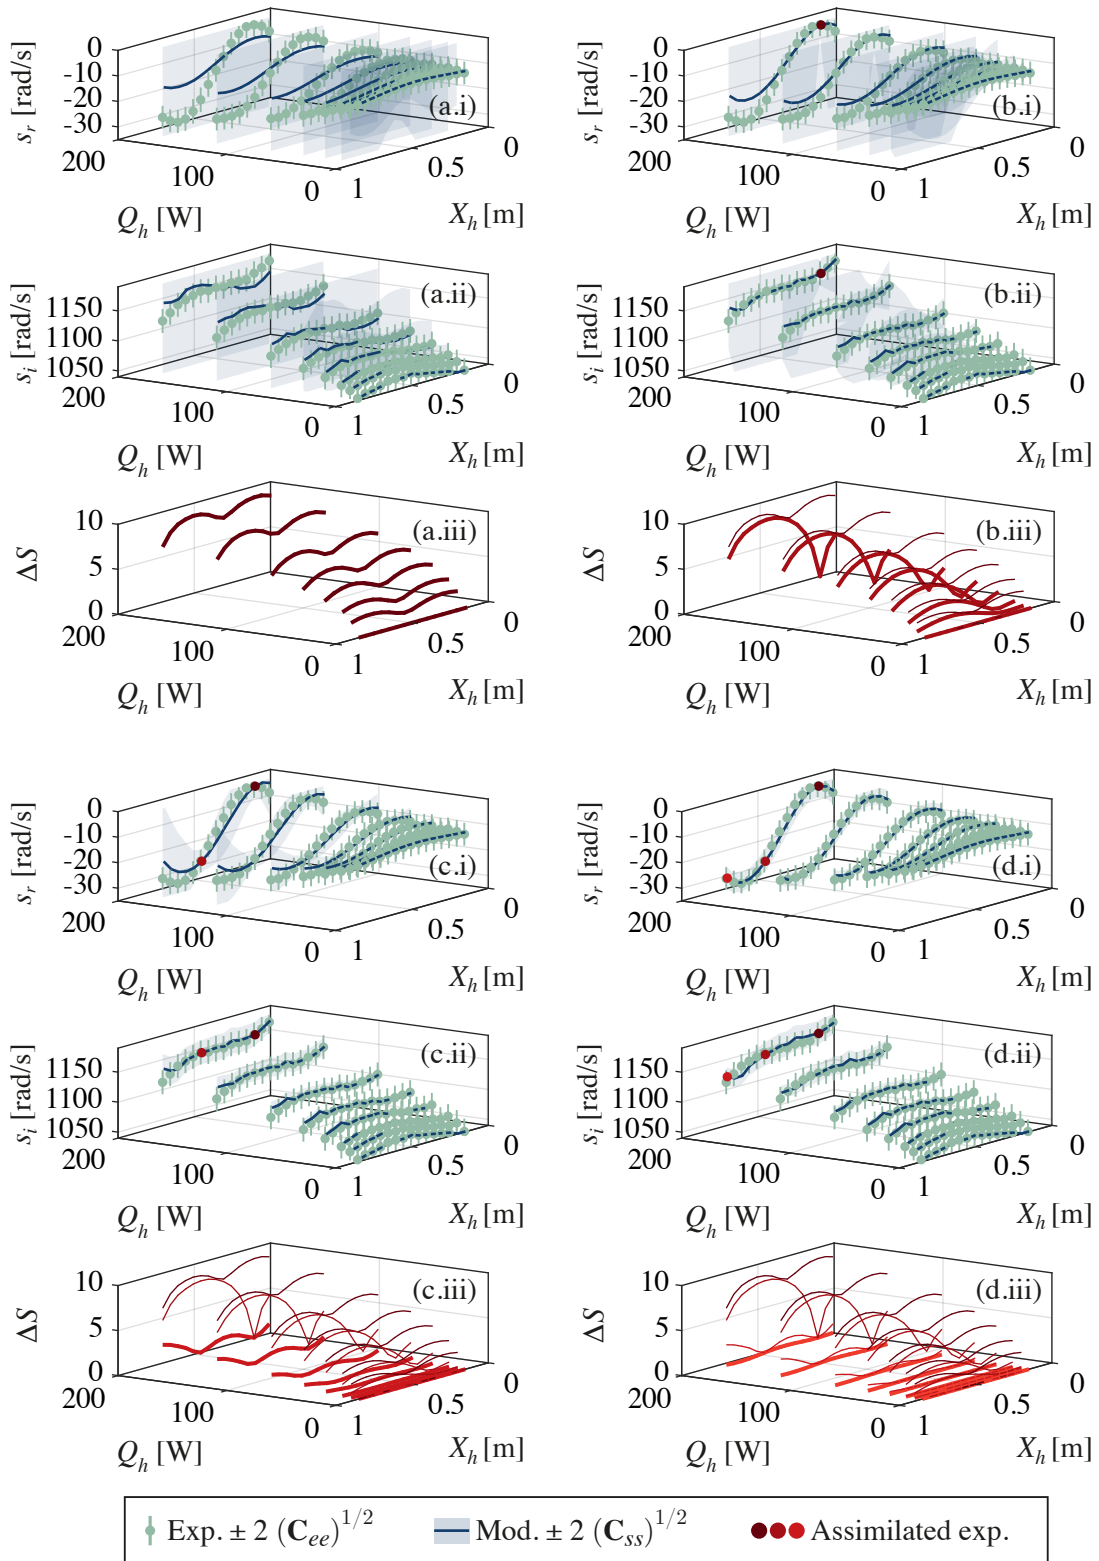


Fig. 3.15 Four steps of Bayesian optimal experiment design for assimilating data from experiments with the heater active. We compare experimental data to model predictions after assimilating: (a) no data, and (b)-(d) the first, second and third observation with maximum information content. For each step we show (i) growth rate, z_r , (ii) angular frequency, z_i , and (iii) information content, ΔS , plotted against heater position, X_h , and heater power, Q_h .

acoustic pressure and velocity is maximum). By contrast, placing the heater at the acoustic velocity node (near the centre of the tube) yields very little information. This is because both the thermoacoustic effect and the viscous dissipation are zero at the velocity node, so an experiment performed there only provides information about the variation of thermal dissipation with heater power.

We assimilate the best experiment and plot the results in Fig. 3.15(b.i)-(b.ii). We see, once again, that the model becomes locally accurate and certain, but that the accuracy and certainty reduce with distance from the assimilated experiment. Unlike in § 3.3.1, the symmetry of the underlying physics is broken by the steady and fluctuating heat release from the heater. As a result, making an observation with the heater placed at one end of the tube no longer provides information about the behaviour of the system when the heater is placed at the other end of the tube. However, we see that making an observation at the highest heater power has also provided information about the system's behaviour at lower heater powers. The model has become confident in its predictions about the system behaviour when the heater is in the same position, but operating at lower power outputs.

We see from Fig. 3.15(b.iii) that the first observation reduces the information we expect to gain from subsequent experiments performed at nearby heater positions, at all heater powers. An experiment performed near the downstream end of the tube now carries slightly less information than it did before the first observation was made. This implies that, while the physics is not fully symmetrical, some of the phenomena may be. We have therefore gained a small amount of information about the behaviour of the system on one end of the tube through an observation at the other end.

The model predictions after assimilating the second most informative experiment are shown in Fig. 3.15(c.i)-(c.ii). The second observation causes a more widespread collapse in uncertainty, although the model is still significantly more certain where data has been assimilated. The model uncertainty is not yet smaller than the experimental uncertainty for all candidate design parameters, so we expect we could still learn more through additional observations.

This is confirmed through Fig. 3.15(c.iii), which shows that the expected information gain has reduced for all heater positions and heater powers. However, there is still meaningful information to be gained from a further experiment performed at the downstream end of the tube (near $X_h = 1$), at the highest heater power.

Assimilating the third most informative experiment yields the model predictions shown in Fig. 3.15(d.i)-(d.ii). The model is now accurate at all pairs of design parameters, regardless of

whether the data was assimilated. Furthermore, the model uncertainty is below the experimental uncertainty everywhere.

We see from Fig. 3.15(d.iii) that the distribution of expected information gain is relatively flat, so there is no clear candidate for further experiments. Furthermore, the magnitude of expected information gain is near-zero everywhere, so there is little value in performing further experiments for the purpose of parameter inference. In the initial dataset we collected data at 120 hot operating points, but we now see that three carefully selected operating points are sufficient, which is just 2.5% of the full dataset.

While the chosen experiments are still somewhat intuitive, even a skilled experimentalist might struggle to select the three optimal experiments in this case. Even though the model remains relatively simple, the task of selecting optimal experiments for learning the values of multiple parameters with coupled effects becomes increasingly challenging as the number of parameters increases.

We now contrast the learning rate achieved by selecting the best and worst experiments at each step, shown in Fig. 3.16. The plot of Shannon entropy *vs.* number of experimental observations in Fig. 3.16(a) shows that selecting the best experiments leads to a rapid reduction of Shannon entropy, corresponding to a rapid learning of the parameter values. By comparison, selecting the worst experiments leads to a very gradual reduction in Shannon entropy, corresponding to gradual learning of the parameter values. Most strikingly, we see that selecting the worst experiments begins by using a set of experiments which carry no information at all. These correspond to the experiments in row H-0 of Tab. 3.1, which were performed with zero heater power. These experiments naturally provide no information about (i) the thermoacoustic effect, or (ii) how the heater visco-thermal dissipation changes with heater power.

This is reinforced in Fig. 3.16(b) which shows that the three best experiments provide significant information gain, following which we gain little information from the subsequent experiments. Selecting the worst experiments yields a small information gain throughout the data set, apart from the first few cold experiments, which provide zero information. Comparing the estimated and achieved information gains, we see that the first best observation provided more information than was predicted by Eq. (2.16). The under-prediction is due to the neglected data-dependent term in Eq. (2.3), because the second order sensitivities are small but not negligible in this case. After assimilating just one observation, however, the data-discrepancy reduces such that the data-dependent term becomes negligible for subsequent observations, and the estimated and actual information gains are almost identical.

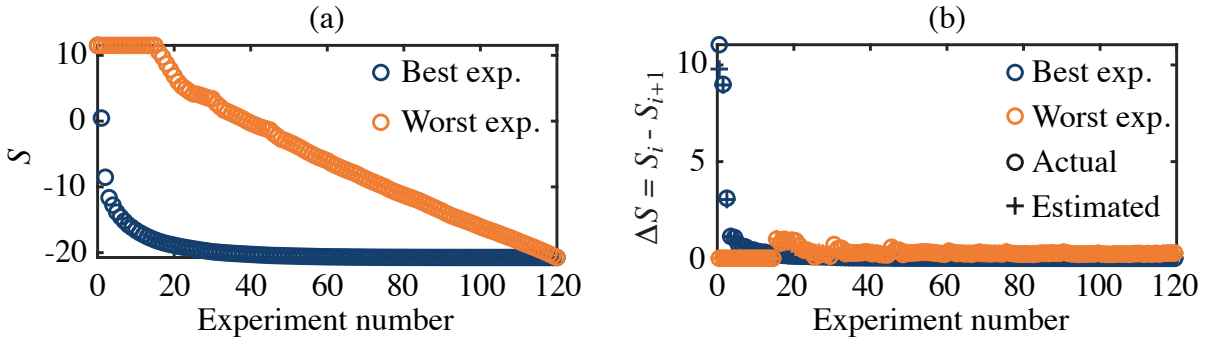


Fig. 3.16 Comparison of learning rate when sequentially assimilating the best *vs* worst experiments. Plot (a) shows how the Shannon entropy of the parameter probability distribution decreases as additional data are assimilated. Assimilating the best experiment at each step (blue) leads to a rapid decrease in Shannon entropy, until the model uncertainty falls below the experimental uncertainty. Assimilating the worst experiment at each step (orange) leads to a gradual decrease in Shannon entropy throughout the data set. Plot (b) shows the information gained from each experiment, which is quantified from the change in Shannon entropy before and after the data was assimilated. We show the information gain estimated before the data is assimilated, using Eq. (2.16) (+), as well as the actual achieved information gain, calculated after the data is assimilated (\circ).

3.3.2 Optimal sensor placement for parameter inference

As a demonstration of Bayesian optimal sensor placement, we will identify the optimal microphone configuration to determine the characteristics of the empty tube. The original dataset was collected with a redundant array containing eight microphones. This allows us to demonstrate optimal sensor placement by assimilating only data from the most informative microphones, using the remaining microphones for validation. In this case, the design parameter vector, \mathbf{x} , contains the sensor locations, and the measurement vector, \mathbf{z} , contains the eigenvalue and the complex pressure at the sensor locations. The model has two complex parameters corresponding to the upstream and downstream reflection coefficients, and a third real parameter corresponding to a correction factor to the boundary layer dissipation model.

As was seen in Fig. 3.6, it is not possible to infer unique values of the model parameters for the empty tube ($|R|$ and η) on this rig because the microphone array could not be phase-calibrated sufficiently accurately. In this section we will artificially phase-calibrate the microphones using the model output as ground truth. This effectively makes the phase data synthetic, but allows us to demonstrate optimal sensor placement without being hampered by the shortcomings of this specific rig. At the same time, we remove the assumption that the upstream and downstream reflection coefficients are identical, which was made in § 3.2. This presents a more challenging test of optimal sensor placement, and will become relevant in

Chapters 4 and 5, where we study a rig with different upstream and downstream end terminations. This means that the model for the empty tube has five unknown parameters: $|R_u|$, $\angle R_u$, $|R_d|$, $\angle R_d$, and η .

Fig. 3.17(a.i)-(a.ii) shows the model predictions of the pressure eigenmode, after the growth rate and natural frequency data has been assimilated. The model already makes reasonable predictions about the pressure eigenmode, even though no pressure measurements have been assimilated. However, the uncertainty in the pressure predictions is large because we have not yet gained enough information about the parameter values. The posterior joint distributions between each pair of parameters are shown in Fig. 3.18(a), from which we see that there is a strong correlation between the values of $|R_u|$ and $|R_d|$, and between $\angle R_u$ and $\angle R_d$. With five unknown parameters, observations of growth rate and natural frequency only provide enough information to learn the correlation between these two pairs of parameters, and the remaining parameter uncertainties remain unchanged from their prior values.

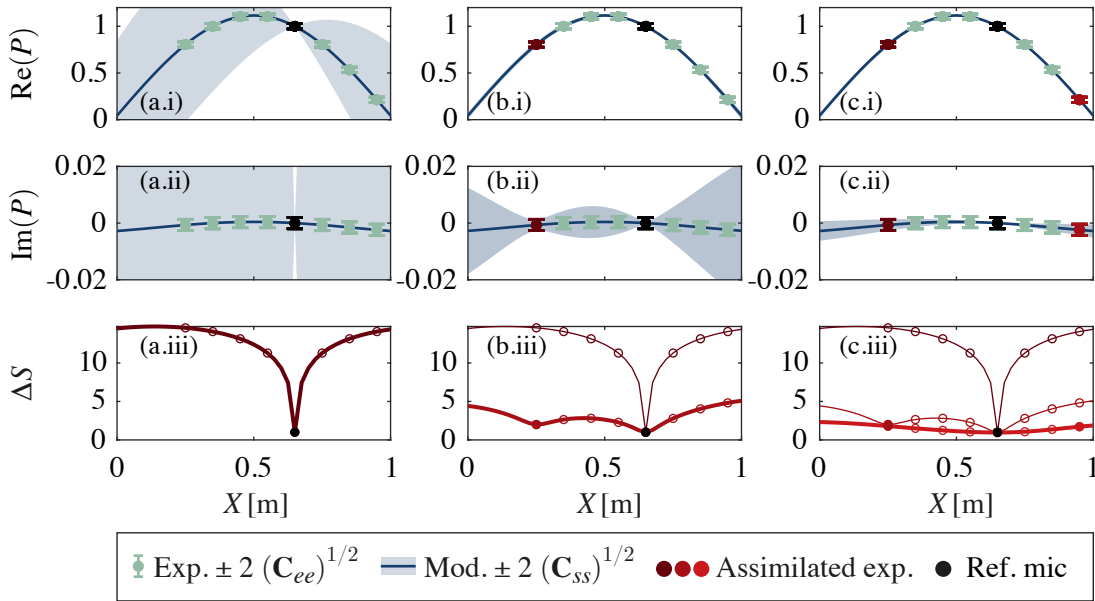


Fig. 3.17 Three stages of optimal sensor placement: (a) reference mic only, (b) one additional mic, and (c) two additional mics. We plot (i) the real component of the pressure, $\text{Re}(P)$, vs axial position in the tube, X , (ii) the imaginary component of pressure, $\text{Im}(P)$, and (iii) the expected information gain, ΔS , from a microphone placed at any axial position. Predictions are plotted as solid lines, with uncertainties indicated with shaded regions. Available microphone data is plotted in teal in (i)-(ii), and as open circles in (iii). Assimilated microphone data is coloured in shades of red from dark (fewest additional measurements) to light (most additional measurements).

Fig. 3.17(a.iii) shows the information that is expected to be gained from a microphone placed at any position along the length of the tube. The locations where microphone data are

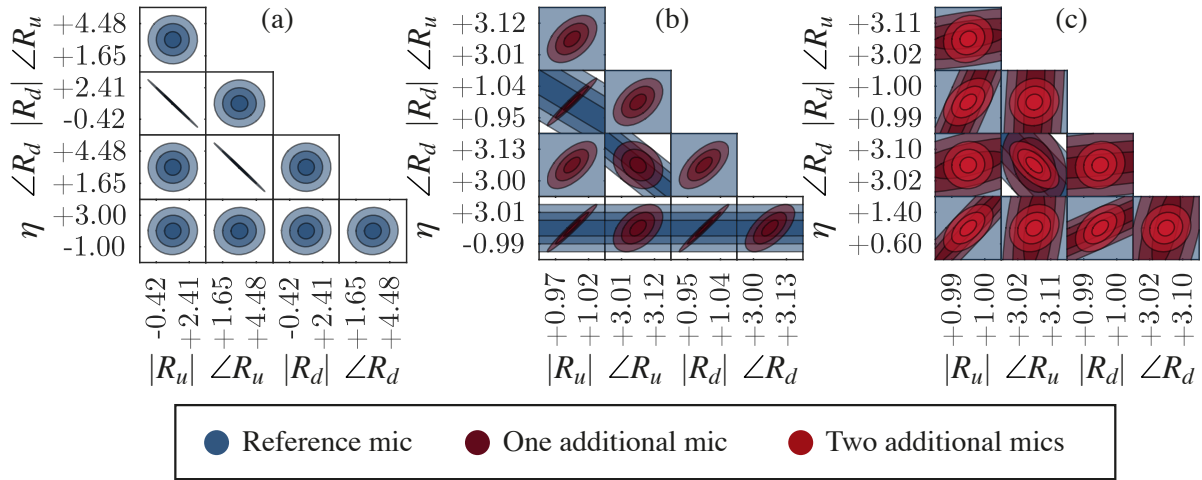


Fig. 3.18 Posterior joint probability distributions after three stages of optimal sensor placement: (a) reference mic only, (b) one additional mic, and (c) two additional mics. The joint distribution between pairs of parameters is indicated in each frame using contours of one, two and three standard deviations from the mean. The parameters are the absolute value and angle of the upstream and downstream reflection coefficients, R_u and R_d , and the boundary layer dissipation correction factor, η .

available are marked with circles. This reveals that there is no microphone installed in the optimal location. Fortunately, however, the best available microphone carries only slightly less information than a microphone placed at the optimal location.

We assimilate the data from the best available microphone and plot the new model predictions in Fig. 3.17(b.i)-(b.ii). We see from 3.17(b.i) that $\text{Re}(P)$ is now well-defined, with error bars that aren't distinguishable on the scale of the plot. The uncertainty in $\text{Im}(P)$, shown in Fig. 3.17(b.ii), has also reduced significantly. Adding a second microphone allows us to estimate the position of the pressure nodes, which identifies a unique pair of values for $\angle R_u$ and $\angle R_d$ and collapses the uncertainty in the prediction of $\text{Re}(P)$. This is reflected in Fig. 3.18(b), where we see that the correlation between $\angle R_u$ and $\angle R_d$ has been removed. However, there is still a three-way correlation between $|R_u|$, $|R_d|$ and η , seen in Fig. 3.18(b), which requires more information to further reduce the uncertainty. (Note that this is equivalent to the correlation that could not be removed in Fig. 3.6). We see from Fig. 3.17(b.iii) that we expect to gain further information from a sensor placed near the downstream end of the tube.

Fig. 3.17(c.i)-(c.ii) shows the updated predictions after assimilating the data from the third microphone. The uncertainty in both the real and imaginary parts of the pressure eigenmode predictions is now below the experimental uncertainty at all available microphone locations. We see from Fig. 3.18(c) that using the information from the three best microphones results in a posterior parameter distribution with predominantly uncorrelated parameter values. In other

words, we have inferred unique parameter values with small error bars using the minimum number of sensors. We therefore see that, out of the array of eight microphones installed in the rig, five were redundant. Furthermore, the installed microphone locations were not optimal, and more information could have been gained if the upstream microphone had been installed closer to the upstream boundary.

We once again demonstrate the effectiveness of this method by comparing the information gained by sequentially assimilating the best and worst microphone data. Fig. 3.19(a) shows the pressure predictions in (a.i)-(a.ii), and expected information gain in (a.iii) as the data from each microphone is assimilated in best-to-worst order. Fig. 3.19(b) shows the same process when the microphone data is assimilated in worst-to-best order. We see that assimilating the best microphone data first causes a rapid collapse in uncertainty as the first two additional microphones are added, following which the remaining microphones add negligible information. In comparison, adding the microphones in reverse order leads to a gradual reduction in uncertainty.

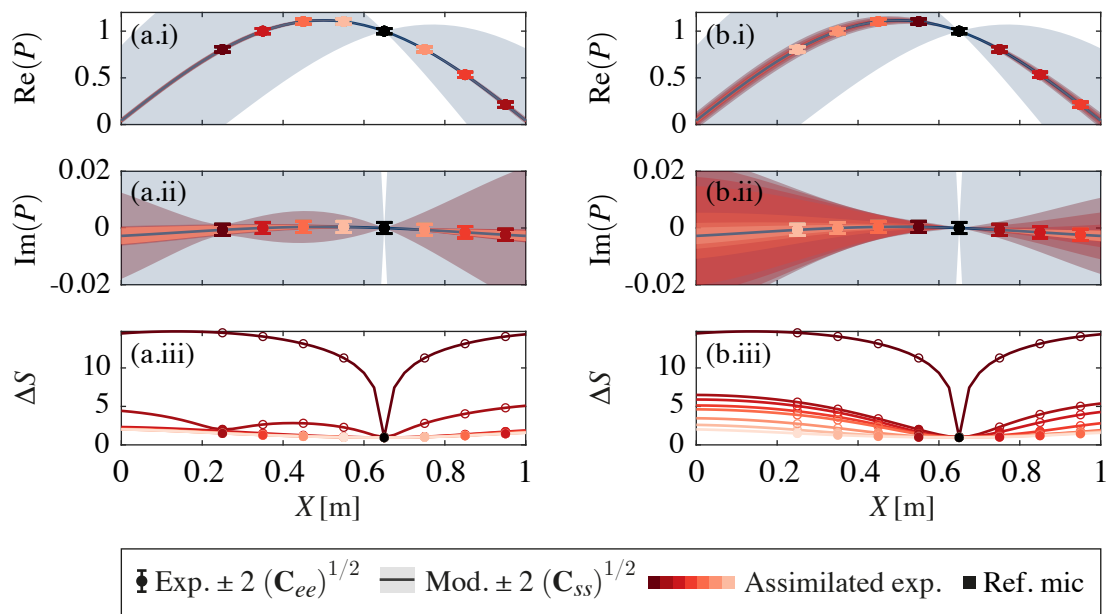


Fig. 3.19 Comparison between sequentially assimilating all available mic data. At each step, we select the (a) best and (b) worst microphones. Figures show (i) the real component of the pressure, $\text{Re}(P)$, vs axial position in the tube, X , (ii) the imaginary component of pressure, $\text{Im}(P)$, and (iii) the expected information gain, ΔS , from a microphone placed at any axial position. Predictions are plotted as solid lines, with uncertainties indicated with shaded regions. Uncertainties after assimilating only the reference mic are shaded in blue. Uncertainties after assimilating additional mics are coloured with shades of red from dark (fewest additional measurements) to light (most additional measurements).

This is reinforced by Fig. 3.20 which shows (a) the Shannon entropy and (b) the change in Shannon entropy as data from each subsequent microphone is assimilated. Fig. 3.20(a) shows that assimilating the microphone data in best-to-worst order reduces the Shannon entropy more rapidly than worst-to-best order. Fig. 3.20(b) confirms that the two best additional microphones provide significant amount of information, following which the subsequent microphones provide negligible information. In comparison, when assimilating the worst microphones first, the first additional microphone provides a significant amount of information, following which each subsequent microphone contributes a small but meaningful amount of information.

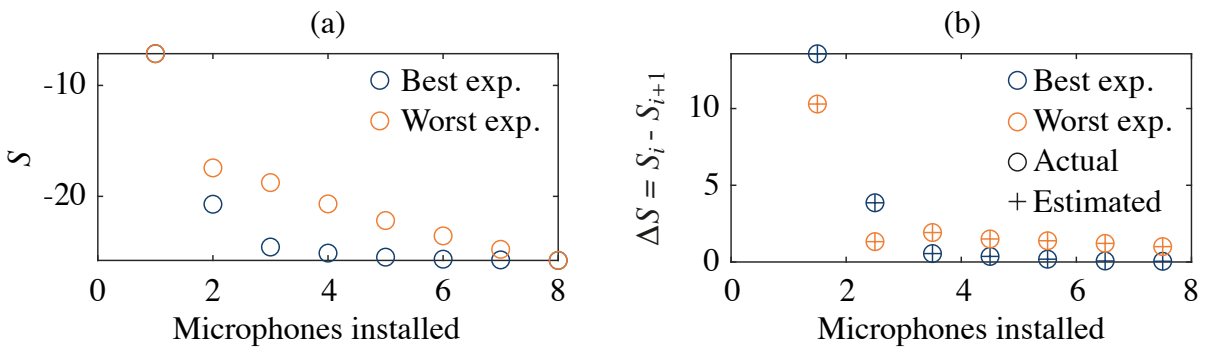


Fig. 3.20 Comparison of learning rate when sequentially selecting data from the best vs worst microphone positions. (a) Shows how the Shannon entropy of the parameter probability distribution decreases as additional mic data are assimilated. Assimilating the mic data from best to worst is shown in blue, and worst to best in orange. (b) Shows the information gained from each microphone, which is quantified from the change in Shannon entropy before and after the mic data was assimilated. We show the information gain estimated before the data is assimilated, using Eq. (2.16) (+), as well as the actual achieved information gain, calculated after the data is assimilated (o).

3.3.3 Optimal design for model comparison

We now consider the third task of identifying the experiment designs that allow us to most easily discriminate between multiple candidate models. For this demonstration we return to the problem of modelling the fluctuating heat release rate. We introduce two additional models for the purpose of comparison with the baseline model, which was trained in § 3.3.1. The first additional model, which we label model A, is a physics-based model with additional parameters that allow the heater feedback to vary with heater position. There is no physical basis for this, so we expect the model comparison to favour the baseline model. The second additional model, which we label model B, is a multivariate polynomial with the minimum degree capable of describing the data. This model disregards the physics entirely, so we once again expect the

baseline model to be favoured. The baseline model has 6 unknown parameters, model A has 14 unknown parameters and model B has 16 unknown parameters.

We begin by identifying and assimilating the optimal experiments to train each of the models using the process described in § 3.3.1. For fair comparison, we combine all the optimal experiments and train all models on the same set of experiments. We see in Fig. 3.21 that the predictions of both model A and model B are comparable to the baseline model's predictions. This makes it difficult to identify the best model by inspection. We would like to use Bayesian model comparison to identify the best model, so we sequentially identify and assimilate the experiments which maximize the discrimination between the candidate models. For clarity, we demonstrate this by separately comparing models A and B to the baseline model, but it can also be done for arbitrarily many models at once.

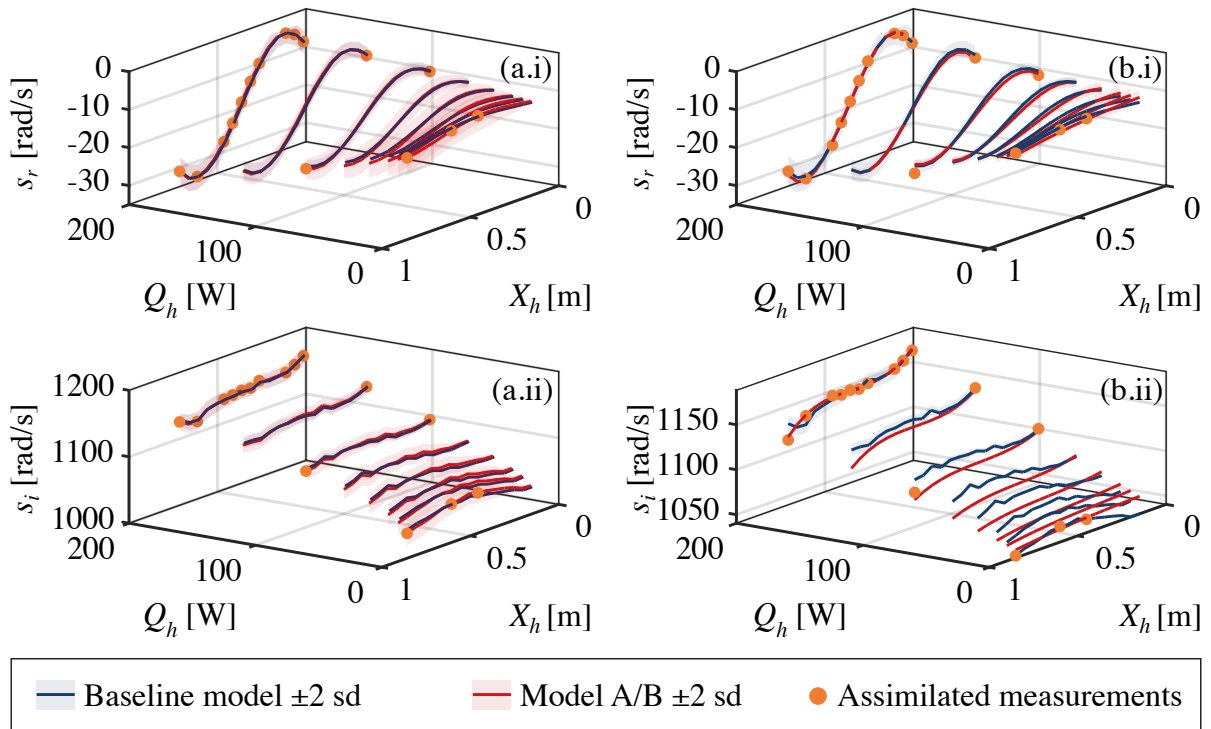


Fig. 3.21 Predictions of three candidate models for the fluctuating heat release rate. We compare (i) growth rate, z_r , and (ii) angular frequency, z_i , predictions produced by the baseline model (blue) against those produced by (a) model A (red) and (b) model B (red). The experiments selected to train the models are shown with orange markers.

We sequentially assimilate each experiment in the data set, selecting both the best and worst experiments for model comparison. The results are shown in Fig. 3.22(a) for comparing model A against the baseline model, and in Fig. 3.22(b) for comparing model B against the baseline

model. We see from Fig. 3.22(a.i) that model A is only slightly less probable than the baseline, while Fig. 3.22(b.i) shows that model B is substantially less probable than the baseline.

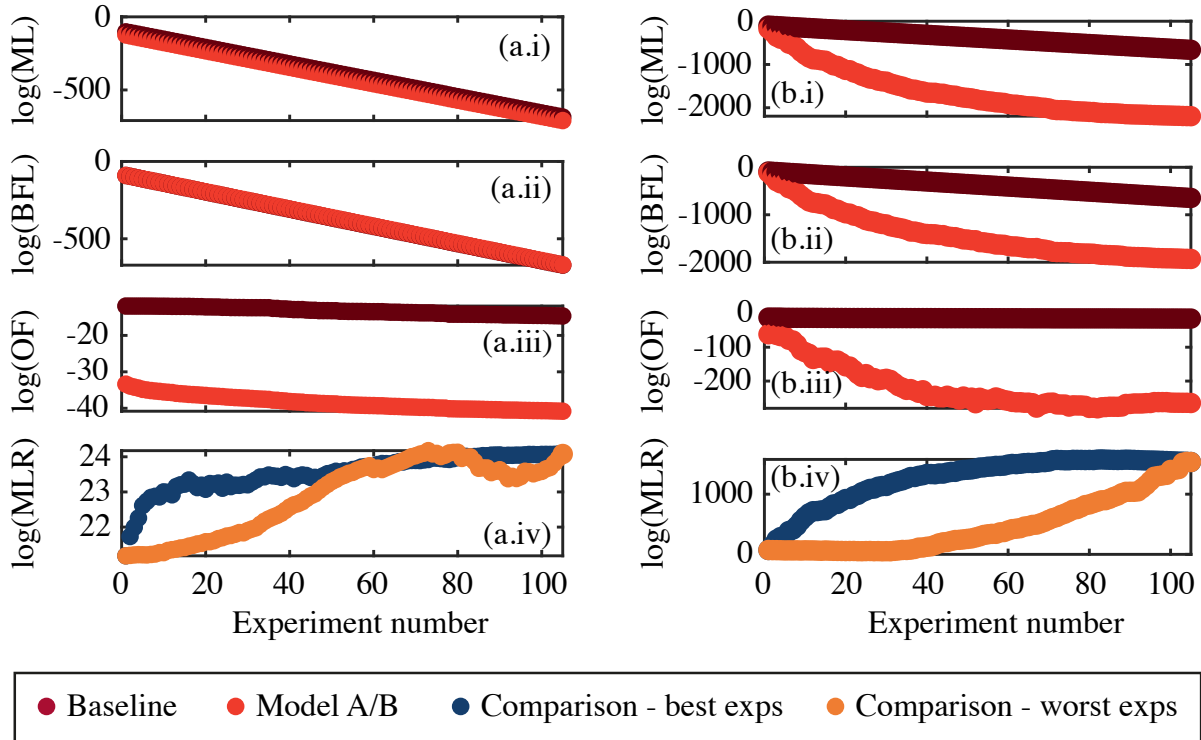


Fig. 3.22 Sequential assimilation of best and worst experiments for model comparison. Four model comparison metrics are plotted against the number of experiments assimilated. The metrics are (i) the log marginal likelihood, $\log(\text{ML})$, as well as its two components: (ii) log best fit likelihood, $\log(\text{BFL})$, and (iii) log Occam factor, $\log(\text{OF})$, and (iv) the log of the ratio of marginal likelihoods between two models, $\log(\text{MLR})$. In (a) we compare the baseline model (dark red) to model A (light red), and in (b) we compare the baseline model (dark red) to model B (light red). In (i)-(iii) we only show the results produced by selecting the best experiment at each step. In (iv) we show the results produced by selecting the best experiments (blue), as well as the worst experiments (orange). A positive $\log(\text{MLR})$ means that the baseline model is preferred.

Comparing the best fit likelihoods between model A and the baseline model, shown in Fig. 3.22(a.ii), we see that they are essentially identical. This tells us that both models fit the data equally well, which is evident in Fig. 3.21(a). This should be expected because model A is identical to the baseline model but has additional unnecessary parameters. The penalty of these additional parameters is seen by comparing the Occam factors in Fig. 3.22(a.iii), which shows that model A has a very small Occam factor (very negative $\log(\text{OF})$). The smaller Occam factor penalizes model A, resulting in a smaller marginal likelihood.

As mentioned in § 2.2.3, when choosing optimal experiments to discriminate between two models, we choose experiments where (i) the model predictions maximally disagree, and

(ii) the model uncertainties maximally disagree. Choice (i) results in a penalty in the best fit likelihood of one of the models, and (ii) results in a penalty in the Occam factor in one of the models. With the predictions of model A being almost identical to those of the baseline model, we cannot rely on the first criteria. Instead, each experiment is selected based on how much it would increase the Occam penalty on the more complex model. We see this in Fig. 3.22(a.iii), which shows that the $\log(\text{OF})$ of model A decreases rapidly at first, and gradually plateaus as the experiments with high discriminatory information are used up. As a result, we see from Fig. 3.22(a.iv) that, when selecting the best experiments, the marginal likelihood ratio increases rapidly at first and then plateaus. By comparison, selecting the worst experiments leads to a more gradual increase in model discrimination.

Comparing the best fit likelihoods between model B and the baseline model, shown in Fig. 3.22(b.ii), we see that model B is increasingly penalized as more data is added. This is because the polynomial model cannot describe the data as well as the baseline model, which is physics-based. Similarly, Fig. 3.22(b.iii) shows that the Occam factor penalty of model B becomes increasingly negative as more data is added, because model B is the more complex model. In this case the utility function must balance the choice of penalizing the model based on its fit to the data, or penalizing the model based on its complexity. Figures 3.22(b.ii-iii) show that both of these penalties increase rapidly at first and then plateau, indicating that the utility function is able to identify the experiments that maximize the discrimination from both penalties. As a result, we see from Fig. 3.22(b.iv) that, when selecting the best experiments, the marginal likelihood ratio increases rapidly at first and then plateaus. By comparison, selecting the worst experiments leads to a more gradual increase in model discrimination, with most of the discriminatory information added by the later experiments.

3.3.4 Note on application

In this chapter we demonstrate experiment designs that have strictly adhered to the optimal design approach. We note, however, that this approach tends to select experiment designs that are concentrated at the extremes of the parameter space. Experience has shown that training models on a narrow portion of the parameter space can lead to systematic errors going undiscovered. Therefore, a strict implementation of this approach may not be appropriate for models that have not been tested on densely sampled datasets to eliminate systematic or structural error in the data or model.

A sensible approach to make the process more robust is to use the ‘exploration vs. exploitation’ concept from Bayesian optimization [134, 135]. In this approach, an extra term is added to the utility function, which encourages the selection of experiments in regions of the parameter space that have not yet been explored. With this modification, the algorithm initially selects experiments with high information content, and then gradually switches to exploring the parameter space.

We also demonstrate optimal experiment design by selecting the single best experiment at each iteration. In practical situations, it may be more convenient to identify and conduct a batch of promising experiments, assimilate the data, and then identify the next batch. The current framework allows us to select arbitrarily many of the ‘best’ experiments at each iteration, so we are able to conduct experiments in batches. In the early stages of parameter inference, however, not all experiments in the batches will be optimal. This is because the second order terms that we neglected when defining the utility functions in Chapter 2 may be significant if the model discrepancy is large, which is likely during the early stages of training. However, as the model becomes more accurate, the information content within the batches will tend towards the optimum. For models that are linear in the parameters, the utility functions are exact, and the entire experimental campaign can be planned in one shot.

Finally, this chapter has not addressed the task of selecting a stopping criterion. That is a case-dependent problem and so a suitable criterion must be selected based on the problem at hand. In some cases it may be sensible to continue collecting data or placing sensors until a budget is exhausted, while in others it may be sensible to stop when the model uncertainty or expected information gain falls below a threshold.

3.4 Conclusion

In this chapter, we use the hot wire Rijke tube to demonstrate the full set of tools that the Bayesian framework provides. We first perform automated experiments to collect a dataset for training. With this dataset, we use Bayesian parameter inference, uncertainty quantification and model comparison to construct a quantitatively accurate model of the rig. We then use Bayesian optimal experiment design to assess how much data is strictly required for several inference tasks.

In § 3.2 we assemble a physics-based model of the hot wire Rijke tube, component by component. For each component we propose several candidate models and infer the unknown parameters of each candidate model. We rank the candidate models based on their marginal

likelihood and select the model with the highest marginal likelihood, given the data. The selected component model is always the simplest model that contains all the physics necessary to explain the data. The final assembled model is therefore as small as possible, quantitatively accurate, and physically interpretable. Further, the model extrapolates successfully because it is physics-based. This is a significant improvement on other attempts in the literature e.g. [123].

During this process, we critically evaluate several assumptions that are commonly applied when modelling the electrically heated Rijke tube, such as neglecting the visco-thermal damping of the boundary layer and neglecting the drag and blockage of the heater. We find that, when these assumptions are made, it becomes impossible to achieve quantitative agreement between the model and the data over the full operating range. The resulting models, therefore, are found to be significantly less likely than models that include these phenomena. This reveals one of the most powerful features of the Bayesian framework: the researcher is forced to carefully consider the underlying physics. At each step, the quality of the assumptions we make are rigorously evaluated with quantitative measures.

In § 3.3 we revisit the densely sampled dataset used to train the model in § 3.2, and use it to demonstrate Bayesian optimal experiment design. We consider three specific questions an experimentalist may face: (i) which experiments would provide the maximum information about the unknown parameters of a model, (ii) where should the sensors be placed to provide maximum information about the unknown parameters, and (iii) which experiments would maximize the discrimination between candidate models? We show that by using Bayesian optimal experiment design we can significantly reduce the experimental cost by (i) performing only a few, highly informative experiments, and (ii) using fewer, carefully placed sensors. This is an important result because, while our rig is cheap to operate, industrial rigs are expensive to operate. There is therefore a large financial incentive to minimize the number of experiments conducted on these rigs.

While other studies in the literature have attempted to apply Bayesian optimal experiment design to engineering problems, they have all found it to be computationally expensive within the frameworks they have chosen [97, 99, 96, 98, 136]. We have shown that, within the adjoint-accelerated Bayesian inference framework, we can obtain a good approximation of the optimal experiments for negligible computational cost, provided the posteriors are approximately Gaussian. We have demonstrated this on real experimental data, where other studies on the literature have solely relied on synthetic data, which is free of systematic error. While our framework is computationally cheap at runtime, it comes with the additional upfront cost of

creating the adjoint code. Adjoint codes are, however, widely available in the field of fluid dynamics.

Chapter 4

Ducted laminar conical flames

This chapter contains excerpts from a conference paper and a journal paper on generating a quantitatively-accurate model of the thermoacoustics of a ducted laminar conical flame using Bayesian parameter inference and model comparison [[137](#), [138](#)] ([\[138\]](#) is currently under review). This work was performed independently by Matthew Yoko. The experiments in this chapter were carried out on a rig that was designed and commissioned for the purpose of this study.

4.1 Introduction

A natural progression from the electrically heated Rijke tube is the ducted laminar conical flame. This retains the simple structure of the chamber, a vertical duct, but replaces the electric heater with a laminar conical flame. The thermoacoustic behaviour of these flames has been extensively studied both experimentally [139, 140, 37, 49, 141] and analytically [142, 41, 143–147].

The mechanism through which thermoacoustic oscillations occur in a ducted laminar conical flame is well understood. The acoustic perturbations give rise to convective velocity perturbations, which traverse the length of the flame at a velocity near the mean flow velocity. These convective perturbations distort the flame front, which in turn modulates the heat release rate. The heat release rate fluctuations then couple back to the acoustic field, giving rise to thermoacoustic oscillations. It is commonly assumed that the convective perturbations occur because the acoustic perturbations give rise to hydrodynamic instabilities [142, 41, 143, 147]. This is contested by one recent study [146], which shows that the convective velocity perturbation is the result of flame-flow interaction. Resolving this interaction requires the system to be modelled to extreme detail, which quickly becomes computationally expensive, and can result in the introduction of an unmanageable number of modelling parameters.

It is common to bypass the complexity of the flame dynamics by modelling the fluctuating heat release rate as a response to a velocity perturbation at some reference location near the base of the flame, either using a simple $n - \tau$ model or a flame transfer function:

$$\mathcal{F} = \frac{Q'/\bar{Q}}{u'/\bar{u}} \quad (4.1)$$

where \mathcal{F} is the (frequency dependent) flame transfer function, Q is the heat release rate, and u is the velocity at a reference location near the base of the flame. The fluctuating quantities are denoted as \bullet' and mean quantities are denoted as $\bar{\bullet}$.

The flame transfer function has been shown to change sensitively with changes in operating condition [47, 48], changes to the confinement of the flame [49, 50] and when the flame is combined with other flames [37, 43]. It is therefore beneficial to determine the flame transfer function with the flame *in-situ*.

Flame transfer functions are typically obtained from direct experimental measurements. These measurements require (i) a means of measuring acoustic velocity at the reference location near the base of the flame, and (ii) a means of measuring the heat release rate fluctuations.

The velocity is typically measured using a hot wire anemometer [47, 140, 148] or via optical methods [41, 149, 150]. The heat release rate fluctuations are typically measured by optical methods [41, 47, 140, 148–150]. None of these measurement techniques are suitable for *in-situ* measurements in a practical combustor, because they either rely on delicate instruments being mounted in a harsh operating environment, or on optical access which is typically limited or unavailable in a practical combustor. In some cases, the acoustic velocity is inferred from multiple microphone measurements distributed along the length of the rig [42, 43]. While this is technically possible to implement on practical combustors, it may not always be possible to infer the acoustic velocity upstream of the flame using pressure signals from microphones mounted on the casing wall.

By contrast, the growth rate and natural frequency of thermoacoustic oscillations can be obtained from even a single microphone mounted far from the flame, either through harmonic forcing [148], or by using data-driven methods to extract the linear growth rate and natural frequency from combustion noise or limit cycle data [151, 152]. It is therefore desirable to be able to infer the fluctuating heat release rate from these observations alone.

One method for indirectly obtaining flame transfer functions from pressure measurements has been demonstrated on a laboratory burner [153], and more recently on the Rolls-Royce SCARLET thermoacoustic test rig [154, 155]. This approach uses the two-source method [156], which requires specially designed experimental rigs with two sets of loudspeakers or sirens to force the rig from either end, and multiple microphones distributed along the length of the rig. Acoustic pressure measurements are collected from four experiments. In the first two experiments, the cold rig is forced harmonically at various frequencies from the upstream end, followed by the downstream end. In the next two experiments the flame is ignited, and the rig is again forced at the same frequencies from the upstream end, followed by the downstream end. The resulting data is then post-processed to extract (i) the characteristics of the cold rig, and (ii) the flame transfer function. This method provides an estimate for the flame transfer function, but does not quantify the uncertainty in this estimate.

A few recent studies have applied data-driven methods to infer the parameters of a fluctuating heat release rate model from pressure time series data [157–159]. In [157, 158], a non-probabilistic approach is used to infer the parameters. The authors use an optimization algorithm to minimize the discrepancy between the model and data, although they do not consider the uncertainties, or the resulting uncertainties in the inferred parameters. In [159], a frequentist approach is used to infer the fluctuating heat release rate from pressure time

series data. In the frequentist framework, the authors are able to quantify the uncertainty in the inferred parameters. They cannot, however, exploit prior knowledge or evaluate the marginal likelihood in order to compare candidate models. Gant et al. [159] demonstrate their method on synthetic data generated by their model. While this is a powerful tool for evaluating and demonstrating an inference framework, it does not allow the researcher to evaluate how the method handles a systematic mismatch between the model and the data, which is always present when assimilating experimental data into a model.

In this chapter, we apply our Bayesian inference framework to infer the flame transfer functions of a series of conical flames from pressure observations. We use Bayesian model comparison to choose the best model for the fluctuating heat release rate from a set of candidate models. We then infer the most probable flame transfer function for each flame, and rigorously quantify the uncertainties in each of the flame transfer functions.

4.1.1 Experimental configuration

The experimental rig comprises a premixed laminar conical flame inserted into a vertical duct, illustrated in Fig. 4.1. The lower end of the duct is fixed to a plenum chamber, through which co-flow air is supplied, and the upper end is open to the atmosphere. A photograph of the rig is provided in Fig. B.3 in Appendix B.

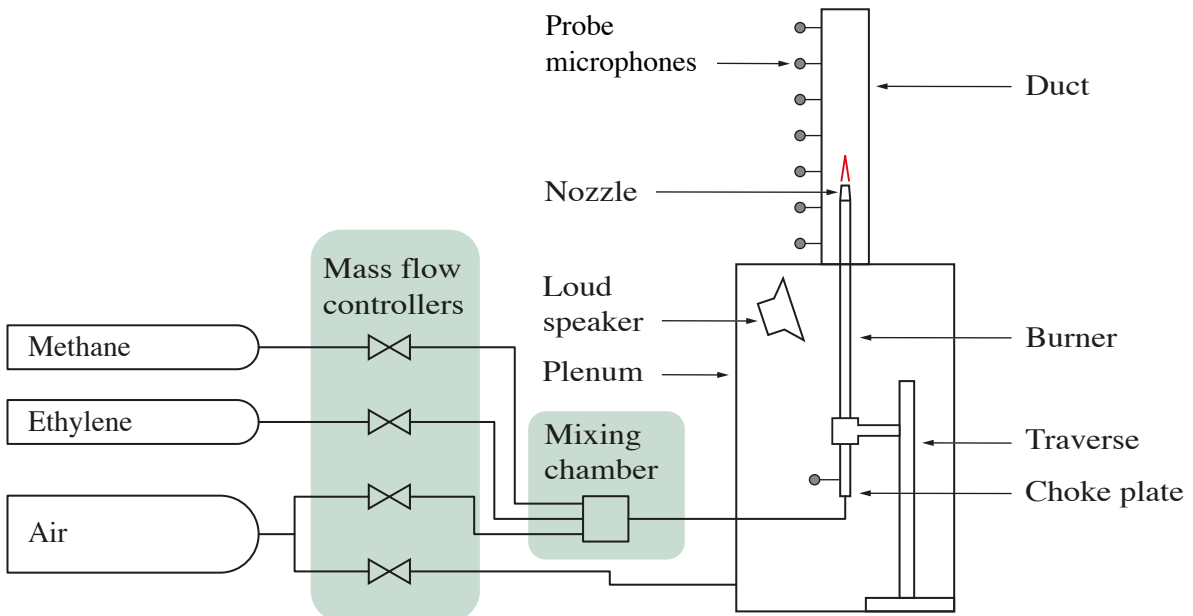


Fig. 4.1 Diagram of the ducted flame rig.

The duct is a 0.8 m long section of quartz tube with an internal diameter of 75 mm. The duct joins the plenum via a machined flange. The flange provides an airtight seal and an acoustic termination without any internal steps. There are 24 ports along the length of the duct, which allow for instrument access to the internal flow. The ports are fitted with airtight plugs when not in use.

The plenum is a fibreboard box with dimensions 1 m × 0.6 m × 0.6 m. The interior is lined with acoustic treatment to damp acoustic oscillations. Air is fed into the plenum via a mass flow controller to provide a constant flow of cool air through the duct. This keeps the duct and instrumentation at an acceptable temperature, and flushes the combustion products out of the rig.

The burner is a 0.85 m long section of brass tubing with an internal diameter of 14 mm. The outlet of the burner is fitted with a nozzle that is chosen such that the system can become thermoacoustically unstable. At the injection plane, the nozzle diameter is 9.35 mm.

The burner is fuelled by a mixture of methane and ethylene. This mixture allows a wide range of thermoacoustic behaviour to be explored by altering the shape of the flame through changes in the unstretched flame speed. The premixed air and fuel are metered using a bank of mass flow controllers (Bronkhorst EL-FLOW). From the mass flow controllers, the gases are fed into a mixing plenum to minimize the potential for equivalence ratio fluctuations. The mixture is then fed to the base of the burner, which is fitted with a choke-plate to decouple the supply lines from the acoustic fluctuations in the rig. Like the duct, the burner tube has eight ports for instrument access to the internal flow.

The burner is mounted to an electrically-driven traverse so that the vertical position of the burner inside the duct can be controlled. We are therefore able to explore changes in (i) flame position, (ii) flame shape (through changes in fuel composition) and (iii) mean heat release rate (through total fuel flow rate and fuel composition).

A loudspeaker (Visaton FRS 8) mounted within the plenum is used to force or stabilize the system, depending on whether the system is linearly stable or unstable. The loudspeaker is driven by a power amplifier (IMG STA-500). A phase-shift amplifier (custom design) is used to actively stabilize the system when it is unstable.

Instrumentation

The acoustic pressure is recorded using eight probe microphones (G.R.A.S. 40SA), with the probes inserted through the ports in the duct and burner walls. Seven microphones are

distributed along the duct with 100 mm spacing, starting at 100 mm from the upstream end. The eighth microphone is inserted through a port 100 mm downstream of the base of the burner tube. The microphones are each connected to an amplifier (G.R.A.S. 12AA), which is then connected to the acquisition system.

The rig is instrumented with 24 K-type thermocouples with a diameter of 0.4 mm. The diameter of the thermocouples was reduced from the hot wire Rijke tube to improve the accuracy of the measurements. Eight thermocouples are installed within the plenum to monitor the ambient conditions, another eight thermocouples are inserted through ports in the duct to monitor the internal gas temperature, and the final eight thermocouples are bonded to the duct's outer wall to monitor the duct temperature. The thermocouples are connected to the acquisition system through three 8-channel thermocouple amplifiers (Omega TC-08).

Automation & Acquisition

The experiment automation and acquisition is handled by a National Instruments controller (PXIe-8840) and acquisition card (PXIe-6358). The controller runs a LabVIEW code to automatically control the experiments. The code runs through a pre-determined test matrix of burner positions and air/fuel flow rates. For each test, the code (i) sets the burner position and air/fuel flow rates, (ii) waits for the system to reach thermal equilibrium, (iii) determines if the flame is linearly stable or self excited, and (iv) conducts 75 tests to determine the growth rate and natural frequency, the procedure of which differs depending on the system stability.

The microphone data is sampled at 50 kHz for the forcing tests, and 300 kHz for the sound speed tests. This is a significant improvement over the system used in the hot wire Rijke tube experiments, which was limited to a maximum sampling rate of 62.5 kHz. By sampling at 300 kHz for the sound speed tests, we can resolve a narrower acoustic impulse, which allows us to more accurately determine the speed of sound in the system. Furthermore, the new system scans all channels simultaneously, eliminating the acquisition delay that was found in the hot wire rig.

4.1.2 Experimental procedure & data processing

As with the hot wire Rijke tube, for each flame and burner position we alternate between two tests: a forcing test and a sound speed test. When the flame is thermoacoustically stable, we perform a forcing test using the same procedure that was used in the hot wire Rijke tube (§ 3.1.2). When the flame is thermoacoustically unstable, we use the active stabilization

with a phase-shift amplifier to stabilize the system. We then begin recording data from the microphones while the system is stable. After 3 seconds we terminate the stabilization, and record the pressure data as the thermoacoustic oscillations grow to a limit cycle. We record data for 15 seconds, following which we re-activate the stabilization and perform the sound speed test. The procedure for the sound speed test is identical to that described in § 3.1.2.

The data processing for the linearly stable flames and the sound speed test is the same as that described in § 3.1.2. For unstable flames, we must extract the linear growth rate from the nonlinear growth to a limit cycle. This growth has previously been modelled using a damped harmonic oscillator [148, 160]:

$$\ddot{\psi} - \psi(2s_r - \kappa\psi^2) + s_i^2\psi = 0 \quad (4.2)$$

where ψ is the state variable, s_r is the linear growth rate, κ is the strength of the nonlinearity, and s_i is the natural frequency. This system has oscillating solutions with an amplitude that grows like [160]:

$$\frac{dA}{dt} = s_r A + \frac{\kappa}{8} A^3 \quad (4.3)$$

where A is the amplitude of the oscillations. We obtain the linear growth rate, s_r , (and the nonlinearity, κ) by assimilating the amplitude of the time series pressure data into Eq. 4.3. This procedure proved more consistently accurate than the more common approach of performing linear regression on the initial portion of the growth, which requires ad-hoc decisions about when the growth transitions from linear to non-linear. The fit of the non-linear model to the data is shown in Fig. 4.2, with the corresponding linear model overlaid.

4.1.3 Experiment design

We investigate 24 different flames, which are selected to explore a wide range of thermoacoustic behaviour. Previous analytical studies of conical flames have shown that the thermoacoustic behaviour is determined by (i) the convective time delay, $\tau_c = L_f/\bar{u}$, which is the time taken for a perturbation travelling at the bulk velocity, \bar{u} , to travel along the length of the flame, L_f , and (ii) the mean heat release rate of the inner cone, \bar{Q} . We therefore parameterize the flames based on these two quantities.

We split the 24 flames into six groups of four flames, and select the composition of each flame such that the convective time delay is constant within each group and the mean heat




Fig. 4.2 The non-linear model (Eq. 4.3) is fitted to the pressure time series data from the unstable flames, allowing us to obtain the linear growth rate from observations of the non-linear growth. We show the fit of the non-linear model to the measured envelope of the pressure data, and the resulting linear model. This is shown for (a) the normalized pressure envelope, and (b) the log-normalized pressure envelope.

release rate varies. The convective time delays range from 9.5 ms to 17 ms in 1.5 ms increments, and the mean heat release rates range from 375 W to 600 W in 50 W increments. These flames produce thermoacoustic behaviour ranging from strongly damped, to neutral, to strongly driven.

The flow rates required to achieve the desired convective time delays and heat release rates are calculated using Cantera [161] and a simple linear model for a steady conical flame [162]. The linear model over-predicts the flame lengths, and therefore the convective time delays. We therefore verify the actual convective time delays experimentally by measuring the length of the steady flames from flame images. The flame properties are summarized in Table A.1 in Appendix A, and images of the steady flames are shown in Fig. 4.3.

We repeat the experiments for each flame 75 times to quantify the random error, which is much larger than it is in the hot wire rig. This is because small fluctuations in the co-flow slightly perturb the flame, causing slight stochastic variation in the measured growth rates. We perform more experiments than in the hot wire rig to better quantify, and average out, this effect.

4.1.4 Thermoacoustic model of the ducted laminar conical flame

We adapt the model described in § 3.1.4 to model the new rig. In this case we no longer need the base flow model for two reasons. Firstly, the rig was designed such that we can impose a steady co-flow using a mass flow controller, so we do not need to model the buoyant flow as we




Fig. 4.3 Processed steady flame images from the 24 flames. Images are grouped and artificially coloured according to their approximate convective time delay, τ_c . Each convective time delay is studied at four mean heat release rates, \bar{Q} . Flames with low mean heat release rate are shown in darker shades and flames with high mean heat release rate are shown in lighter shades.

do in the hot wire rig. Secondly, the higher sampling rate of the new acquisition system allows us to obtain accurate sound speed measurements directly from the pressure time series data.

The thermoacoustic network model was adapted to handle multiple coupled acoustic networks. This allows us to model the acoustic fields in both the duct and the burner tube, and their interaction. The code was also updated to automatically handle area changes so that blockage effects could be accounted for without having to manually apply linear feedback. The revised model is shown schematically in Fig. 4.4.

4.2 Assimilating data into the thermoacoustic model

In this section we use Bayesian parameter inference and model comparison to construct a quantitatively accurate thermoacoustic network model of the ducted laminar conical flame. We once again infer the unknown model parameters sequentially to ensure that each inference problem is well-posed. We begin by characterizing the sources of acoustic damping in the cold rig. We then introduce the flame and infer the parameters of the fluctuating heat release rate models.




Fig. 4.4 Diagram of the acoustic network model used in this study. The unknown model parameters are: R_* , the reflection coefficients at the boundaries, η_* , the strengths of the visco-thermal damping, and \mathcal{F} , the transfer function from velocity perturbations to heat release rate fluctuations.

4.2.1 Characterization of the cold rig

The model of the cold rig has nine unknown parameters, which comprise three complex reflection coefficients and three visco-thermal damping strengths. The reflection coefficients are from the upstream and downstream ends of the duct, R_u and R_d , and the base of the burner, R_b . We note that in this rig, R_u and R_d will be different because the upstream end terminates in a flanged connection to the plenum, while the downstream end is open to atmosphere. The three visco-thermal damping strengths correspond to the boundary layers on (i) the internal wall of the duct, η_d , (ii) the external wall of the burner, η_{be} , and (iii) the internal wall of the burner, η_{bi} . As in Chapter 3, the parameters η_* are multiplicative factors applied to an analytical model for visco-thermal damping [130], which in turn calculates the local linear feedback coefficients k_{**} . If the analytical model is correct, then we should find that $\eta_* = 1$ after the data has been assimilated.

We perform four sets of cold experiments, which we refer to as C1-C4. In C1 we use the decay rate and natural frequency of the empty duct to infer R_u , R_d and η_d . As with the hot wire Rijke tube, we could not repeatably calibrate the relative phase measurements sufficiently accurately to infer all the cold parameters simultaneously with weak priors. In this case,

however, we have five unknown parameters compared to the three in the hot wire rig. This makes it necessary to provide strong prior information for some parameters to avoid highly correlated posteriors. In Chapter 3 we found that the posterior expected value for η was always very close to unity. For this reason, we set the prior $\eta_d = 1$ and assign a small uncertainty to this value. We supply prior information about R_u and R_d from analytical models for the reflection of acoustic energy at flanged [163] and unflanged [131] terminations. We assign a large uncertainty to these priors because the models assume infinitely long ducts, infinitely thin walls, and infinite flanges, which are not representative of our rig.

In C2 we traverse a dummy burner through the rig. The dummy burner is a solid rod with the same exterior dimensions as the burner. From this set of experiments we assimilate R_u , R_d , η_d and η_{be} . We use the posterior values and uncertainties from C1 as the priors for R_u , R_d and η_d . We inflate the uncertainty in R_u , because we expect the upstream reflection coefficient to change due to the obstruction of the dummy burner. Similarly to C1, we assign a prior of $\eta_{be} = 1$ with small uncertainty.

In C3 we traverse the actual burner through the rig, but with a rigid plug in the base. We now assimilate all six parameters, but with prior information for R_u , R_d , η_d and η_{be} provided by the posterior from C2. The prior for η_{bi} is set to 1, and R_b is set to the theoretical value for a hard boundary. We place a low uncertainty on the value of R_b because the plug is effectively a rigid wall.

Finally, in C4 we traverse the burner through the tube with the choke plate in place, while supplying the burner with a sufficient mass flow rate of air for the choke plate to be choked. We again assimilate all six parameters, with prior information for R_u , R_d , η_d , η_{be} and η_{bi} provided by the posterior from C3. The prior for R_b is set to the theoretical value for a choked boundary, with large uncertainty.

The results of the characterization of the cold rig are shown in Fig. 4.5. The experimental observations are compared to both the prior and posterior model predictions. We see that the prior predictions are qualitatively accurate, but not quantitatively accurate because of small errors in the unknown parameters. The posterior predictions are quantitatively accurate and have defined confidence bounds. If the errors in the model of the cold rig were not removed, they would be incorporated into the fluctuating heat release rate parameters later on.

The prior and posterior joint distributions are shown graphically in Fig. 4.6. Each disc shows the joint distribution between a pair of parameters. The grey discs indicate the prior covariance, the orange discs indicate the covariance after assimilating the C1 data, the dark blue

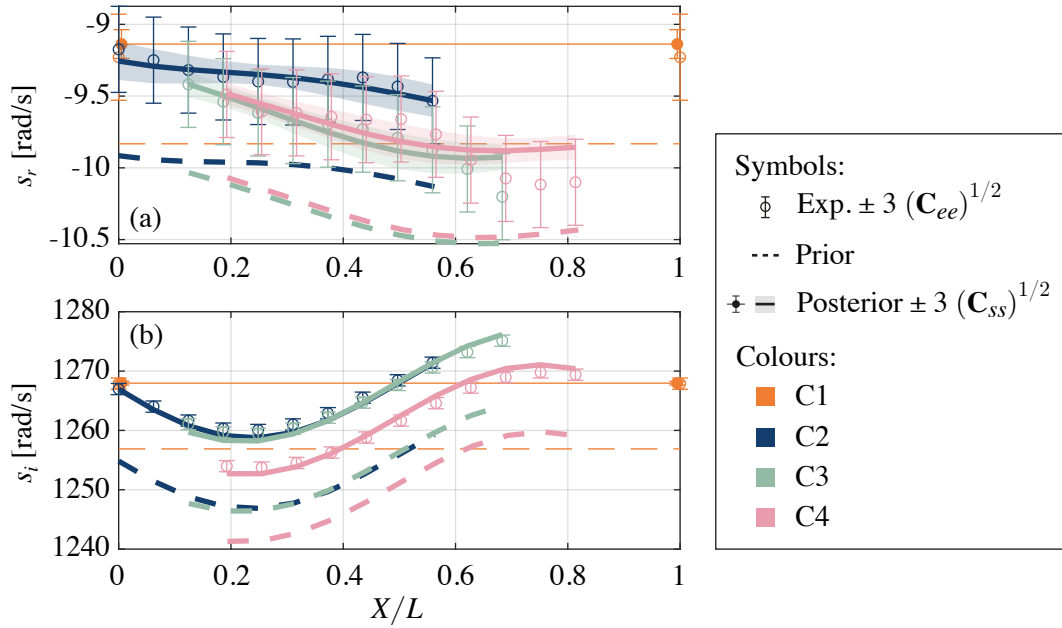


Fig. 4.5 Comparison of experimental measurements and model predictions of (a) growth rate, s_r and (b) angular frequency, s_i , plotted against normalized burner exit location, X/L for the three sets of cold characterization experiments. Experimental measurements are plotted (circles) with a confidence bound of 3 standard deviations. Prior model predictions are plotted (dashed lines) without confidence bounds. Model predictions after data assimilation are plotted (solid lines) with a confidence bound of 3 standard deviations.

discs indicate the covariance after assimilating the C2 data, the teal discs indicate the covariance after assimilating the C3 data, and the pink discs indicate the covariance after assimilating the C4 data.

In general, we see that the discs shrink as data is assimilated, because the parameter uncertainty reduces. We also see that the discs move away from the prior expected value. Both of these show that information is gained when data is assimilated [100]. The uncertainties in the η_* parameters do not, however, change considerably. This is because we had high confidence in the model and therefore set tight priors on η_* .

We also see from Fig. 4.6 that the posterior expected values for some parameters change as data from each subsequent experiment is assimilated. Most of these changes are small and can be attributed to random error in the experiments, which were all performed on different days. Two changes, however are clearly systematic. The first of these can be seen in the prediction of $\text{Im}(R_u)$ after C2-C4 are assimilated vs the C1 prediction. Recall that the C1 experiment was conducted on the empty duct, while C2-C4 had the burner in place. We therefore attribute this change to the disturbance of the burner on the upstream boundary, which causes a change in

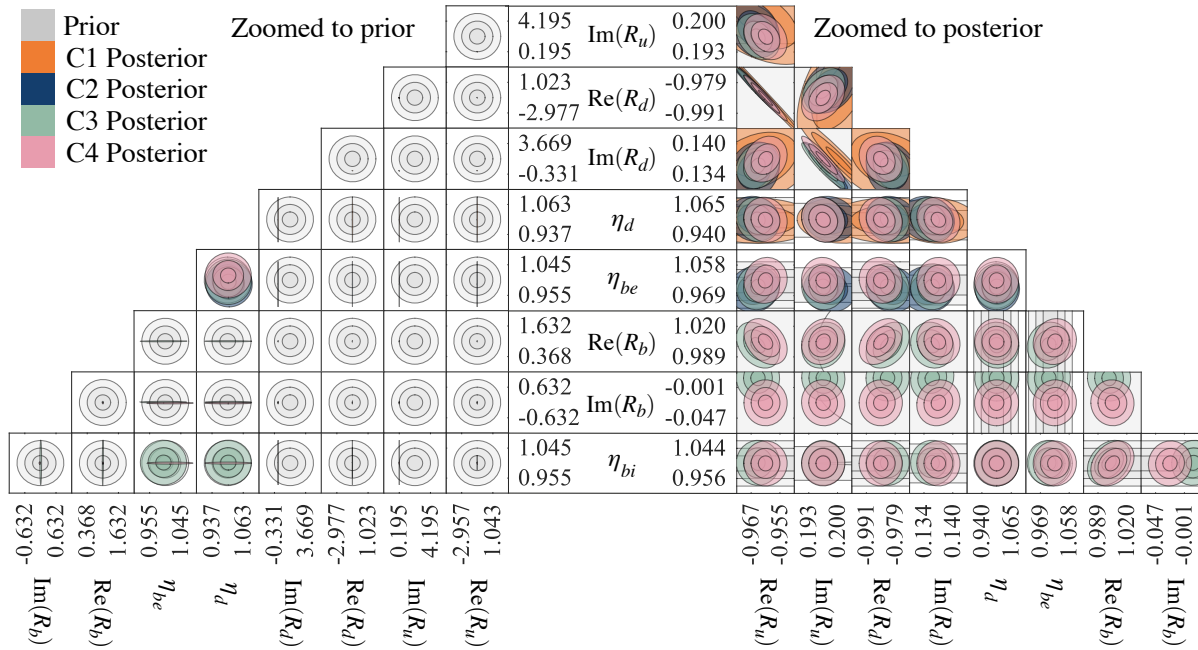


Fig. 4.6 Joint probability distributions, zoomed to the prior (left) and posterior (right), after assimilating data from the C1-C4 experiments. Each disc shows the joint distribution between a pair of parameters. The three rings represent one, two and three standard deviations from the mean.

$\text{Im}(R_u)$. The second can be seen in the prediction of $\text{Im}(R_b)$ after C4 is assimilated vs when C3 is assimilated. The C3 experiment used a burner with a brass plug in the base, while the C4 experiment had the choke plate installed. We therefore expect to see a slight change in R_b between these two experiments.

Finally, we see that some parameters could not be decorrelated, as was the case in the hot wire rig. These are $\text{Re}(R_u)$ and $\text{Re}(R_d)$, and $\text{Im}(R_u)$ and $\text{Im}(R_d)$. It is not possible to further decorrelate the parameters without more accurate measurements of the pressure phase.

4.2.2 Assimilating heat release rate models

With the acoustic damping of the cold rig carefully characterized, we can now assimilate models for the fluctuating heat release rate of the flame, which drives or damps the acoustic oscillations depending on its phase relative to the pressure [124]. We fix the parameters of the cold rig to the values inferred in the previous subsection, except for the downstream reflection coefficient, which we expect to vary significantly with the increase in temperature when the flame is introduced. To account for this, we calculate a correction factor to the model for the

reflection coefficient that was proposed by Levine and Schwinger [131]. This allows us to use the corrected Levine-Schwinger model to calculate R_d when the flame is introduced.

We begin by comparing several models for the fluctuating heat release rate in our rig using experimental observations from three flames. We then infer the most probable parameters of the selected model using experimental observations from all 24 flames shown in Fig. 4.3.

Selecting a model for the fluctuating heat release rate

The fluctuating heat release rate is modelled as a feedback mechanism from the acoustic velocity into the energy equation, which we label k_{eu_f} [109]. We propose a model for k_{eu_f} in the form of a typical flame transfer function:

$$k_{eu_f} = \frac{\gamma - 1}{\gamma} \frac{\bar{Q}}{\bar{p}\bar{u}A} \mathcal{F} \quad (4.4)$$

where γ is the ratio of specific heats, \bar{Q} is the mean heat release rate of the flame, \bar{p} is the mean pressure at the injection plane, \bar{u} is the mean velocity at the injection plane, and A is the cross-sectional area of the duct at the injection plane. \mathcal{F} is the complex-valued flame transfer function (Eq. (4.1)), which relates fluctuations in velocity, u' , to fluctuations in heat release rate, Q' . The magnitude, $|\mathcal{F}|$, is the thermoacoustic gain, and the angle, $\angle \mathcal{F}$ is the thermoacoustic phase delay.

We infer the most probable flame transfer function from experimental observations of the growth rate, frequency and Fourier-decomposed pressure. We begin by traversing three of the 24 flames through the duct. For this initial study, we choose the three flames with the shortest convective time delay and lowest mean heat release rate. These flames remain linearly stable at all burner positions but present different thermoacoustic decay rates. We assume that the flame transfer function should not depend on the position of the burner in the duct so, for each flame, we seek a single flame transfer function that is valid for all burner positions.

At any burner position, the flames are exposed to two distinct acoustic velocity perturbations: that from the acoustic field within the burner tube, and that from the acoustic field in the duct. We test two models from the literature and propose two new models. Model 1 considers the blockage of the burner tube but neglects the acoustic field inside the burner tube and assumes that the flame reacts only to the velocity perturbation in the duct [141, 144, 164, 165]. Model 2 includes both acoustic fields, but assumes that the flame reacts only to the velocity perturbation in the burner tube. This is based on studies that have measured flame transfer functions of laminar conical flames *ex-situ* [37, 140, 166], and assumes that these results apply when the

flame is *in-situ*. We propose models 3 and 4, which include both acoustic fields and assume that the flame reacts to both sources of velocity perturbation. In model 3 the flame reacts to both sources of velocity perturbation, with a different gain and a different phase delay for each source. In model 4 the flame reacts to both sources of velocity perturbation, with a different gain but the same phase delay for each source.

The four models are shown graphically in Fig. 4.7. Models 1 and 2 have two real parameters: the gain and phase delay of the flame transfer function, \mathcal{F} . Model 3 has four real parameters: the gain and phase delay of two flame transfer functions, \mathcal{F}_d and \mathcal{F}_b . Model 4 has three real parameters: two gains, $|\mathcal{F}_b|$ and $|\mathcal{F}_d|$ and a single phase delay, $\angle\mathcal{F}_b = \angle\mathcal{F}_d$. In models 3 and 4, the subscripts *b* and *d* refer to the burner and duct respectively.




Fig. 4.7 Four models for the fluctuating heat release rate of a ducted conical flame. (a) Model 1: the flame reacts to the velocity perturbation in the duct alone. The acoustics in the burner are not modelled. (b) Model 2: the flame reacts to the velocity perturbation in the burner alone. (c) Model 3: the flame reacts to the velocity perturbations in both the duct and the burner with different gains and phase delays. (d) Model 4: the flame reacts to the velocity perturbations in both the duct and the burner with different gains, but the same phase delay. \hat{Q} and \hat{u} refer to the normalized perturbations, \ast'/\bar{x} , \mathcal{F} is the flame transfer function, and the subscripts *b* and *d* refer to the burner and duct respectively.

We assimilate the data into each model to find the most probable flame transfer functions. The posterior model predictions for all four models are compared against experimental observations in Fig. 4.8. We see from the results of model 1, shown in Fig. 4.8(a), that neglecting the acoustic field in the burner tube leads to a model that cannot fit the data. Most prominently, it is clear from Fig. 4.8(a.i) that a flame transfer function based on the duct velocity perturbations must predict zero thermoacoustic effect when the flame is placed at the duct's velocity node, which is just downstream of $X/L = 0.4$. This effect is clearly not observed in the data, which

shows a strong thermoacoustic effect when the burner is placed at the duct's velocity node. Further, we see from Fig. 4.8(a.ii) that model 1 cannot predict the frequency correctly because the effect of the acoustic field in the burner tube has been neglected. The inferred total uncertainty, $(C_{tt})^{1/2}$, has been plotted as the data error bars, while the parametric uncertainty has been plotted as the model error bars. We see that the uncertainty is large for model 1 because of the structural error in the model.

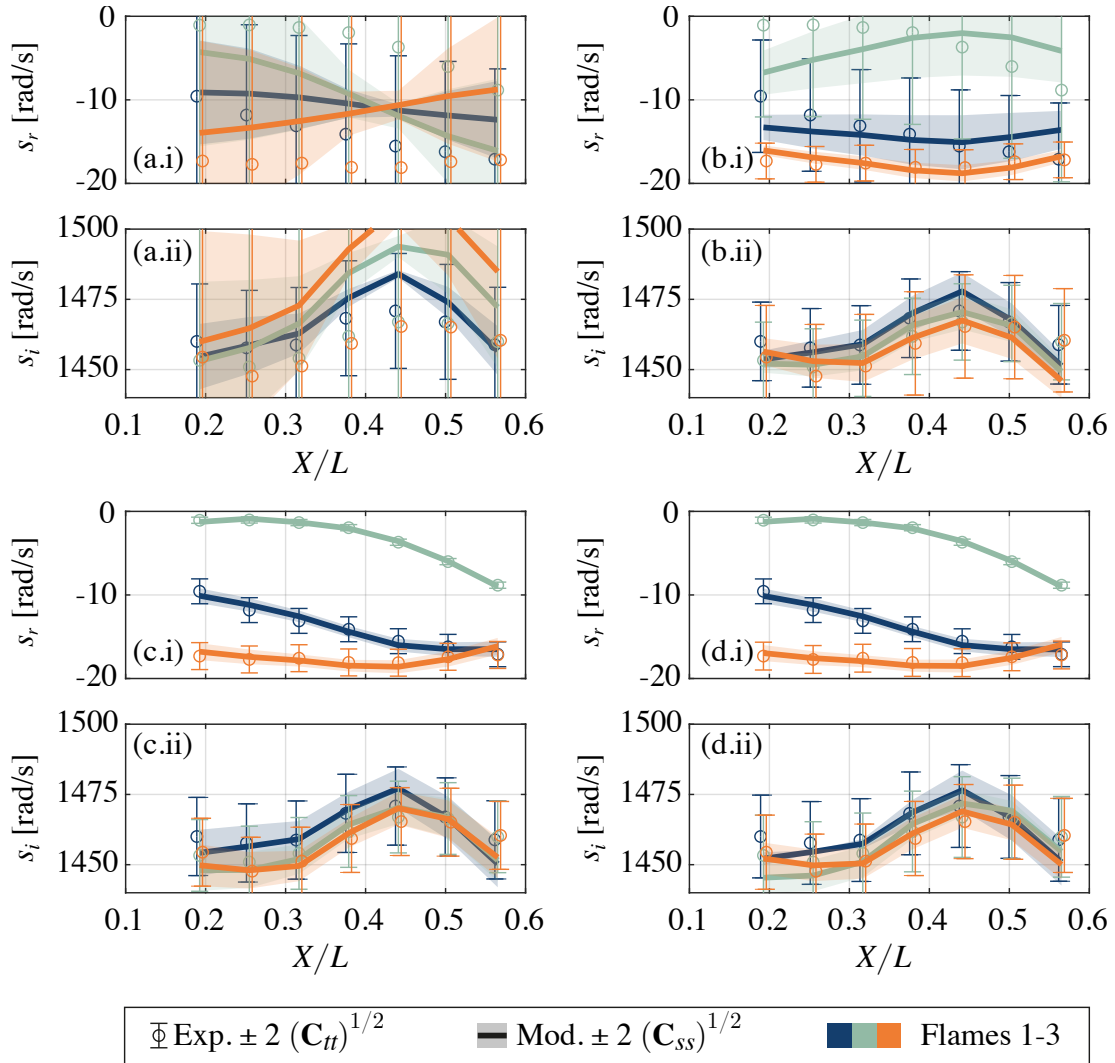


Fig. 4.8 Posterior model predictions and experimental measurements of (i) growth rate, s_r , and (ii) angular frequency, s_i , plotted against normalized burner position, X/L for three different flames. Model predictions are plotted as solid lines with the shaded region indicating the parametric uncertainty. Experimental measurements are plotted as circular markers with error bars indicating the random and inferred systematic uncertainty. The results for each of the three flames are shown in different colours, which correspond to the colours in Fig. 4.3. The posterior model predictions of (a) model 1, (b) model 2, (c) model 3, and (d) model 4 are shown.

We see from Fig. 4.8(b) that, while model 2 fits the data better than model 1, it suffers from a similar limitation. Model 2 must predict steadily decreasing thermoacoustic effect as the burner approaches the duct pressure nodes, which are near $X/L = 0$ and 1. The pressure fluctuations in the duct give rise to the acoustic field in the burner tube, so when the burner is placed at the duct pressure node, the acoustic field in the burner tube vanishes, along with the heat release rate fluctuations. This can be seen in Fig. 4.8(b.i) where the model predictions converge towards a common growth rate as the burner approaches either end of the duct. It is clear from the data, however, that the thermoacoustic effect does not vanish as the burner approaches the pressure node, as can be seen from the wide spread in growth rate measurements at $X/L = 0.2$. We notice from Fig. 4.8(b.ii) that including the burner tube acoustic field in the model allows the model to make more accurate frequency predictions. Finally, while the inferred uncertainty is smaller than for model 1, it is still large because of the structural error in the model.

Motivated by the shortcomings of models 1 and 2, we propose model 3 to allow the flame to react to both sources of velocity perturbation. We see from Fig. 4.8(c) that model 3 fits the data well for all three flames, at all burner positions, and that the inferred uncertainty is small. However, from the phenomenology of the problem we expect that each flame should react with a single characteristic time delay, regardless of the source of the perturbation. We therefore propose model 4 which enforces this constraint. We see from Fig. 4.8(d) that model 4 also fits the data well for all three flames, at all burner positions, and the inferred uncertainty remains small.

While models 1 and 2 are easy to discard, it is more difficult to discriminate between models 3 and 4, so we use Bayesian model comparison to rank the models and identify the best one. The model ranking metrics are summarized in Fig. 4.9. Comparing the log-marginal likelihoods of each of the models, we see that models 3 and 4 are substantially more probable than models 1 and 2, with model 4 being marginally more probable than model 3. This is consistent with our expectations based on the phenomenology of the problem. Models 1 and 2 are simple and therefore have smaller Occam penalties, but they fit the data poorly and are therefore penalized by small best fit likelihoods. By comparison, models 3 and 4 fit the data well and therefore have large best fit likelihoods, which outweigh the Occam penalty from increased complexity. While model 3 fits the data slightly better than model 4, the additional complexity of model 3 is not justified by the improvement in fit, and so model 4 is the most probable model given our experimental data.




Fig. 4.9 Model ranking metrics for four candidate models. The best fit likelihood (BFL) measures how well the model fits the data. The Occam factor (OF) penalizes the model based on its parametric complexity. The marginal likelihood (ML) is the overall evidence for a given model, and is the product of the BFL and the OF (i.e. $\log(\text{ML}) = \log(\text{BFL}) + \log(\text{OF})$). The model with the largest marginal likelihood is the most likely model, given the experimental data.

Finally, in Fig. 4.10 we compare the inferred uncertainty to the known uncertainty, which was estimated based on the mean random error for all tests. Random error is the dominant source of quantifiable experimental uncertainty in this rig. We see that the inferred uncertainty in both growth rate and frequency for models 1 and 2 is significantly larger than the known uncertainty, indicating either systematic error in the experiments or structural error in the model. By comparison, the inferred uncertainty for models 3 and 4 is comparable to the known uncertainty. This suggests that the systematic error in models 1 and 2 is due to structural error in the models, rather than systematic measurement error, because it has been eliminated in models 3 and 4.

Inferring the parameters of the fluctuating heat release rate model

We now apply the most probable model to all 24 flames. As a reminder, the flames are categorized in six groups of four flames, where the flames in each of the six groups have the same convective time delay but varying mean heat release rate. Each flame is traversed from 0.2 m to 0.35 m from the duct inlet, in 0.05 m increments. The experimental results are

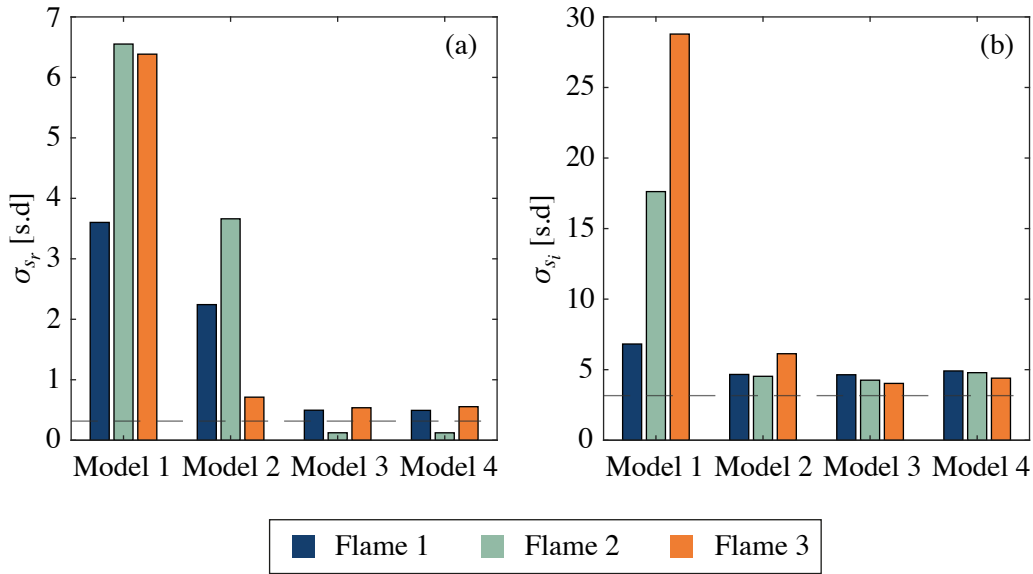


Fig. 4.10 Inferred uncertainties for each flame, modelled by each of the four candidate models. The uncertainty in the (a) growth rate, σ_{s_r} , and (b) frequency, σ_{s_f} , is shown in units of standard deviations. The dashed line represents the known uncertainty, which was estimated based on the random error and precision of the measurement chain.

shown in Fig. 4.11, from which we see that the chosen flame parameterization has produced a convenient basis for exploring thermoacoustics in conical flames. Changing the convective time delay changes the thermoacoustic behaviour, while changing the power mainly changes the strength of the thermoacoustic effect. The data includes neutral flames (blue and orange), driving flames (teal, red and yellow) and damping flames (pink). This allows us to test our inference framework on a wide range of flame dynamics.

We use the experimental data to infer the parameters of model 4, chosen in the previous section. The posterior model predictions are compared with the experimental data for all 24 flames at 4 flame positions in Fig. 4.12. We see that the model predictions are within the experimental uncertainty bounds for all the flames at all positions, except for the frequency prediction of the highest power flame in group 6 (see Fig. 4.12 (f.ii)). For this experiment the model over-predicts the frequency by 2.4 Hz, which is less than 1% of the measured value. We should expect increased error in the frequency predictions for longer flames, because the frequency predictions are sensitive to the sound speed field, which becomes poorly approximated in the 1D network model for longer flames.

The results from Fig. 4.12 are repeated in Fig. 4.13, but are grouped according to flame power rather than convective time delay, and the axis scales have been matched between the plots. This makes the model fit less clear, but highlights some important physical trends. Firstly,

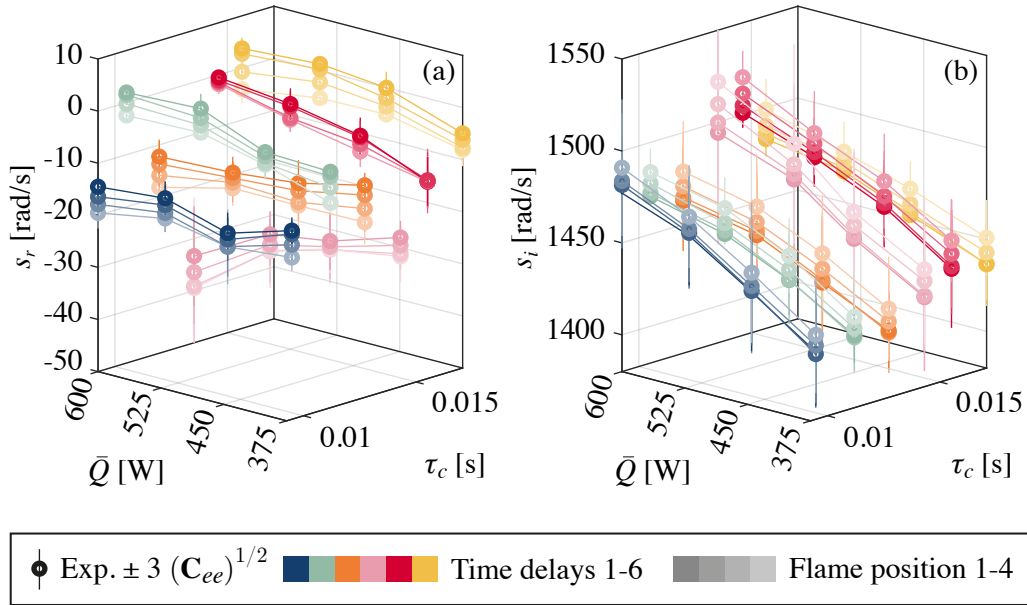


Fig. 4.11 Experimental measurements of (a) growth rate, s_r , and (b) angular frequency, s_i , plotted against the flame convective time delay, τ_c and mean heat release rate, \bar{Q} . The experimental data points are shown with circular markers, with vertical lines representing a confidence interval of 3 standard deviations. A thin connecting line has been added between experimental data points as a visual aid. The results for each of the four burner positions are shown, with darker shades representing lower burner positions and lighter shades representing higher burner positions. The results are coloured according to the flame groups, which correspond to the colours in Fig. 4.3.

the growth rate plots emphasize the fact that increasing the flame power while keeping the convective time delay constant strengthens the thermoacoustic effect. Secondly, it is clear that several flames display the same thermoacoustic behaviour, as seen by the overlapping growth rate measurements/predictions. We should therefore expect that these flames have similar flame transfer functions.

We see from Figs. 4.12 and 4.13 that, although the model was selected based on the lowest power flames from groups 1-3, it remains accurate at higher powers and longer convective time delays once the correct model parameters are found. This demonstrates the power of a physics-based, data-driven modelling approach. Once the best model is selected, it can be applied to cases well outside the range of the data used to select the model. This is particularly useful for thermoacoustic systems because the model selection process can be carried out using data from low power experiments, which are cheaper and safer to conduct, and then applied to higher power cases using only a few experimental observations to find the most probable model parameters.

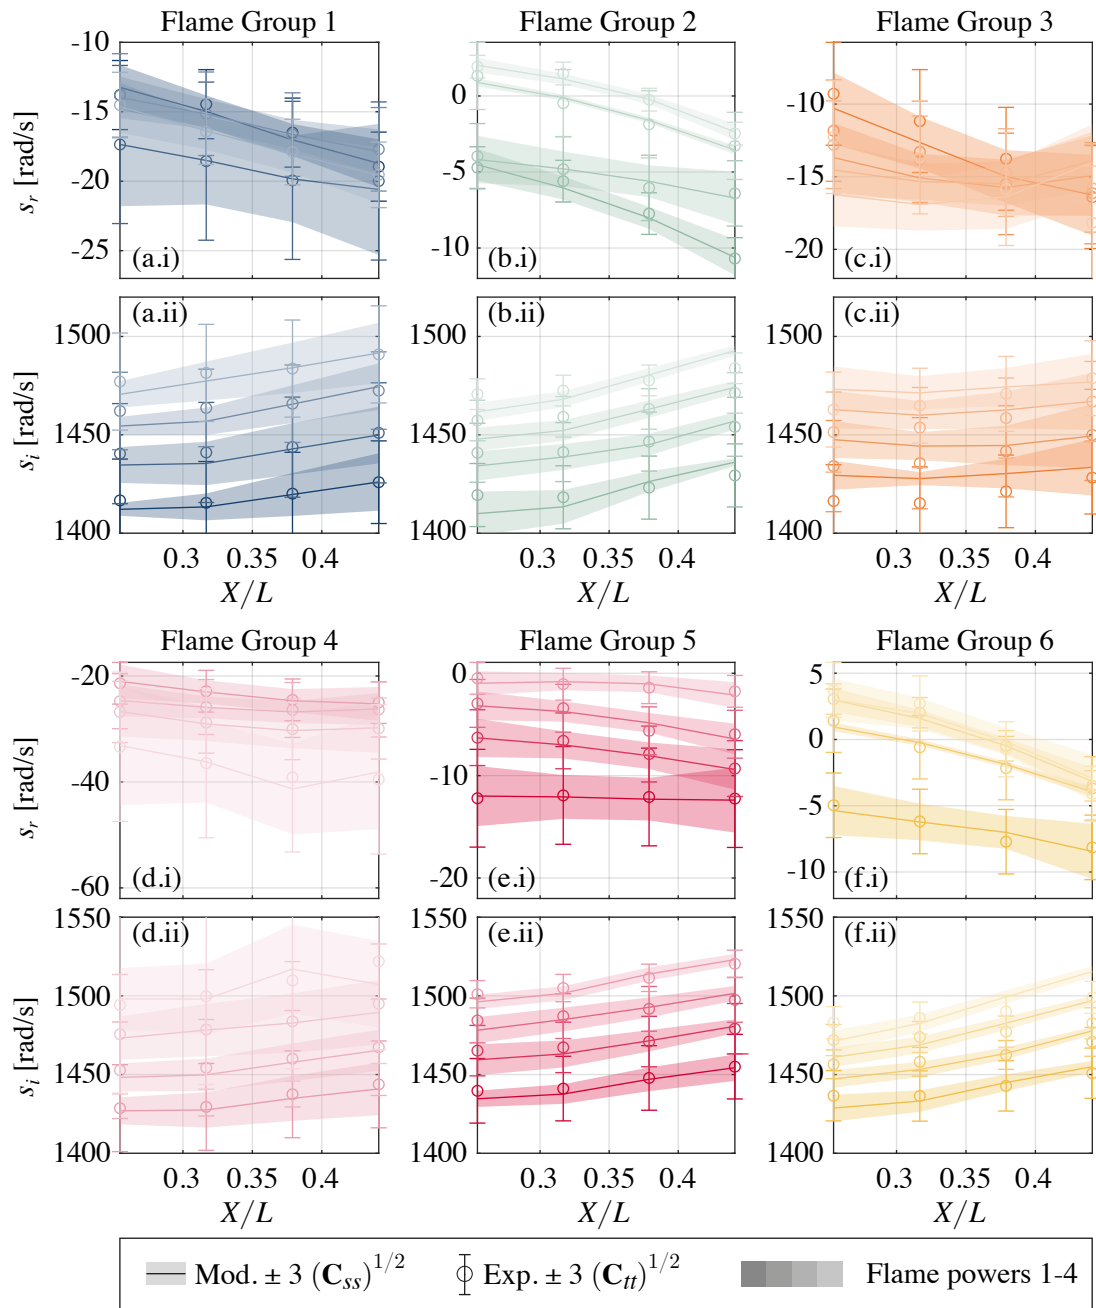


Fig. 4.12 Posterior model predictions and experimental measurements of (i) growth rate, s_r , and (ii) angular frequency, s_i , plotted against normalized flame position, X/L . Model predictions (solid lines) are plotted with confidence bounds of 3 standard deviations. Frames (a)-(f) show each of the six groups of flames with constant convective time delay, coloured according to the colours in Fig. 4.3. Darker shades correspond to lower flame powers, and lighter shades correspond to higher flame powers.

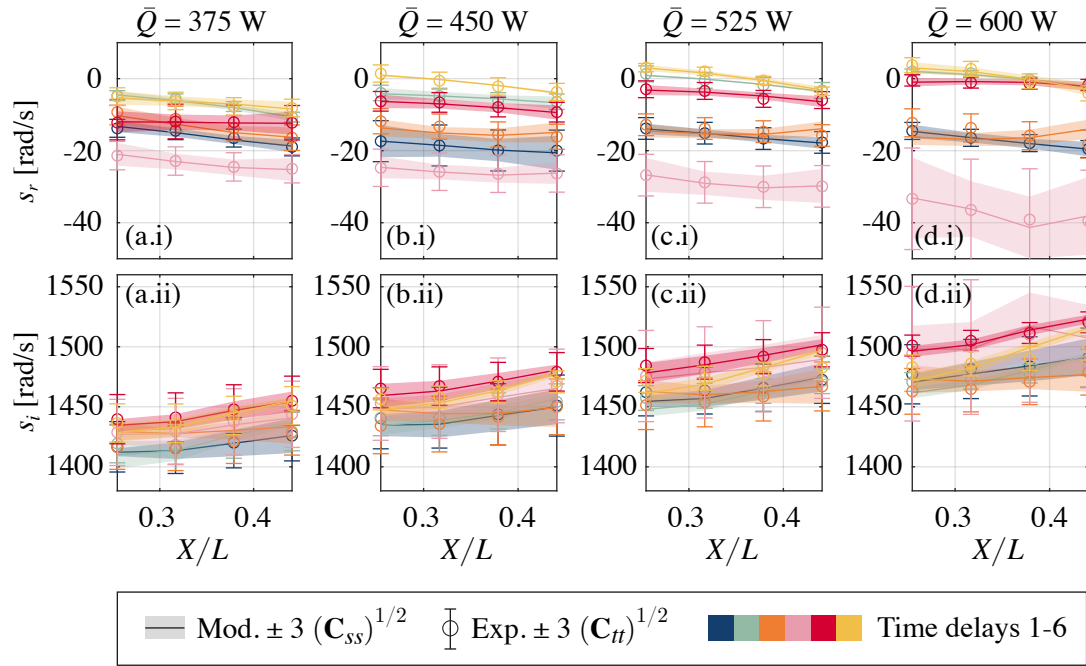


Fig. 4.13 Posterior model predictions and experimental measurements of (i) growth rate, s_r , and (ii) angular frequency, s_i , plotted against normalized flame position, X/L . The model predictions are shown as solid lines with a shaded patch representing the confidence bounds. The experimental results are shown with circular markers, with vertical lines representing confidence bounds. Frames (a)-(d) show the results for each of the four flame powers. The results for each of the six convective time delays are shown with different colours, corresponding to those in Fig. 4.3.

We have shown that the inference process results in a quantitatively accurate model, but it is equally important that the inferred flame transfer functions are physically meaningful. We focus on the flame transfer function between heat release rate fluctuations and velocity perturbations in the burner tube, because this is most commonly discussed in the literature. In Fig. 4.14, we plot the 24 inferred flame transfer functions for internal perturbations on polar axes with confidence bounds of 2 standard deviations. First, we see that the flames are appropriately placed on the polar plot, with driving flames (teal, red and yellow) clustered together with roughly 90° phase lag between the velocity perturbation and the heat release rate fluctuation. This corresponds to heat release rate fluctuations peaking at moments of maximum acoustic pressure, as we expect for thermoacoustically driving flames. Similarly, the damping flame (pink) has a phase lag of 270° , corresponding to heat release rate fluctuations peaking at moments of minimum acoustic pressure. The neutral flames (blue and orange) have different phase delays, with the orange flames having a phase lag of 180° and the blue flames having a phase lag of 270° , but both fall in the neutral region of the polar plot.

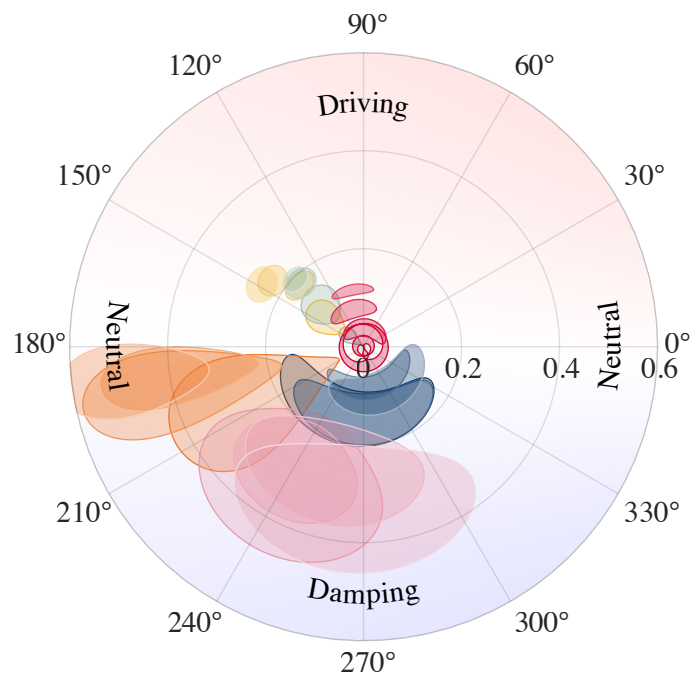


Fig. 4.14 Polar plot of the inferred flame transfer functions for internal perturbations for all 24 flames. The gain is shown on the radial axis, and phase delay on the angular axis. The shaded areas represent a confidence region of 2 standard deviations. The colours correspond to those in Fig. 4.3, with darker shades representing lower flame powers and lighter shades representing higher flame powers. The red-white-blue contour in the background represents the effect of flame transfer function gain and phase on the instability growth rate, where red represents positive growth rates, white represents no growth and blue represents negative growth rates.

The polar plot also shows that the uncertainty in the inferred flame transfer functions depends on two factors: (i) the flame behaviour and (ii) the measurement uncertainty. We see in Fig. 4.14 that the neutral flames (blue and orange) have large uncertainties. This is because the thermoacoustic effect is weak, and therefore difficult to observe from pressure measurements alone. By contrast, the driving flames (teal, red and yellow) have smaller uncertainties, because the thermoacoustic effect is strong and therefore easy to observe. The damping flame (pink) has a large uncertainty even though the thermoacoustic effect is strong. This is because the oscillations decay quickly, meaning that the decay rate and natural frequency must be measured from few oscillations, which increases the measurement uncertainty. It is convenient that we have high certainty in the behaviour of driving flames, because these are typically of most interest to designers.

Finally, we check the validity of the inferred flame transfer functions by comparing them to directly measured values. We did not directly measure the flame transfer function in our

experiments, so instead we compare the inferred flame transfer functions to direct measurements from similar systems in the literature. No experimental studies in the literature have measured the response of a flame to forcing from outside the burner tube. We can therefore only compare the inferred flame transfer functions between heat release rate and velocity perturbations from within the burner tube to those from the literature. Cuquel et al. [49] have shown that for conical flames, flame confinement only affects the flame transfer function for confinement ratios (burner radius / duct radius) above 0.44. Our rig has a confinement ratio of 0.125, so we expect that we can compare the inferred flame transfer function for internal velocity perturbations to those directly measured on unconfined flames.

The results of the comparison are plotted in Fig. 4.15. We show results from three experimental studies [49, 139, 167] and one analytical model [143]. The experimental studies all considered unconfined, premixed, laminar, conical flames forced through the burner tube. The burner of Kornilov [139] was similar to that in the current study, while the burners of Schuller et al. [167] and Cuquel et al. [49] had a diameter of roughly double that in the current study. The analytical model of Schuller et al. [143] considered an unconfined, premixed, laminar conical flame of arbitrary diameter.

We plot the gain and phase of the flame transfer function for internal perturbations against reduced frequency in Fig. 4.15. We use the same definition for reduced frequency as Schuller et al. [167]: $\omega_* = s_i R / (S_L [1 - (S_L / \bar{u})^2]^{1/2})$, where s_i is the frequency of oscillations, R is the burner radius at the injection plane, S_L is the unstretched laminar flame speed and \bar{u} is the bulk velocity in the burner tube.

The flame transfer function gains are compared in Fig. 4.15(a). Considering only the experimental data taken from literature, we note that despite the similarity of the experimental configurations, the measured flame transfer functions vary significantly. While the gain measurements agree fairly well at low reduced frequencies, there is significant spread in the measurements between reduced frequencies of about 7 and 20. Considering the spread in the direct measurements, we see that the inferred gains agree reasonably well with the direct measurements for the blue, teal and red flames. The inferred gains for the orange, pink and yellow flames are slightly higher than the direct measurements. The orange flames have a phase delay of 180° , which makes the gain difficult to infer because the growth rate becomes insensitive to gain. This is reflected in the large uncertainty in gain for the orange flames. The pink flames were strongly damping which led to larger experimental error. The increased experimental error was estimated from the variance in 75 experimental observations, which

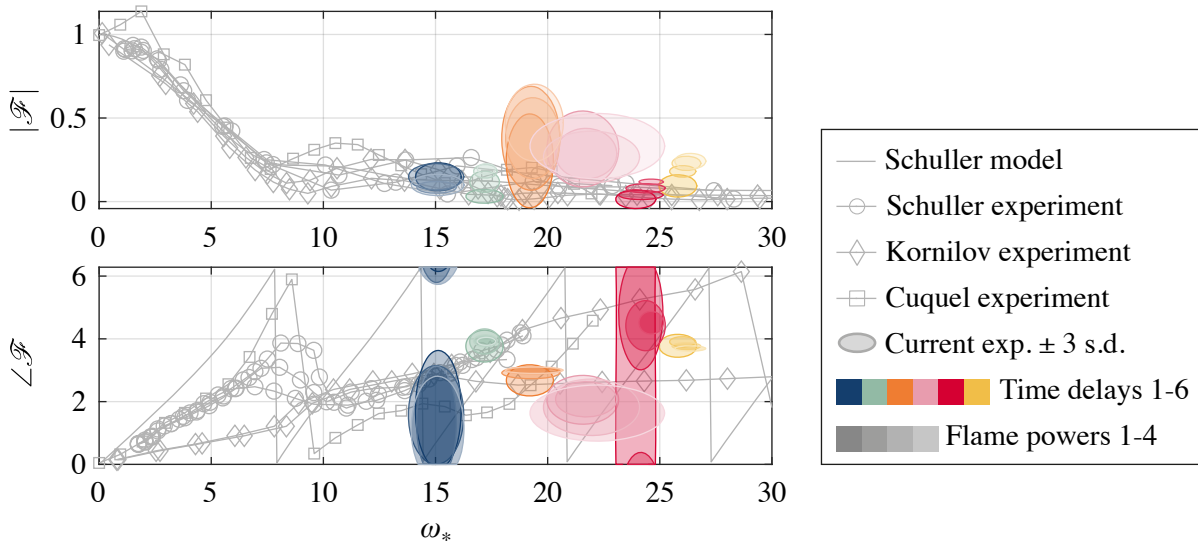


Fig. 4.15 Comparison of the inferred flame transfer functions for internal perturbations (colours) with direct measurements (lines with symbols) and an analytical model (line) from the literature. The flame transfer function (a) gain, $|\mathcal{F}|$, and (b) phase delay, $\angle\mathcal{F}$, is plotted against the reduced frequency, ω_* . The inferred flame transfer functions are shown as ellipses indicating a confidence interval of 3 standard deviations, with colours corresponding to those in Fig. 4.3. There is good agreement for the inferred gain when (i) the thermoacoustic effect was strong and therefore easily observable, and (ii) the experiments had low systematic error. A larger discrepancy is therefore expected for (i) the orange flames because they were thermoacoustically neutral, and (ii) the pink and yellow flames because they contained unquantified systematic error. The direct phase measurements (grey lines) do not agree with each other, even though those experiments were similar to each other, indicating that the phase is highly sensitive to the experimental configuration. The inferred phase measurements (colours) are similarly scattered.

appears to underestimate the error because the uncertainty bounds do not cover the direct measurements. The yellow flames also have a component of unquantified systematic error, which is likely to come from the error in approximating the sound speed field in the 1D network model for these long flames. In the case of both the pink and yellow flames, the systematic error could be estimated if a suitable model for the flame transfer function were available.

The flame transfer function phases are compared in Fig. 4.15(b). Considering the experimental data first, we note that the phase measurements show almost no agreement at any of the reduced frequencies. We therefore cannot expect that the phases inferred from our experiments should show any meaningful agreement with the direct measurements from the literature. The variability of the direct phase measurements is particularly problematic due to the severe sensitivity of the thermoacoustic behaviour to the phase delay [19]. Had we directly applied the flame transfer functions from any of the studies from the literature in our model, we would not have been able to produce a quantitatively accurate predictions of the system

behaviour. This motivates the approach of inferring flame transfer functions *in-situ*, rather than relying only on results in the literature or even direct measurements taken from the unconfined burner.

4.2.3 Conclusion

In this chapter, we introduce additional complexity to the system by replacing the electric heater with a laminar conical flame. For the cold rig, this introduces the challenge of generating a quantitatively accurate model of a system with coupled acoustic ducts. For the hot rig it introduces the complexity of richer heat release rate dynamics, and more random variation in the experimental data. We once again perform automated experiments to collect the data, which we assimilate into physics-based models of the thermoacoustic system.

Using Bayesian inference, we infer the flame transfer functions of 24 flames *in-situ* from pressure measurements alone. The state-of-the-art for measuring flame transfer functions typically requires optical access to the flame, which is often not available in industrial combustion chambers. The flame must therefore be characterized outside the chamber, where it is not exposed to the confinement, heat loss and flame-flame interactions that are present within the chamber, to which the flame transfer function has been shown to be sensitive [37, 43, 49, 50]. The ability to obtain flame transfer functions with the flame *in-situ* from practically achievable measurements is therefore a valuable result.

We verify the inferred flame transfer functions by comparing them to direct measurements from similar rigs in the literature. The direct measurements themselves were somewhat inconsistent for the gain, and totally inconsistent for the phase delay. Nonetheless, we show that the inferred quantities are generally within the expected range, and any significant discrepancies are easily explained.

While a handful of previous studies have calculated flame transfer functions without optical access, none has assessed their uncertainties, and therefore they tend to be over-confident in their results [153–159]. We rigorously quantify the uncertainties in the inferred flame transfer functions and find, as expected, that the flame transfer functions are most precise if (i) the thermoacoustic effect is strong, and (ii) the measurement uncertainty is small. This can help to guide future experiments on industrial rigs.

More generally, we once again demonstrate that the Bayesian framework forces the researcher to adhere rigorously to both the physics and the experimental data. This often reveals shortcomings in existing models. In this chapter we find that the experimental data cannot be

explained if the heat release rate depends only on velocity perturbations in one of the ducts, which is a common assumption in the literature. We find that the data contains strong evidence that the heat release rate depends instead on the velocity perturbations in both the duct and the burner tube.

Chapter 5

Ducted turbulent conical flames

This chapter contains excerpts from a conference paper on inferring flame transfer functions of turbulent conical flames in a duct [[168](#)]. This work was performed independently by Matthew Yoko.

5.1 Introduction

In Chapter 4, we applied our Bayesian framework to laminar conical flames in a duct. This allowed us to infer the most probable values of the flame transfer function for 24 flames, with the flames *in-situ*, using only pressure measurements. Unlike other methods of obtaining flame transfer functions from pressure time series data [157–159], our method rigorously quantifies the uncertainties in the inferred flame transfer functions.

The method we have used makes no assumptions about the structure of the flame, and so it could theoretically be applied to arbitrarily complex flames with turbulence and swirl. We showed, however, that the method is sensitive to the experimental uncertainty. This raises questions about whether the method can still work for turbulent flames, which (i) produce additional broadband noise, and (ii) exhibit stochastic thermoacoustic behaviour, both of which would increase the experimental uncertainty. To test this, we extend the work in Chapter 4 to turbulent conical flames. This allows us to assess the influence of turbulent noise without having to significantly redesign the existing burner.

In Chapter 4 we validated our inferred flame transfer functions by comparing them with direct measurements from the literature. The results were somewhat inconclusive, especially due to the wide inconsistency between phase delay measurements in the literature. In the current chapter, we improve on this by directly measuring the fluctuating heat release rates to provide a more rigorous validation of the method.

5.1.1 Experimental configuration

The experimental rig is slightly modified from Chapter 4, as illustrated in Fig. 5.1. We only report the changes compared with § 4.1.1. These are as follows:

Firstly, the burner is fitted with a turbulence generation grid, which is a stainless steel disk with a thickness of 0.3 mm, perforated with 19 holes of diameter 1.5 mm. The grid is positioned 35 mm upstream of the injection plane. Secondly, we install a hot wire probe at the upstream end of the duct to obtain an absolute reference for the phase of the velocity fluctuations. Finally, we use a high speed camera (Phantom Miro M310) fitted with a zoom lens (Nikkor 80-300 mm) to capture footage of the steady and perturbed flames in order to directly quantify the heat release rate fluctuations.

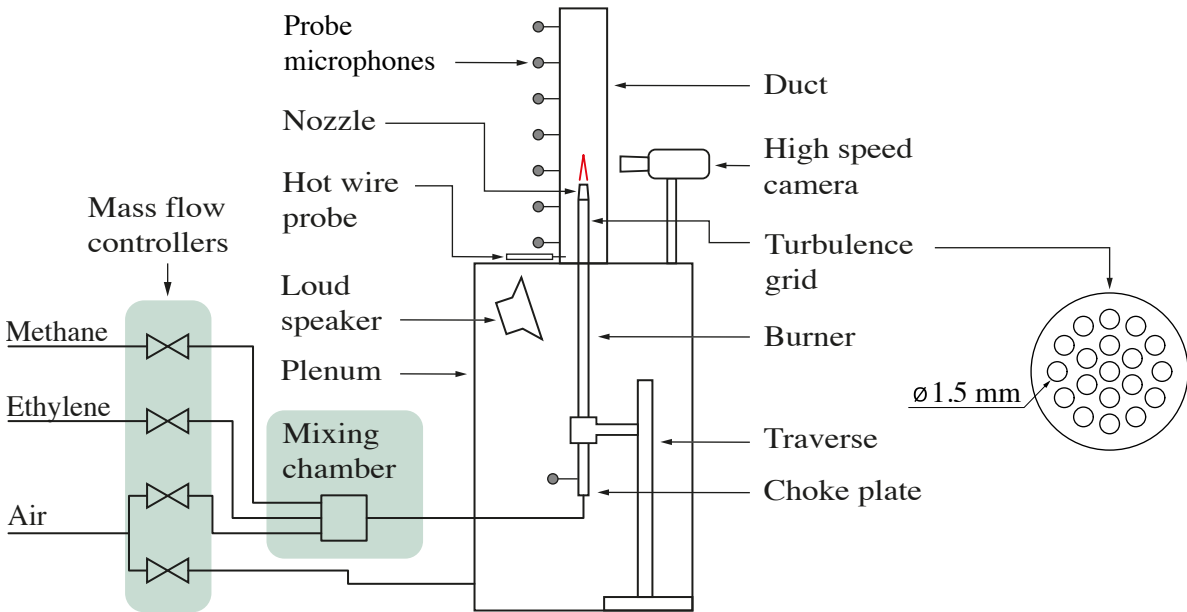


Fig. 5.1 Diagram of turbulent ducted flame rig.

5.1.2 Direct measurements of the fluctuating heat release rate

We synchronize the camera trigger to the acoustic forcing, and capture images of 100 thermoacoustic cycles at 20 phase angles (2000 frames in total). By phase-averaging the frames, we isolate the coherent perturbations caused by the acoustic forcing and remove the stochastic turbulent perturbations, as illustrated in Fig. 5.2. The camera exposure is fixed to $150 \mu\text{s}$ for all tests. The video files are large, so we only record three experiments per flame.

The high speed camera captures the unfiltered emission of the flame in the visible range, which we use as a proxy for the heat release rate. It is not generally guaranteed that the unfiltered light emission can be related to the heat release rate. Several studies have, however, shown the unfiltered light emission to be a reasonable approximation for heat release rate for various premixed flames [169–172], many of which were similar to (or more complex than) the flames we study in this chapter. We therefore assume that $I \propto Q$, where I is the spatially integrated intensity of the image, and Q is the instantaneous heat release rate. The relative fluctuating intensity, $\hat{I} = I'/\bar{I}$, and the relative fluctuating heat release rate, $\hat{Q} = Q'/\bar{Q}$, are therefore equivalent.

We perform a windowed Fourier-decomposition on Q' to obtain $|Q'|$ and $\angle Q'$, where $\angle Q'$ is referenced to the phase of the pressure signal at the reference microphone. The probe microphones induce a small phase delay, which we correct with the absolute phase reference




Fig. 5.2 Instantaneous (top) and phase-averaged (bottom) images of the perturbed flame at ten phase angles (artificial colour).

from the hot wire probe. We use the mean of the Fourier-decomposed signal to compute the expected value of \hat{Q} . We use the standard deviation to estimate the random error in the direct measurements, which is primarily due to the stochastic turbulent fluctuations.

5.1.3 Flame properties

In this chapter we study 15 turbulent conical flames. We control the flame properties by varying the fuel and air flow rates using the mass flow controllers. We showed in Chapter 4 that the flame power predominantly modulates the strength of the thermoacoustic effect, while the convective time delay alters the thermoacoustic behaviour. We therefore choose to study flames with smaller increments in convective time delay, but equivalent power. The 15 flames all have a heat release rate of 1 kW in the inner cone. We choose to control the heat release rate of the inner cone because we observed that, in our rig, the outer cone does not respond to acoustic forcing, and therefore does not significantly contribute to the thermoacoustic oscillations.

The main flow and combustion characteristics of the 15 flames are illustrated in Fig. 5.3. The flames are equispaced on the 1 kW iso-contour, extending from the maximum flow rate that the mass flow controllers can deliver down to lean blow-off. The benefit of only studying rich flames is that we can study a wide range of equivalence ratios with minimal changes in injection velocity. We can therefore study a wide range of flame lengths and flame shapes with almost constant turbulence intensities. The main flow and turbulence properties of the 15 flames are summarized in Table 5.1. Further details are given in Table A.2 in Appendix A.




Fig. 5.3 Properties of the 15 flames studied in this chapter. The flames studied are presented as blue dots on axes of the air and fuel flow rate commands sent to the mass flow controllers. The injection velocity and equivalence ratio are overlaid as labelled iso-contours. The inner cone heat release rate is shown as coloured contours in the background, with the 1 kW iso-contour highlighted as a solid blue line.

Table 5.1 Summary of the properties of the 15 flames studied in this chapter.

Property	Units	Value / Range
Equivalence ratio	-	1.09-1.42
Injection velocity	m/s	4.58-5.00
Volumetric air flow rate	L_n/min	15.7-16.7
Volumetric fuel flow rate	L_n/min	0.72-1.00
Inner cone heat release rate	W	1000
Total heat release rate	W	1100-1600
Reynolds number	-	2800-3000
Turbulence intensity	-	5.94-6.02%

5.1.4 Thermoacoustic model of the ducted flame rig

The thermoacoustic network model is the same as that described in § 4.1.4. For convenience, we repeat the model diagram in Fig. 5.4. The only modification to the rig that may impact the modelling is the introduction of the turbulence grid, which could be a source of visco-thermal dissipation. This is modelled in the same manner as the cold electric heater in Chapter 3.




Fig. 5.4 Diagram of the acoustic network model used in this study. The unknown model parameters are: R_* , the reflection coefficients at the boundaries, η_* , the strengths of the visco-thermal damping, and \mathcal{F} , the transfer function from velocity perturbations to heat release rate fluctuations.

5.2 Inferring flame transfer functions of turbulent conical flames from pressure measurements

In this section we apply the Bayesian inference framework to infer the fluctuating heat release rate of the 15 flames from pressure measurements. As before, we begin by characterizing the sources of acoustic damping in the cold rig, following which we introduce the flames.

5.2.1 Characterizing the cold rig

In Chapter 4, we assimilated nine parameters describing the characteristics of the cold rig. We used six parameters for the three complex reflection coefficients. We used a further three parameters for the strength of the visco-thermal damping on (i) the inner wall of the duct, (ii) the outer wall of the burner and (iii) the inner wall of the burner. The posterior values of the three visco-thermal damping strengths were, however, quite similar, so in the current chapter we apply a single value for the strength of this damping to all three surfaces. This reduces the number of parameters used to describe the cold rig from nine to seven, which raises the overall marginal likelihood of the cold rig model.

We perform three sets of cold experiments to infer the seven unknown parameters, which we label C1-C3. In C1, we once again harmonically force the empty duct to infer R_u , R_d and η . We supply weak prior information for the reflection coefficients using analytical models for the reflection at flanged [163] and unflanged [131] duct terminations. We set a tight prior of $\eta = 1$, because the model makes assumptions which are reasonable for our rig, and so we expect it to be accurate.

In C2 we introduce the burner and traverse it through the rig while supplying a mass flow rate of air sufficient to choke the choke plate. We use the data to infer R_b and to update (i) the upstream reflection coefficient, including the disturbance of the burner, and (ii) the strength of the visco-thermal damping accounting for both the duct and burner walls. We supply the posteriors inferred from the C1 experiments as priors for the C2 experiments, but inflate the uncertainty in R_u and η to allow the parameters to update based on the new evidence. We supply a weak prior for R_b , using the theoretical value for a choked boundary.

In C3 we install the turbulence grid and repeat the C2 experiments. This was intended to inform a model for the damping of the turbulence grid, but we found that the turbulence grid had a negligible impact on the decay rate and natural frequency, suggesting that the damping is negligible. We had expected the damping to be small because the turbulence grid (i) is very thin (0.3 mm), (ii) has a large open area ratio (43%), and (iii) is placed near a velocity node in the burner.

The results of the cold rig characterization are shown in Fig. 5.5. As in Chapter 4, we see that the prior model is qualitatively accurate in each case, but not quantitatively accurate because of small errors in the parameter values. After we infer the parameters from the data, the model predictions match the data to within experimental uncertainty, and the uncertainty bounds on the posterior model predictions are small, meaning that we have high confidence in

the posterior parameter values. While the errors in the prior model predictions may be small, it is important that we can accurately model the cold rig. Any errors in the cold rig model will be incorporated into the flame transfer function in the next step, making the inferred flame transfer functions inaccurate.

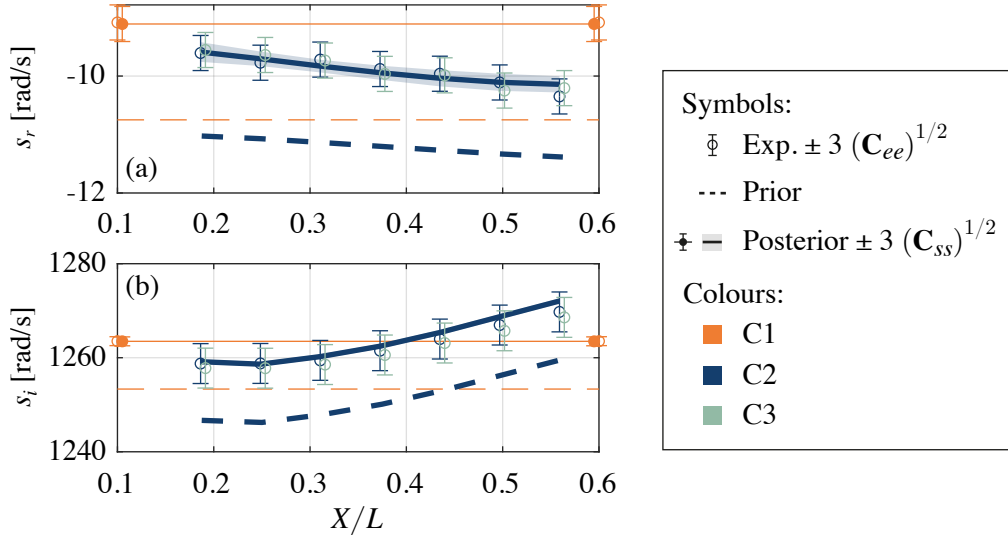


Fig. 5.5 Experimental measurements of (a) growth rate and (b) angular frequency plotted against burner exit location for the three sets of cold characterization experiments. Prior model predictions are plotted (dashed lines) without confidence bounds. Model predictions after data assimilation are plotted (solid lines) with a confidence bound of 2 standard deviations.

The prior and posterior joint parameter probability distributions are shown graphically in Fig. 5.6. Each set of axes plot the joint distribution between a pair of parameters. The discs represent regions of one, two and three standard deviations, centred around the expected value. As in Chapters 3 and 4, it was not possible to fully disentangle the values of $\text{Re}(R_u)$ and $\text{Re}(R_d)$, and the values of $\text{Im}(R_u)$ and $\text{Im}(R_d)$. The posterior uncertainties are, however, acceptably small, even with the uncertainty due to the correlation.

5.2.2 Assimilating heat release rate models from pressure data

With an accurate model of the cold acoustics, any changes in the system behaviour when the flame is introduced can be attributed to the flame transfer function. We assume that the cold rig parameters do not change when the flame is introduced, apart from R_d , which we expect to change with temperature. As before, we account for this by calculating a correction factor to the Levine-Schwinger model for the reflection coefficient [131]. When the flame is introduced, we use the corrected model to calculate the downstream reflection coefficient of the hot duct.




Fig. 5.6 Prior and posterior joint parameter probability distributions after assimilating data from the cold experiments. Each disc shows the joint distribution between a pair of parameters. The three rings represent one, two and three standard deviations, centred around the expected value. The upper and lower triangles show the same information zoomed to the prior 3 s.d. bound (lower triangle) and the posterior 3 s.d. bound (upper triangle).

We model the fluctuating heat release rate as a velocity-dependent source in the energy equation of the acoustic network model using the local linear feedback, k_{euf} . This is a complex number, which we infer from data. As a reminder, the feedback strength is related to the typical flame transfer function by:

$$\mathcal{F} = \frac{Q'/\bar{Q}}{u'/\bar{u}} = \frac{\gamma}{\gamma-1} \frac{\bar{p}\bar{u}}{\bar{Q}} k_{euf} \quad (5.1)$$

where \mathcal{F} is the complex-valued flame transfer function, which relates fluctuations in velocity, u' , to fluctuations in heat release rate, Q' . The fluctuations in velocity and heat release rate are normalized by the mean bulk values, \bar{u} and \bar{Q} . γ is the ratio of specific heats, and \bar{p} is the mean pressure at the injection plane.

We infer the parameters of the fluctuating heat release rate model for each of the 15 flames individually, using only observations of the growth (or decay) rate and the natural frequency of oscillations. We repeat this with the flames in two axial positions within the duct: $x/L = 0.25$ and $x/L = 0.5$, where x is the axial position along the duct measured from the upstream end, and L is the length of the duct. For each of the flames, at each flame position, we infer two parameters, which are the real and imaginary parts of k_{eu_f} . We use real-imaginary form, because it leads to a well-posed optimization problem with a single global optimum.

We compare the posterior model predictions to the experimental observations in Fig. 5.7, which plots the growth rate and natural frequency of oscillations against Strouhal number. We have defined the Strouhal number for the turbulent flames as $St = s_i L_f / \bar{u}$, where s_i is the angular frequency, L_f is the length of the unperturbed flame, and \bar{u} is the mean injection velocity. Although each flame was forced at the natural frequency of the rig, they each have different lengths and injection velocities, so we obtain results over a wide range of Strouhal numbers.




Fig. 5.7 Comparison of experimental (circles and error bars) and predicted (line and patch) values of (a) growth rate and (b) natural frequency of thermoacoustic oscillations plotted against Strouhal number, where the length scale is the length of the unperturbed flame and the velocity scale is the mean injection velocity. The results for two burner positions are shown: $x/L = 0.25$ in blue and $x/L = 0.5$ in orange.

After inferring the most probable fluctuating heat release rate for each flame, the model predicts the experimental observations exactly. However, we note that in this chapter we

assimilate the fluctuating heat release rate for each flame, at each burner position individually. This gives the model freedom to fit *any* data. Any errors in the data or the cold model would therefore be incorporated into the fluctuating heat release rate parameters in order to make the model fit the data.

We convert the feedback strength, k_{eu_f} , into the more familiar form of a flame transfer function, using Eq. (5.1). We emphasize, however, that for each flame we only obtain the flame transfer function at a single frequency (the observed frequency). The inferred flame transfer functions for the 15 flames at two burner positions are plotted against Strouhal number in Fig. 5.8. The Bode plot in Fig. 5.8 qualitatively resembles the typical flame transfer function for laminar conical flames (see Fig. 4.15, or Refs. [140, 166, 167]). We also note that the flame transfer functions inferred from the two burner positions are similar, which we should expect because the flame transfer function should not depend on burner position. We do not expect them to be identical, however, because it was not possible to keep the forcing amplitude, u'/\bar{u} consistent between the two tests. This was not a problem in chapter 4, because we always maintained the forcing amplitude within the linear regime. In this chapter, however, we increased the forcing amplitude to maximize the signal-to-noise ratio of the data.




Fig. 5.8 Flame transfer functions inferred from pressure data. (a) Gain and (b) phase delay are plotted against Strouhal number for 15 flames at two burner positions: $x/L = 0.25$ (blue) and $x/L = 0.5$ (orange). The errorbars denote a region of 2 standard deviations from the expected value.

An alternate representation of the variation of flame transfer function with the flame properties is given in Fig. 5.9, which plots the flame transfer functions on polar axes. For clarity, only the flame transfer functions for the flames at $x/L = 0.25$ are shown, because the results for $x/L = 0.5$ are similar.




Fig. 5.9 Flame transfer functions for the 15 flames at $x/L = 0.25$ plotted on polar axes, with the gain plotted on the radial axis and the phase delay on the angular axis. The patches denote a region of 1 standard deviation. The contour in the background represents the effect of the gain and phase of the flame transfer function on the growth rate of oscillations, where red represents increased growth rates, blue represents reduced growth rates and white represents no change in growth rates.

In Chapter 4 we noted that, for laminar conical flames, the uncertainty of the inferred flame transfer functions was largest for (i) neutral flames, and (ii) strongly damping flames. We attributed (i) to the fact that neutral flames produce a weak thermoacoustic effect, so it is difficult to infer the flame transfer function from observations of the thermoacoustic effect alone, and we attributed (ii) to the fact that strongly damping flames produce quickly decaying oscillations, which increases the experimental uncertainty in the growth rate and natural frequency. For turbulent flames, however, we see that the uncertainty in the flame transfer function is dominated by experimental uncertainty. As was noted at the beginning of this chapter, the experimental uncertainty for turbulent flames is larger because (i) the turbulence produces additional broadband noise, which reduces the signal-to-noise ratio of the data, and (ii) the turbulence perturbs the flame shape, causing variation in the thermoacoustic effect,

which introduces additional random error into the experimental data. We still see, however, that we obtain precise estimates of \mathcal{F} for the driving flames, which are of the most interest in industrial systems.

Finally, we see from both Figs. 5.8 and 5.9 that the uncertainty in $\angle\mathcal{F}$ is much lower than the uncertainty in $|\mathcal{F}|$. This is because the predictions of the thermoacoustic model are most sensitive to the phase delay [109]. In the forward modelling problem this is typically a challenge, because small errors in $\angle\mathcal{F}$ can cause large errors in the model predictions. In the inference problem, however, this is beneficial because it provides a more precise posterior parameter estimation.

5.2.3 Validation of inferred fluctuating heat release rate

We validate the inferred quantities by comparing the posterior fluctuating heat release rate against that measured directly using the high speed camera. The validation results are shown in Fig. 5.10. We plot the magnitude and phase of the relative fluctuating heat release rate, $\hat{Q} = Q'/\bar{Q}$, against Strouhal number, St , for the 15 flames at both burner positions.

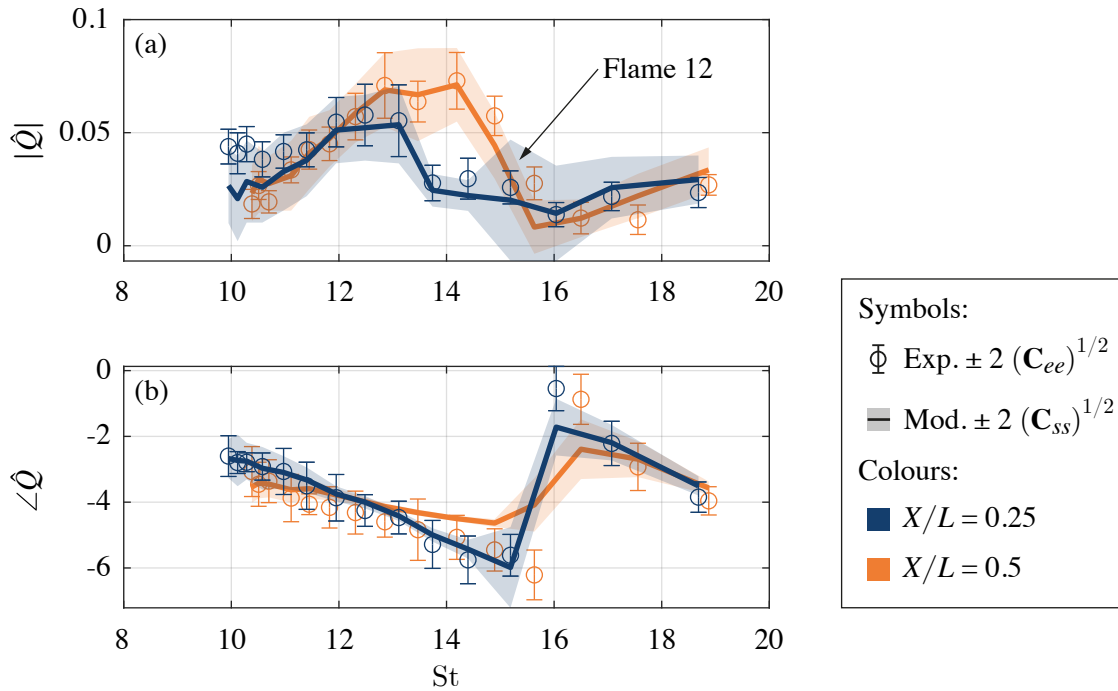


Fig. 5.10 (a) Magnitude of the relative heat release rate fluctuations, $|\hat{Q}|$, (b) the phase of the relative heat release rate fluctuations, $\angle\hat{Q}$ as functions of Strouhal number, St . Comparison of the direct measurements (circles with error bars) with the inferred quantities (lines and patches) for the flame positions $x/L = 0.25$ (blue) and $x/L = 0.5$ (orange).

We see from Fig. 5.10(a) that the magnitude of the inferred fluctuating heat release rate compares well with the direct measurements for all 15 flames at both burner positions. For the burner position $x/L = 0.25$, the data includes both linearly stable and self excited flames. The system was unstable for Strouhal numbers near 14, where the inferred and measured values for $|\hat{Q}|$ is substantially lower than when the burner was at $x/L = 0.5$. This is because the acoustic amplitude reached during the limit cycle was lower than the forcing amplitude, producing smaller heat release rate fluctuations.

For both burner positions, the data covers a wide range of forcing amplitudes between 6% and 25%. This covers amplitudes that typically result in both linear and nonlinear thermoacoustic responses, indicating that this framework could be used to infer flame describing functions, if data at multiple forcing frequencies and amplitudes were available.

The uncertainty in the inferred heat release rate magnitude is generally only slightly larger than the uncertainty of the direct measurements, except for flame 12, where the uncertainty in the inferred heat release rate is much larger than the direct measurement. The large uncertainty in flame 12 is, however, due to larger experimental uncertainty. When the burner is at $x/L = 0.25$, flame 12 produces a growth rate near zero (see Fig. 5.7). The system alternates between linearly stable and self excited, which causes a large variation in the experimental results. This was not captured by the direct measurements, because we only recorded three of the 100 experiments, all of which happened to capture similar behaviour. The similar uncertainties in the inferred and measured heat release rates suggests that the increased random error introduced by turbulence affects the direct measurements and inferred quantities to a similar extent.

Fig. 5.10(b) shows the phase of the fluctuating heat release rate, measured relative to the pressure at the reference microphone. We see that the inferred quantities compare well with the direct measurements, except for when the magnitude of the heat release rate fluctuations are small. In this case we can expect larger errors in both the inferred quantities and direct measurements, because it is generally challenging to identify the phase of a low amplitude, noisy signal.

Unlike the magnitude, the uncertainty in the inferred phase is smaller than that in the direct measurements. This is because (i) the direct measurement of the relative phase between two noisy signals is prone to error, and (ii) the sensitivity of the model to the phase delay results in a more precise inferred quantity. The uncertainty in the inferred heat release rate phase does, however, increase when the heat release rate magnitude is small, as we would expect.

5.3 Conclusion

In this chapter, we further increase the complexity of the system by studying turbulent conical flames. The results of Chapter 4 revealed that the uncertainty in the inferred flame transfer function grows with experimental uncertainty. Turbulence introduces large random uncertainty through (i) additional broadband noise, and (ii) stochastic variation in the thermoacoustic behaviour. It is therefore important to test how the framework handles turbulent flames, as they are most prevalent in industry.

We show that, in spite of the larger random error introduced by turbulent fluctuations, the Bayesian framework still performs well. We are still able to converge to the most probable parameters and generate quantitative agreement between the model and the data. As expected, we find that the uncertainty in the inferred flame transfer functions increases as a result of the larger experimental uncertainty. However, we still obtain acceptably precise estimates of the flame transfer functions for driving flames. These are the most dangerous, so it is useful that we are able to characterize them with the highest precision.

The set of flames studied in this chapter include both linearly stable and self excited flames. Each linearly stable flame was forced at the peak amplitude that could be driven by the loudspeaker, which produced a range of amplitudes that typically result in both linear and nonlinear thermoacoustic responses. We are able to infer the flame transfer function in each of these cases, demonstrating that the method is not limited to simple cases like stable flames with small amplitude perturbations.

We directly measure the fluctuating heat release rate, and use the direct measurements to validate the inferred quantities. We show that the expected values of the inferred fluctuating heat release rates compare well to the direct measurements for all 15 flames. Further, we see that the direct measurements themselves suffer large random uncertainty due to the stochastic variation introduced by turbulence. As a result we find, unexpectedly, that the inferred fluctuating heat release rate magnitudes have similar precision to the direct measurement, while the phase delays are inferred with more precision than the direct measurement.

PART III:

CONCLUSION

Chapter 6

Adjoint-accelerated Bayesian inference in thermoacoustics

6.1 Summary of main results

In Chapter 3 we use the hot wire Rijke tube to demonstrate the full set of tools that the Bayesian framework provides. We first perform automated experiments to collect a dataset for training. With this dataset, we use Bayesian parameter inference, uncertainty quantification and model comparison to construct a quantitatively accurate model of the rig. We then use Bayesian optimal experiment design to assess how much data is strictly required for several inference tasks.

In § 3.2 we apply our adjoint-accelerated Bayesian inference framework to generate a quantitatively accurate model of the hot wire Rijke tube. The final assembled model contains the fewest possible parameters, is quantitatively accurate over the full operating range, and is physically interpretable. Further, the model extrapolates successfully because it is physics-based. This is a significant improvement on other attempts in the literature e.g. [123].

Following the Bayesian framework forced us to critically evaluate several assumptions typically made in the literature. We find that many of these assumptions are not supported by the data. This reveals one of the most powerful features of the Bayesian framework: the researcher is forced to carefully consider the underlying physics. At each step, the quality of the assumptions we make are rigorously evaluated with quantitative measures.

In § 3.3 we revisit the densely sampled dataset used to train the model in § 3.2, and use it to demonstrate Bayesian optimal experiment design. We show that by using Bayesian optimal

experiment design we can significantly reduce the experimental cost by (i) performing only a few, highly informative experiments, and (ii) using fewer, carefully placed sensors. This is an important result because, while our rig is cheap to operate, it can be prohibitively expensive to collect data on industrial rigs. There is therefore a large financial incentive to minimize the number of experiments conducted on these rigs.

In Chapter 4 we introduce additional complexity to the system by replacing the electric heater with a laminar conical flame. For the cold rig, this introduces the challenge of generating a quantitatively accurate model of a system with coupled acoustic ducts. For the hot rig it introduces the complexity of richer heat release rate dynamics, and more random variation in the experimental data. We once again perform automated experiments to collect the data, which we assimilate into physics-based models of the thermoacoustic system.

Using Bayesian inference, we infer the flame transfer functions of 24 flames *in-situ* from pressure measurements alone. If the flame's response is sensitive to its environment, which it often is, this is preferable to measuring the flame transfer functions *ex-situ* using optical methods.

We verify the inferred flame transfer functions by comparing them to direct measurements from similar rigs in the literature. The direct measurements themselves were somewhat inconsistent for the gain, and totally inconsistent for the phase delay. Nonetheless, we show that the inferred quantities are generally within the expected range, and any significant discrepancies are easily explained.

We rigorously quantify the uncertainties in the inferred flame transfer functions and find, as expected, that the flame transfer functions are most precise if (i) the thermoacoustic effect is strong, and (ii) the measurement uncertainty is small. This can help to guide future experiments on industrial rigs.

As in Chapter 3, we evaluate modelling assumptions that have been made in other studies in the literature. We again find that the data does not support assumptions that have been commonly used in other studies.

In Chapter 5 we test our inference framework on turbulent conical flames to see how the framework handles the large random uncertainty introduced by turbulence through (i) additional broadband noise, and (ii) stochastic variation in the thermoacoustic behaviour. We show that the Bayesian framework still succeeds in the presence of the additional random experimental

error introduced by turbulent fluctuations. While the larger error increases the uncertainty in the inferred flame transfer functions, we still obtain acceptably precise estimates for driving flames. These are the most dangerous, so it is useful that we are able to characterize them most precisely.

We improve on the validation attempt made in Chapter 4 by directly measuring the fluctuating heat release rate of each flame. We show that the inferred quantities compare well to the direct measurements for all 15 flames. Further, we see that the direct measurements themselves suffer large random uncertainty due to the stochastic variation introduced by turbulence. As a result we find that, unexpectedly, the fluctuating heat release rate magnitudes are inferred with similar precision to the direct measurement, while the phase delays are inferred with more precision than the direct measurement.


6.2 Outlook

Overall, this thesis demonstrates that the Bayesian framework provides a powerful set of tools for combining experimental data with numerical models. Within this framework, the often separate silos of experimentation and modelling become strongly mutually beneficial. The modeller benefits from the work of the experimentalist by using experimental data to improve the accuracy of physics-based models. The experimentalist benefits from the work of the modeller by using the model to design optimal experiments, significantly reducing the experimental effort and cost.

Unlike many other machine learning frameworks, our framework produces data-driven models that are robust, interpretable, and capable of extrapolating beyond the observed dataset. It is therefore a natural choice for data-driven solutions in engineering and physics. While this framework is most readily applied to reduced order models, it extends naturally to more complex problems such as CFD simulations. Within our group it has been used to learn the boundary conditions of a Navier–Stokes problem from noisy MRV data of cardio-vascular flows [173], and to infer the parameters of a RANS turbulence model from LES data for a turbulent jet flame [174]. The framework is general and can be applied to a wide range of problems, provided you have (i) a source of data, (ii) a physics-based model, and (iii) the derivative of the model with respect to the parameters, e.g. with an adjoint code. Adjoint or automatically-differentiated codes are readily available in the field of fluid mechanics because they have been widely used for shape optimization. This work is a further example of the utility of adjoint codes. Automatically differentiated codes can be similarly useful, but care must be

taken when coding the forward problem to ensure the code is differentiable. Careful application of the adjoint approach always produces a well-posed differentiable code.

In the work on the hot wire Rijke tube, we were able to construct a general model of the system because it was possible to propose a simple model for the heater dynamics. The dynamics of even the simplest flames, however, require detailed modelling that is beyond the scope of this thesis. This has been the subject of work carried out in parallel, where we have assimilated data from video footage of perturbed laminar flames into a G-equation model of the flame [175]. If this work is carried further, it will be possible to construct a quantitatively accurate model of a burner that is capable of predicting the flame transfer function at an arbitrary operating condition and excitation frequency. This, combined with the acoustic network models presented in this thesis, would form a powerful design tool that could be used to inform more elegant and robust interventions to thermoacoustic instability with fewer prototyping iterations.

For more complex flames, such as the turbulent swirl flames found in gas turbine combustors, a simple G-equation model will not be suitable. For these cases it would be most productive to assimilate data from experiments or LES into RANS simulations to learn the most probable parameters of the turbulence and combustion models [174]. It is, as yet, unknown if a RANS model captures enough of the physics to be capable of becoming quantitatively accurate. The predictions could, nevertheless, be improved through the application of Bayesian inference, while simultaneously introducing a computationally efficient method for uncertainty quantification to a field that could greatly benefit from it. 

It may also be productive to investigate the use of the distributed time delay [40] as a simple model for the flame response. With this model, a few parameters (the discrete time delays) can be used to reconstruct the full frequency response of the flame. We could propose prior probabilities for several time delays based on our understanding of the physical phenomena in the rig, e.g. vortex shedding from an obstruction upstream of the flame, or equivalence ratio fluctuations in the injector. We could then assimilate data to get more precise estimates of the time delays, and use systematic uncertainty quantification and model comparison to check whether we have captured all the important time delays. The resulting model should be capable of predicting the system behaviour at unobserved frequencies, and perhaps even nearby operating conditions, provided the flow-physics doesn't change significantly. The model could then be used to find optimal interventions to thermoacoustic instability, provided these interventions don't change the time delays, thereby invalidating the flame model. As

an example, it could possibly be used to identify points of maximum receptivity to a passive control system [176, § 3.2.2].

Even without detailed flame modelling, we believe that the results of Chapters 4 and 5 can be immediately beneficial to studies on industrial test rigs. This would improve on the experimental methods currently in use in one of two ways. Where flame transfer functions were previously obtained *ex-situ* through optical means, they can now be obtained *in-situ*. Where they were obtained *in-situ* using the multi-source method, our method can quantify the uncertainty in the inferred quantities, which is a valuable input for decision-making. Work is already underway on applying this framework to the Rolls–Royce SCARLET industrial test rig [177].

When applying the framework to industrial systems, optimal experiment design will become invaluable due to the extreme cost of data collection. Because of the low computational cost of this framework, it would be possible to develop an automated experiment system that recursively (i) uses the model to identify an optimal experiment, (ii) conducts the optimal experiment, and (iii) assimilates the data into the model to update the parameter predictions. This would significantly reduce the effort and cost of experimentation, making it feasible to apply Bayesian inference in an industrial setting.

References

- [1] B Higgins. On the Sound Produced by a Current of Hydrogen Gas Passing Through a Tube. *Journal of Natural Philosophy Chemistry and the Arts*, 1, 1802.
- [2] C Sondhauss. Über die Schallschwingungen der Luft in erhitzten Glasröhren und in gedeckten Pfeifen von ungleicher Weite. *Ann. Phys.*, 155:1–34, 1850.
- [3] P. L. Rijke. Notice of a new method of causing a vibration of the air contained in a tube open at both ends. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 17(116):419–422, 1859. doi: 10.1080/14786445908642701.
- [4] J.W. Rayleigh. The explanation of certain acoustical phenomena. *Nature*, 1878.
- [5] F Culick. Unsteady motions in combustion chambers for propulsion systems. Technical report, Res. Technol. Organ., North Atl. Treaty Organ., Washington, DCs, 2006.
- [6] Ying Huang and Vigor Yang. Dynamics and stability of lean-premixed swirl-stabilized combustion. *Progress in Energy and Combustion Science*, 35(4):293–364, 2009. doi: 10.1016/j.pecs.2009.01.002.
- [7] Timothy C. Lieuwen and Vigor Yang. *Combustion Instabilities In Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modeling*. American Institute of Aeronautics and Astronautics, 2006.
- [8] T. Poinso. Prediction and control of combustion instabilities in real engines. *Proceedings of the Combustion Institute*, 36(1):1–28, 2017. doi: 10.1016/j.proci.2016.05.007.
- [9] H. C. Mongia, T. J. Held, G. C. Hsiao, and R. P. Pandalai. Challenges and Progress in Controlling Dynamics in Gas Turbine Combustors. *Journal of Propulsion and Power*, 19(5):822–829, 2003. doi: 10.2514/2.6197.
- [10] Tim Lieuwen and Keith McManus. Introduction: Combustion Dynamics in Lean-Premixed Pre-vaporized (LPP) Gas Turbines. *Journal of Propulsion and Power*, 19(5):721, 2003. doi: 10.2514/2.6171.
- [11] Joseph C. Oefelein and Yang Vigor. Comprehensive review of liquid-propellant combustion instabilities in f-1 engines. *Journal of Propulsion and Power*, 9(5):657–677, 1993. doi: 10.2514/3.23674.
- [12] Maria A. Heckl. Active control of the noise from a rijke tube. *Journal of Sound and Vibration*, 124(1):117–133, 1988. doi: 10.1016/S0022-460X(88)81408-1.

- [13] Ann P. Dowling and Aimee S. Morgans. Feedback control of combustion oscillations. *Annual Review of Fluid Mechanics*, 37(Putnam 1971):151–182, 2005. doi: 10.1146/annurev.fluid.36.050802.122038.
- [14] J. R. Seume, N. Vortmeyer, W Krause, et al. Application of active combustion instability control to a heavy duty gas turbine. *ASME Journal of Engineering for Gas Turbines and Power*, 120(4), 1998.
- [15] J.M. Cohen, N.M. Rey, C.A. Jacobson, and T.J. Anderson. Active control of combustion instability in a liquid-fueled low-NO_x combustor. *Journal of Engineering for Gas Turbines and Power*, 121(2), 1999.
- [16] Bayu Dharmaputra, Sergey Shcherbanev, Bruno Schuermans, and Nicolas Noiray. Thermoacoustic stabilization of a sequential combustor with ultra-low-power nanosecond repetitively pulsed discharges. *Combustion and Flame*, 258:113101, 2023. doi: 10.1016/j.combustflame.2023.113101.
- [17] Ammar M. Alkhalifa, Abdulrahman Alsalem, Davide Del Cont-Bernard, and Deanna A. Lacoste. Active control of thermoacoustic fluctuations by nanosecond repetitively pulsed glow discharges. *Proceedings of the Combustion Institute*, 39(4):5429–5437, 2023. doi: 10.1016/j.proci.2022.06.013.
- [18] Deanna A. Lacoste. Flames with plasmas. *Proceedings of the Combustion Institute*, 39(4):5405–5428, 2023. doi: 10.1016/j.proci.2022.06.025.
- [19] Matthew P Juniper and R.I Sujith. Sensitivity and Nonlinearity of Thermoacoustic Oscillations. *Annual Review of Fluid Mechanics*, 50:661–689, 2018.
- [20] José G. Aguilar and Matthew P. Juniper. Thermoacoustic stabilization of a longitudinal combustor using adjoint methods. *Physical Review Fluids*, 5(8):1–26, 2020. doi: 10.1103/PhysRevFluids.5.083902.
- [21] D. Noh, E. Karlis, S. Navarro-Martinez, et al. Azimuthally-driven subharmonic thermoacoustic instabilities in a swirl-stabilised combustor. *Proceedings of the Combustion Institute*, 37(4):5333–5341, 2019. doi: 10.1016/j.proci.2018.07.090.
- [22] Guillaume Vignat, Ermanno Lo Schiavo, Davide Laera, et al. Dynamics of spray and swirling flame under acoustic oscillations: A joint experimental and LES investigation. *Proceedings of the Combustion Institute*, 38(4):6015–6024, 2021. doi: 10.1016/j.proci.2020.05.054.
- [23] E. Lo Schiavo, D. Laera, E. Riber, et al. On the impact of fuel injection angle in Euler–Lagrange large eddy simulations of swirling spray flames exhibiting thermoacoustic instabilities. *Combustion and Flame*, 227:359–370, 2021. doi: 10.1016/j.combustflame.2021.01.009.
- [24] O. Schulz, U. Doll, D. Ebi, et al. Thermoacoustic instability in a sequential combustor: Large eddy simulation and experiments. *Proceedings of the Combustion Institute*, 37(4): 5325–5332, 2019. doi: 10.1016/j.proci.2018.07.089.

- [25] Daniel Fredrich, William P. Jones, and Andrew J. Marquis. Thermo-acoustic Instabilities in the PRECCINSTA Combustor Investigated Using a Compressible LES-pdf Approach. *Flow, Turbulence and Combustion*, 106(4):1399–1415, 2021. doi: 10.1007/s10494-020-00177-3.
- [26] Zhi X. Chen, Nedunchezian Swaminathan, Marek Mazur, et al. Numerical investigation of azimuthal thermoacoustic instability in a gas turbine model combustor. *Fuel*, 339 (December 2022), 2023. doi: 10.1016/j.fuel.2023.127405.
- [27] Shigeru Tachibana, Kinya Saito, Takeshi Yamamoto, et al. Experimental and numerical investigation of thermo-acoustic instability in a liquid-fuel aero-engine combustor at elevated pressure: Validity of large-eddy simulation of spray combustion. *Combustion and Flame*, 162(6):2621–2637, 2015. doi: 10.1016/j.combustflame.2015.03.014.
- [28] G. Staffelbach, L. Y.M. Gicquel, G. Boudier, and T. Poinsot. Large Eddy Simulation of self excited azimuthal modes in annular combustors. *Proceedings of the Combustion Institute*, 32 II(2):2909–2916, 2009. doi: 10.1016/j.proci.2008.05.033.
- [29] Pierre Wolf, Gabriel Staffelbach, Laurent Y.M. Gicquel, et al. Acoustic and Large Eddy Simulation studies of azimuthal modes in annular combustion chambers. *Combustion and Flame*, 159(11):3398–3413, 2012. doi: 10.1016/j.combustflame.2012.06.016.
- [30] A Giusti and E Mastorakos. Turbulent Combustion Modelling and Experiments: Recent Trends and Developments. *Flow, Turbulence and Combustion*, 103:847–869, 2019.
- [31] Ann P. Dowling and Simon R. Stow. Acoustic Analysis of Gas Turbine Combustors. *Journal of Propulsion and Power*, 19(5):751–764, 2003. doi: 10.2514/2.6192.
- [32] Simon R Stow and Ann P. Dowling. Thermoacoustic oscillations in an annular combustor. In *Proceedings of the ASME Turbo Expo*. ASME, 2001.
- [33] Luigi Crocco and Sin-i Cheng. Theory of combustion instability in liquid propellant rocket motors. Technical report, Princeton Univ Nj, 1956.
- [34] H. J. Merk. An analysis of unstable combustion of premixed gas. In *Symposium (International) on Combustion*, pages 500–512, 1957. doi: 10.1016/S0082-0784(57)80067-8.
- [35] A P Dowling. Nonlinear self-excited oscillations of a ducted flame. *Journal of Fluid Mechanics*, 346:271–290, 1997.
- [36] A. P. Dowling. A kinematic model of a ducted flame. *Journal of Fluid Mechanics*, 394: 51–72, 1999. doi: 10.1017/S0022112099005686.
- [37] D. Durox, T. Schuller, N. Noiray, and S. Candel. Experimental analysis of nonlinear flame transfer functions for different flame geometries. *Proceedings of the Combustion Institute*, 32 I(1):1391–1398, 2009. doi: 10.1016/j.proci.2008.06.204.
- [38] Bruno B.H. Schuermans, Wolfgang Polifke, and Christian Oliver Paschereit. Modeling transfer matrices of premixed flames and comparison with experimental results. *Proceedings of the ASME Turbo Expo*, 2, 1999. doi: 10.1115/99-GT-132.

- [39] W. Polifke, A. Poncet, C. O. Paschereit, and K. Döbbeling. Reconstruction of acoustic transfer matrices by instationary computational fluid dynamics. *Journal of Sound and Vibration*, 245(3):483–510, 2001. doi: 10.1006/jsvi.2001.3594.
- [40] Wolfgang Polifke. Modeling and analysis of premixed flame dynamics by means of distributed time delays. *Progress in Energy and Combustion Science*, 79, 2020. doi: 10.1016/j.pecs.2020.100845.
- [41] Sébastien Ducruix, Daniel Durox, and Sébastien Candel. Theoretical and experimental determination of the flame transfer function of a laminar premixed flame. *Proceedings of the Combustion Institute*, 28:765–773, 2000.
- [42] Bruno Schuermans, Felix Guethe, and Wolfgang Mohr. Optical transfer function measurements for technically premixed flames. *Journal of Engineering for Gas Turbines and Power*, 132(8):1–8, 2010. doi: 10.1115/1.3124663.
- [43] Eirik Æsøy, Thomas Indlekofer, Francesco Gant, et al. The effect of hydrogen enrichment, flame-flame interaction, confinement, and asymmetry on the acoustic response of a model can combustor. *Combustion and Flame*, 242, 2022. doi: 10.1016/j.combustflame.2022.112176.
- [44] R Kaess, T Poinso, and W Polifke. Determination of the stability map of a premix burner based on flame transfer functions computed with transient CFD. *Proceedings of the European Combustion Meeting*, (January):1–6, 2009.
- [45] Xingsi Han, Jingxuan Li, and Aimee S. Morgans. Prediction of combustion instability limit cycle oscillations by combining flame describing function simulations with a thermoacoustic network model. *Combustion and Flame*, 162(10):3632–3647, 2015. doi: 10.1016/j.combustflame.2015.06.020.
- [46] Fabien Dupuy, Marco Gatti, Clément Mirat, et al. Combining analytical models and LES data to determine the transfer function from swirled premixed flames. *Combustion and Flame*, 217:222–236, 2020. doi: 10.1016/j.combustflame.2020.03.026.
- [47] M. Gatti, R. Gaudron, C. Mirat, et al. A comparison of the transfer functions and flow fields of flames with increasing swirl number. *Proceedings of the ASME Turbo Expo*, 4B-2018:1–12, 2018. doi: 10.1115/GT2018-76105.
- [48] Håkon T. Nygård and Nicholas A. Worth. Flame transfer functions and dynamics of a closely confined premixed bluff body stabilized flame with swirl. *Journal of Engineering for Gas Turbines and Power*, 143(4):1–10, 2021. doi: 10.1115/1.4049513.
- [49] A. Cuquel, D. Durox, and T. Schuller. Scaling the flame transfer function of confined premixed conical flames. *Proceedings of the Combustion Institute*, 34(1):1007–1014, 2013. doi: 10.1016/j.proci.2012.06.056.
- [50] Luis Tay-Wo-Chong and Wolfgang Polifke. Large eddy simulation-based study of the influence of thermal boundary condition and combustor confinement on premix flame transfer functions. *Journal of Engineering for Gas Turbines and Power*, 135(2):1–9, 2013. doi: 10.1115/1.4007734.

- [51] Chandrachur Bhattacharya, Jacqueline O'Connor, and Asok Ray. Data-driven Detection and Early Prediction of Thermoacoustic Instability in a Multi-nozzle Combustor. *Combustion Science and Technology*, 00(00):1–32, 2020. doi: 10.1080/00102202.2020.1820495.
- [52] Ushnish Sengupta, Carl E. Rasmussen, and Matthew P. Juniper. Bayesian machine learning for the prognosis of combustion instabilities from noise. *Proceedings of the ASME Turbo Expo*, 4A-2020:1–10, 2020. doi: 10.1115/GT2020-14904.
- [53] Sudeepta Mondal, Najah F. Ghalyan, Asok Ray, and Achintya Mukhopadhyay. Early Detection of Thermoacoustic Instabilities Using Hidden Markov Models. *Combustion Science and Technology*, 191(8):1309–1336, 2019. doi: 10.1080/00102202.2018.1523900.
- [54] 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):1, 2019. doi: 10.1103/PhysRevApplied.11.064034.
- [55] Tatsuya Hashimoto, Hajime Shibuya, Hiroshi Gotoda, et al. Spatiotemporal dynamics and early detection of thermoacoustic combustion instability in a model rocket combustor. *Physical Review E*, 99(3):1–7, 2019. doi: 10.1103/PhysRevE.99.032208.
- [56] Marcin Rywik, Axel Zimmermann, Alexander J Eder, et al. Spatially resolved modeling of the nonlinear dynamics of a laminar premixed flame with a multilayer perceptron - convolution autoencoder network. *Proceedings of the ASME Turbo Expo*, pages 1–12, 2023.
- [57] Pritthi Chattopadhyay, Sudeepta Mondal, Asok Ray, and Achintya Mukhopadhyay. Dynamic Data-Driven Combustor Design for Mitigation of Thermoacoustic Instabilities. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 141(1):1–7, 2019. doi: 10.1115/1.4040210.
- [58] Taraneh Sayadi and Peter J. Schmid. Frequency response analysis of a (non-)reactive jet in crossflow. *Journal of Fluid Mechanics*, 922:1–21, 2021. doi: 10.1017/jfm.2021.479.
- [59] Jayesh Dhadphale, Vishnu R. Unni, Abhishek Saha, and R. I. Sujith. Neural ODE to model and prognose thermoacoustic instability. pages 1–31, 2021.
- [60] P.S. Laplace. *Theorie Analytique des Probabilites*. Courcier Imprimeur, 1812.
- [61] Stephen M. Stigler. Laplace's 1774 memoir on inverse probability. *Statistical Science*, 1(3):359–363, 1986. doi: 10.1214/ss/1177013620.
- [62] Udo Von Toussaint. Bayesian inference in physics. *Reviews of Modern Physics*, 83(3):943–999, 2011. doi: 10.1103/RevModPhys.83.943.
- [63] Harold Jeffreys. *Scientific Inference*. Cambridge University Press, 3 edition, 1973.
- [64] R. T. Cox. *Probability, frequency and rational expectation*, 1946.
- [65] E.T. Jaynes. Bayesian Methods: General Background. In *Maximum Entropy and Bayesian Methods in Applied Statistics*, pages 1–25. Cambridge University Press, 1984. doi: 10.1017/cbo9780511569678.003.

- [66] Stephen F. Gull. *Bayesian Inductive Inference and Maximum Entropy*, 1988.
- [67] David J C MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- [68] C. R. Jenkins and J. A. Peacock. The power of Bayesian evidence in astronomy. *Monthly Notices of the Royal Astronomical Society*, 413(4):2895–2905, 2011. doi: 10.1111/j.1365-2966.2011.18361.x.
- [69] Eric Thrane and Colm Talbot. An introduction to Bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models. *Publications of the Astronomical Society of Australia*, 36, 2019. doi: 10.1017/pasa.2019.2.
- [70] J. Antoniadis, P. Arumugam, S. Arumugam, et al. The second data release from the European Pulsar Timing Array: III. Search for gravitational wave signals. *Astronomy and Astrophysics*, 678, 2023. doi: 10.1051/0004-6361/202346844.
- [71] Gabriella Agazie, Akash Anumalapudi, Anne M. Archibald, et al. The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background. *The Astrophysical Journal Letters*, 951(1):L8, 2023. doi: 10.3847/2041-8213/acdac6.
- [72] J. P. Huelsenbeck, F. Ronquist, R. Nielsen, and J. P. Bollback. Bayesian inference of phylogeny and its impact on evolutionary biology. *Science*, 294(5550):2310–2314, 2001. doi: 10.1126/science.1065889.
- [73] Darren J. Wilkinson. Bayesian methods in bioinformatics and computational systems biology. *Briefings in Bioinformatics*, 8(2):109–116, 2007. doi: 10.1093/bib/bbm007.
- [74] Rajesh Chowdhary, Jinfeng Zhang, and Jun S. Liu. Bayesian inference of protein-protein interactions from biological literature. *Bioinformatics*, 25(12):1536–1542, 2009. doi: 10.1093/bioinformatics/btp245.
- [75] Campbell R. Harvey and Guofu Zhou. Bayesian inference in asset pricing tests. *Journal of Financial Economics*, 26(2):221–254, 1990. doi: 10.1016/0304-405X(90)90004-J.
- [76] Thomas Flury and Neil Shephard. Bayesian inference based only on simulated likelihood: Particle filter analysis of dynamic economic models. *Econometric Theory*, 27(5):933–956, 2011. doi: 10.1017/S0266466610000599.
- [77] Thomas Lux. Approximate Bayesian inference for agent-based models in economics: a case study. *Studies in Nonlinear Dynamics and Econometrics*, 27(4):423–447, 2023. doi: 10.1515/snde-2021-0052.
- [78] E.S. Epstein. *Statistical inference and prediction in climatology: A Bayesian approach*. Springer, 2016.
- [79] Y. Wang, T. Maeda, K. Satake, et al. Tsunami Data Assimilation Without a Dense Observation Network. *Geophysical Research Letters*, 46(4):2045–2053, 2019. doi: 10.1029/2018GL080930.
- [80] Ian T. Nabney, Dan Cornford, and Christopher K.I. Williams. Bayesian inference for wind field retrieval. *Neurocomputing*, 30(1-4):3–11, 2000. doi: 10.1016/S0925-2312(99)00136-8.

- [81] Tobin Isaac, Noemi Petra, Georg Stadler, and Omar Ghattas. Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet. *Journal of Computational Physics*, 296:348–368, 2015. doi: 10.1016/j.jcp.2015.04.047.
- [82] Jaydeep M. Karandikar, Nam Ho Kim, and Tony L. Schmitz. Prediction of remaining useful life for fatigue-damaged structures using Bayesian inference. *Engineering Fracture Mechanics*, 96:588–605, 2012. doi: 10.1016/j.engfracmech.2012.09.013.
- [83] H. Rappel, L. A.A. Beex, J. S. Hale, et al. A Tutorial on Bayesian Inference to Identify Material Parameters in Solid Mechanics. *Archives of Computational Methods in Engineering*, 27(2):361–385, 2020. doi: 10.1007/s11831-018-09311-x.
- [84] Pinghe Ni, Qiang Han, Xiuli Du, and Xiaowei Cheng. Bayesian model updating of civil structures with likelihood-free inference approach and response reconstruction technique. *Mechanical Systems and Signal Processing*, 164(March 2021):108204, 2022. doi: 10.1016/j.ymssp.2021.108204.
- [85] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine Learning for Fluid Mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020. doi: 10.1146/annurev-fluid-010719-060214.
- [86] Alexandros Kontogiannis, Scott V. Elgersma, Andrew J. Sederman, and Matthew P. Juniper. Joint reconstruction and segmentation of noisy velocity images as an inverse Navier-Stokes problem. *Journal of Fluid Mechanics*, 944:1–36, 2022. doi: 10.1017/jfm.2022.503.
- [87] Michael B. Giles and Niles A. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65(3-4):393–415, 2000. doi: 10.1023/A:1011430410075.
- [88] Paolo Luchini and Alessandro Bottaro. Adjoint equations in stability analysis. *Annual Review of Fluid Mechanics*, 46:493–517, 2014. doi: 10.1146/annurev-fluid-010313-141253.
- [89] J. Reuther, J. J. Alonso, M. J. Rimlinger, and A. Jameson. Aerodynamic shape optimization of supersonic aircraft configurations via an adjoint formulation on distributed memory parallel computers. *6th Symposium on Multidisciplinary Analysis and Optimization*, (January), 1996. doi: 10.2514/6.1996-4045.
- [90] M. Pini, G. Persico, D. Pasquale, and S. Rebay. Adjoint method for shape optimization in real-gas flow applications. *Journal of Engineering for Gas Turbines and Power*, 137(3):1–13, 2015. doi: 10.1115/1.4028495.
- [91] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998. ISBN 9780511802256. doi: 10.1017/cbo9780511802256.011.
- [92] Francesco Garita. *Physics-Based Statistical Learning in Thermoacoustics*. PhD thesis, University of Cambridge, 2021.

- [93] Noemi Petra, James Martin, Georg Stadler, and Omar Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems, part II: stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, 36(4):1525–1555, 2014.
- [94] Robb J. Muirhead. *Aspects of Multivariate Statistical Theory.*, volume 79. 1984. ISBN 0471769851. doi: 10.2307/2288369.
- [95] D. V. Lindley. On a Measure of the Information Provided by an Experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956. doi: 10.1214/aoms/1177728069.
- [96] Xun Huan and Youssef M. Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232(1):288–317, 2013. doi: 10.1016/j.jcp.2012.08.013.
- [97] Kenneth J. Ryan. Estimating Expected Information Gains for Experimental Designs with Application to the Random Fatigue-Limit Model. *Journal of Computational and Graphical Statistics*, 12(3):585–603, 2003. doi: 10.1198/1061860032012.
- [98] Siddhartha Verma, Costas Papadimitriou, Nora Lüthen, et al. Optimal sensor placement for artificial swimmers. *Journal of Fluid Mechanics*, 884, 2019. doi: 10.1017/jfm.2019.940.
- [99] Thomas J. Loredo. Bayesian Adaptive Exploration. In *AIP Conference Proceedings*, volume 330, pages 195–206, 2004. ISBN 0735401829. doi: 10.1063/1.1835214.
- [100] David J. C. MacKay. Information-Based Objective Functions for Active Data Selection. *Neural Computation*, 4(4):590–604, 1992. doi: 10.1162/neco.1992.4.4.590.
- [101] Dimitrios I. Papadimitriou and Costas Papadimitriou. Optimal sensor placement for the estimation of turbulence model parameters in CFD. *International Journal for Uncertainty Quantification*, 5(6):545–568, 2015. doi: 10.1615/Int.J.UncertaintyQuantification.2015015239.
- [102] Omid Bidar, Sean Anderson, and Ning Qin. A Priori Sensor Placement Strategy for Turbulent Mean Flow Reconstruction Using Parametric Model Perturbations. In *AIAA SciTech Forum*, number January, pages 1–9, 2024.
- [103] S. Kullback and R.A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 1(22), 1951.
- [104] A Sgarro. Informational divergence and the dissimilarity of probability distributions. *Calcolo*, (18), 1981.
- [105] Matthew Juniper and Matthew Yoko. Data Assimilation with Laplace’s Method in Thermoacoustics. In *Symposium on Thermoacoustics in Combustion: Industry meets Academia*, 2021.
- [106] Matthew P Juniper and Matthew Yoko. Generating a physics-based quantitatively-accurate model of an electrically-heated Rijke tube with Bayesian inference. *Journal of Sound and Vibration*, 535(December 2021):117096, 2022. doi: 10.1016/j.jsv.2022.117096.

- [107] Matthew Yoko and Matthew P Juniper. Minimizing the data required to train a physics-based thermoacoustic model. In *29th international congress on sound and vibration*, 2023.
- [108] Matthew Yoko and Matthew P Juniper. Optimal experiment Bayesian inference design with. *Data-Centric Engineering [Under review]*, pages 1–24, 2024.
- [109] Matthew P. Juniper. Sensitivity analysis of thermoacoustic instability with adjoint Helmholtz solvers. *Physical Review Fluids*, 3(11), 2018. doi: 10.1103/PhysRevFluids.3.110509.
- [110] K. T. Feldman, Jr. Review of the literature on Rijke thermoacoustic phenomena. *Journal of Sound and Vibration*, 7(1):83–89, 1968. doi: 10.1016/0022-460X(68)90159-4.
- [111] R L Raun, M W Beckstead, J C Finlinson, and K P Brooks. A review of Rijke tubes, Rijke burners and related devices. *Progress in Energy and Combust Science*, 19:313–364, 1993. doi: 10.1016/0360-1285(93)90007-2.
- [112] G. Bisio and G. Rubatto. Sondhauss and Rijke oscillations - Thermodynamic analysis, possible applications and analogies. *Energy*, 24(2):117–131, 1999. doi: 10.1016/S0360-5442(98)00090-5.
- [113] Boa Teh Chu. On the energy transfer to small disturbances in fluid flow (Part I). *Acta Mechanica*, 1(3):215–234, 1965. doi: 10.1007/BF01387235.
- [114] G. F. Carrier. The mechanics of the Rijke tube. *Quarterly of Applied Mathematics*, 12(4):383–395, 1955. doi: 10.1090/qam/69698.
- [115] M.J. Lighthill. The response of laminar skin friction and heat transfer to fluctuations in the stream velocity. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 224(1156):1–23, 1954. doi: 10.1098/rspa.1954.0137.
- [116] H. J. Merk. Analysis of heat-driven oscillations of gas flows. *Applied Scientific Research*, 6(5-6):402–420, 1957. doi: 10.1007/bf03185045.
- [117] B. J. Bayly. Onset and equilibration of oscillations in general Rijke devices. *Journal of the Acoustical Society of America*, 79(3):846–851, 1986. doi: 10.1121/1.393475.
- [118] Sathesh Mariappan. Theoretical and Experimental Investigation of the Non-Normal Nature of Thermoacoustic Interactions. *Aerospace Engineering*, (December), 2011.
- [119] F. Selimefendigil, S. Föller, and W. Polifke. Nonlinear identification of unsteady heat transfer of a cylinder in pulsating cross flow. *Computers and Fluids*, 53(1):1–14, 2012. doi: 10.1016/j.compfluid.2011.08.012.
- [120] Armin Witte and Wolfgang Polifke. Dynamics of unsteady heat transfer in pulsating flow across a cylinder. *International Journal of Heat and Mass Transfer*, 109:1111–1131, 2017. doi: 10.1016/j.ijheatmasstransfer.2017.02.072.
- [121] M Heckl. Non-linear acoustic effect in the Rijke Tube. *Acoustica*, 72:63–71, 1990.
- [122] Armin Witte. *Dynamics of unsteady heat transfer in pulsating flow across a cylinder*. PhD thesis, T. U. Munich, 2018.

- [123] Konstantin Matveev. Thermoacoustic Instabilities in the Rijke Tube: Experiments and Modeling. *Thesis*, 2003:xiii–1161, 2003.
- [124] John William Strutt Baron Rayleigh. *The theory of sound vol. 2*. 1896.
- [125] Lawrence E Kinsler, Austin R Frey, Alan B Coppens, and James V Sanders. *Fundamentals of acoustics*. 1982.
- [126] C. Nicoli and P. Pelcè. One-dimensional model for the Rijke tube. *Journal of Fluid Mechanics*, 202(83):83–96, 1989. doi: 10.1017/S0022112089001102.
- [127] Gregor Gelbert, Jonas P. Moeck, Christian O. Paschereit, and Rudibert King. Feedback control of unstable thermoacoustic modes in an annular Rijke tube. *Control Engineering Practice*, 20(8):770–782, 2012. doi: 10.1016/j.conengprac.2012.03.016.
- [128] Luca Magri and Matthew P. Juniper. Sensitivity analysis of a time-delayed thermoacoustic system via an adjoint-based approach. *Journal of Fluid Mechanics*, 719: 183–202, 2013. doi: 10.1017/jfm.2012.639.
- [129] Luca Magri, Michael Bauerheim, and Matthew P. Juniper. Stability analysis of thermoacoustic nonlinear eigenproblems in annular combustors. Part I. Sensitivity. *Journal of Computational Physics*, 325:395–410, 2016. doi: 10.1016/j.jcp.2016.07.032.
- [130] H Tjrdeman. On the propagation of sound waves in cylindrical tubes. *Journal of sound and vibration*, 39:1–33, 1974.
- [131] Harold Levine and Julian Schwinger. On the radiation of sound from an unflanged circular pipe. *Physical review*, 73(4):383–406, 1948.
- [132] S.W. Rienstra and A Hirschbert. *An introduction to acoustics*. Eindhoven U. Tech., 2003.
- [133] L. V. King. On the convection of heat from small cylinders in a stream of fluid: determination of the convection constants of small platinum wires with applications to hot-wire anemometry. *Philosophical Transactions of the Royal Society*, 214:373–432, 1914.
- [134] D. Huang, T. T. Allen, W. I. Notz, and N. Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of Global Optimization*, 34(3):441–466, 2006. doi: 10.1007/s10898-005-2454-3.
- [135] U. Sengupta and M. P. Juniper. Thermoacoustic stabilization of combustors with gradient-augmented Bayesian optimization and adjoint models. *Symposium on Thermoacoustics in Combustion, 6-10 September 2021, Munich, Germany*, pages 1–11, 2021.
- [136] Tulay Ercan and Costas Papadimitriou. Bayesian optimal sensor placement for parameter estimation under modeling and input uncertainties. *Journal of Sound and Vibration*, 563 (May), 2023. doi: 10.1016/j.jsv.2023.117844.
- [137] Matthew Yoko and Matthew P Juniper. Data-driven modelling of thermoacoustic instability in a ducted conical flame. In *Symposium on Thermoacoustics in Combustion, 11-14 September 2021, Zurich, Switzerland*, pages 1–12, Zurich, Switzerland, 2023.

- [138] Matthew Yoko and Matthew P Juniper. Adjoint-accelerated Bayesian inference applied to the thermoacoustic behaviour of a ducted conical flame. *Journal of Fluid Mechanics [Under Review]*, 2024.
- [139] Viktor Kornilov. *Experimental Research of Acoustically Perturbed Bunsen Flames*. Number december. 2006. ISBN 9789038630380. doi: 10.6100/IR615531.
- [140] V. N. Kornilov, K. R.A.M. Schreel, and L. P.H. De Goey. Experimental assessment of the acoustic response of laminar premixed Bunsen flames. *Proceedings of the Combustion Institute*, 31 I(1):1239–1246, 2007. doi: 10.1016/j.proci.2006.07.079.
- [141] Dan Zhao and Z. H. Chow. Thermoacoustic instability of a laminar premixed flame in Rijke tube with a hydrodynamic region. *Journal of Sound and Vibration*, 332(14): 3419–3437, 2013. doi: 10.1016/j.jsv.2013.01.031.
- [142] F. Baillot, A. Bourehla, and D. Durox. The characteristics method and cusped flame fronts. *Combustion Science and Technology*, 112:327–350, 1996. doi: 10.1080/00102209608951963.
- [143] T. Schuller, D. Durox, and S. Candel. A unified model for the prediction of laminar flame transfer functions: Comparisons between conical and V-flame dynamics. *Combustion and Flame*, 134(1-2):21–34, 2003. doi: 10.1016/S0010-2180(03)00042-7.
- [144] Maria A. Heckl and M. S. Howe. Stability analysis of the Rijke tube with a Green’s function approach. *Journal of Sound and Vibration*, 305(4-5):672–688, 2007. doi: 10.1016/j.jsv.2007.04.027.
- [145] Thierry Schuller, Thierry Poinsot, and Sébastien Candel. Dynamics and control of premixed combustion systems based on flame transfer and describing functions. *Journal of Fluid Mechanics*, 894, 2020. doi: 10.1017/jfm.2020.239.
- [146] Thomas Steinbacher and Wolfgang Polifke. Convective Velocity Perturbations and Excess Gain in Flame Response as a Result of Flame-Flow Feedback. *Fluids*, 7(2), 2022. doi: 10.3390/fluids7020061.
- [147] Alessandro Giannotta, Stefania Cherubini, Pietro De Palma, and Matthew P. Juniper. The effect of flame curvature and flame base movement on the frequency response of a conical Bunsen flame. *Combustion and Flame*, 259(October 2023):113179, 2024. doi: 10.1016/j.combustflame.2023.113179.
- [148] D. Mejia, M. Miguel-Brebion, and L. Selle. On the experimental determination of growth and damping rates for combustion instabilities. *Combustion and Flame*, 169: 287–296, 2016. doi: 10.1016/j.combustflame.2016.05.004.
- [149] A. L. Birbaud, D. Durox, and S. Candel. Upstream flow dynamics of a laminar premixed conical flame submitted to acoustic modulations. *Combustion and Flame*, 146(3): 541–552, 2006. doi: 10.1016/j.combustflame.2006.05.001.
- [150] Alexis Cuquel, Daniel Durox, and Thierry Schuller. Impact of flame base dynamics on the non-linear frequency response of conical flames. *Comptes Rendus - Mécanique*, 341 (1-2):171–180, 2013. doi: 10.1016/j.crme.2012.11.004.

- [151] Nicolas Noiray. Linear Growth Rate Estimation from Dynamics and Statistics of Acoustic Signal Envelope in Turbulent Combustors. *Journal of Engineering for Gas Turbines and Power*, 139(4), 2017. doi: 10.1115/1.4034601.
- [152] N. Noiray and A. Denisov. A method to identify thermoacoustic growth rates in combustion chambers from dynamic pressure time series. *Proceedings of the Combustion Institute*, 36(3):3843–3850, 2017. doi: 10.1016/j.proci.2016.06.092.
- [153] Christian Oliver Paschereit, Bruno Schuermans, Wolfgang Polifke, and Oscar Mattson. Measurement of transfer matrices and source terms of premixed flames. *Proceedings of the ASME Turbo Expo*, 2, 1999. doi: 10.1115/99-GT-133.
- [154] Nicholas C. W. Treleaven, Andre Fischer, Claus Lahiri, et al. The effects of forcing direction on the flame transfer function of a lean-burn spray flame. In *Proceedings of the ASME Turbo Expo*, 2021.
- [155] Andre Fischer and Claus Lahiri. Ranking of aircraft fuel-injectors regarding low frequency thermoacoustics based on an energy balance method. *Proceedings of the ASME Turbo Expo*, 2021.
- [156] M. L. Munjal and A. G. Doige. Theory for of a Two Source-Location Parameters Element Method of an Experimental Evaluation Four-Pole. *Journal of Sound and Vibration*, 141: 323–333, 1990.
- [157] Abdulla Ghani, Isaac Boxx, and Carrie Noren. Data-driven identification of nonlinear flame models. *Journal of Engineering for Gas Turbines and Power*, 142(12):1–7, 2020. doi: 10.1115/1.4049071.
- [158] Abdulla Ghani and Alp Albayrak. From Pressure Time Series Data to Flame Transfer Functions: A Framework for Perfectly Premixed Swirling Flames. *Journal of Engineering for Gas Turbines and Power*, 145(1):1–9, 2023. doi: 10.1115/1.4055724.
- [159] F. Gant, G. Ghirardo, A. Cuquel, and M. R. Bothien. Delay Identification in Thermoacoustics. *Journal of Engineering for Gas Turbines and Power*, 144(2):1–10, 2022. doi: 10.1115/1.4052060.
- [160] Nicolas Noiray and Bruno Schuermans. Deterministic quantities characterizing noise driven Hopf bifurcations in gas turbine combustors. *International Journal of Non-Linear Mechanics*, 50:152–163, 2013. doi: 10.1016/j.ijnonlinmec.2012.11.008.
- [161] David G Goodwin, Harry K Moffat, Ingmar Schoegl, et al. Cantera: An Object-oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes. <https://www.cantera.org>, 2022.
- [162] Stephen R Turns. *An introduction to combustion*. McGraw-Hill, 3 edition, 2012.
- [163] A. N. Norris and I. C. Sheng. Acoustic radiation from a circular pipe with an infinite flange. *Journal of Sound and Vibration*, 135(1):85–93, 1989. doi: 10.1016/0022-460X(89)90756-6.

- [164] Dan Zhao. Transient growth of flow disturbances in triggering a Rijke tube combustion instability. *Combustion and Flame*, 159(6):2126–2137, 2012. doi: 10.1016/j.combustflame.2012.02.002.
- [165] Kristin M. Kopp-Vaughan, Steven G. Tuttle, Michael W. Renfro, and Galen B. King. Heat release and flame structure measurements of self-excited acoustically-driven premixed methane flames. *Combustion and Flame*, 156(10):1971–1982, 2009. doi: 10.1016/j.combustflame.2009.06.018.
- [166] Alexis Cuquel, Daniel Durox, and Thierry Schuller. Experimental determination of flame transfer function using random velocity perturbations. *Proceedings of the ASME Turbo Expo*, 2011.
- [167] Thierry Schuller, Sébastien Ducruix, Daniel Durox, and Sébastien Candel. Modeling tools for the prediction of premixed flame transfer functions. *Proceedings of the Combustion Institute*, 29(1):107–113, 2002. doi: 10.1016/S1540-7489(02)80018-9.
- [168] Matthew Yoko and Matthew P. Juniper. Inferring flame transfer functions of turbulent conical flames from pressure measurements. In *Proceedings of the ASME Turbo Expo*, London, UK, 2024.
- [169] P. Palies, D. Durox, T. Schuller, and S. Candel. The combined dynamics of swirler and turbulent premixed swirling flames. *Combustion and Flame*, 157(9):1698–1717, 2010. doi: 10.1016/j.combustflame.2010.02.011.
- [170] Jonas P. Moeck, Jean François Bourgouin, Daniel Durox, et al. Nonlinear interaction between a precessing vortex core and acoustic oscillations in a turbulent swirling flame. *Combustion and Flame*, 159(8):2650–2668, 2012. doi: 10.1016/j.combustflame.2012.04.002.
- [171] Jonas P. Moeck, Jean François Bourgouin, Daniel Durox, et al. Tomographic reconstruction of heat release rate perturbations induced by helical modes in turbulent swirl flames. *Experiments in Fluids*, 54(4), 2013. doi: 10.1007/s00348-013-1498-2.
- [172] J Li, D Durox, F Richecoeur, and T Schuller. Analysis of chemiluminescence, density and heat release rate fluctuations in acoustically perturbed laminar premixed flames. *Combustion and Flame*, 162:3934–3945, 2015. doi: 10.1016/j.combustflame.2015.07.031.
- [173] Alexandros Kontogiannis and Matthew Juniper. Inverse problems in magnetic resonance velocimetry: Shape, forcing and boundary condition inference. *American Society of Mechanical Engineers, Fluids Engineering Division (Publication) FEDSM*, 2:1–11, 2021. doi: 10.1115/FEDSM2021-66080.
- [174] Santosh Hemchandra, Anindya Datta, and Matthew P. Juniper. Learning RANS model parameters from LES using Bayesian inference. In *Proceedings of the ASME Turbo Expo*, Boston, Massachusetts, 2023. ISBN 9780791886953. doi: 10.1115/gt2023-102159.
- [175] Alessandro Giannotta, Matthew Yoko, Stefania Cherubini, et al. Bayesian data assimilation of acoustically forced laminar premixed conical flames. In *Symposium on Thermoacoustics in Combustion, 11-14 September 2023, Zurich, Switzerland*, Zurich, Switzerland, 2023.

- [176] Luca Magri. Adjoint Methods as Design Tools in Thermoacoustics. *Applied Mechanics Reviews*, 71(2), 2019. doi: 10.1115/1.4042821.
- [177] Jinquan Zheng, Andre Fischer, Claus Lahiri, et al. Bayesian data assimilation in cold flow experiments on an industrial thermoacoustic rig. In *Proceedings of the ASME Turbo Expo*, 2024.

Appendix A

Properties of the flames studied in Chapters 4 and 5

Table A.1 Summary of the properties of the 24 laminar flames studied. We show the average measured flow rates of air, methane (CH₄) and ethylene (C₂H₄), the equivalence ratio (ϕ), the bulk velocity in the burner tube (\bar{U}), the predicted and measured flame lengths ($L_{f,m}$ and $L_{f,p}$), the predicted and measured convective time delays ($\tau_{c,m}$ and $\tau_{c,p}$), and the inner cone mean heat release rate (\bar{Q}).

Group	Air	CH ₄	C ₂ H ₄	ϕ	\bar{U}	$L_{f,p}$	$L_{f,m}$	$\tau_{c,p}$	$\tau_{c,m}$	\bar{Q}
[-]	[ln/min]	[ln/min]	[ln/min]	[-]	[m/s]	[mm]	[mm]	[ms]	[ms]	[W]
1	6.049	0.325	0.325	1.28	1.75	17.4	16.6	9.9	9.5	374.9
2	6.147	0.348	0.348	1.35	1.79	20.5	20.0	11.5	11.2	374.9
3	6.219	0.364	0.364	1.40	1.82	23.6	23.1	13.0	12.7	374.9
4	6.283	0.379	0.379	1.44	1.84	26.9	26.4	14.6	14.3	374.9
5	6.338	0.391	0.391	1.47	1.86	30.1	28.9	16.2	15.5	374.9
6	6.384	0.401	0.401	1.50	1.88	33.0	32.1	17.5	17.1	374.9
1	7.246	0.387	0.387	1.27	2.10	20.8	19.8	9.9	9.4	450.0
2	7.369	0.416	0.416	1.34	2.15	24.6	23.8	11.5	11.1	449.9
3	7.459	0.436	0.436	1.39	2.18	28.4	27.4	13.0	12.6	449.9
4	7.537	0.454	0.454	1.43	2.21	32.2	30.9	14.6	14.0	449.9
5	7.603	0.468	0.468	1.47	2.24	36.1	34.2	16.2	15.3	449.9
6	7.659	0.481	0.481	1.50	2.26	39.6	37.2	17.6	16.5	449.9
1	8.444	0.449	0.449	1.27	2.45	24.2	23.1	9.9	9.4	525.0
2	8.594	0.484	0.484	1.34	2.51	28.7	27.6	11.5	11.0	524.9
3	8.699	0.508	0.508	1.39	2.55	33.1	31.7	13.0	12.4	524.9
4	8.790	0.529	0.529	1.43	2.58	37.6	36.4	14.6	14.1	524.9
5	8.868	0.546	0.546	1.47	2.61	42.2	39.2	16.2	15.0	524.8
6	8.934	0.561	0.561	1.49	2.64	46.3	43.0	17.6	16.3	524.9
1	9.644	0.512	0.512	1.26	2.80	27.7	25.9	9.9	9.3	600.0
2	9.818	0.553	0.553	1.34	2.86	32.8	30.1	11.5	10.5	599.9
3	9.939	0.580	0.580	1.39	2.91	37.9	34.8	13.0	12.0	599.9
4	10.045	0.604	0.604	1.43	2.95	43.0	39.9	14.6	13.5	599.9
5	10.134	0.624	0.624	1.47	2.99	48.2	43.6	16.2	14.6	599.8
6	10.209	0.641	0.641	1.49	3.01	52.9	47.2	17.6	15.6	599.8

Table A.2 Summary of the properties of the 15 turbulent flames studied. We show the average measured flow rates of air, methane (CH₄) and ethylene (C₂H₄), the equivalence ratio (ϕ), the bulk velocity in the burner tube (\bar{U}), the measured flame length (L_f), the measured convective time delay (τ_c), and the inner cone mean heat release rate (\bar{Q}).

Flame [-]	Air [ln/min]	CH ₄ [ln/min]	C ₂ H ₄ [ln/min]	ϕ [-]	\bar{U} [m/s]	L_f [mm]	τ_c [ms]	\bar{Q} [W]
1	15.742	0.720	0.720	1.09	4.50	30.6	6.8	1000.7
2	15.732	0.740	0.740	1.12	4.51	31.0	6.9	1000.5
3	15.760	0.760	0.760	1.15	4.53	31.5	7.0	1000.4
4	15.810	0.780	0.780	1.17	4.56	32.4	7.1	1000.3
5	15.873	0.800	0.800	1.20	4.58	33.5	7.3	1000.2
6	15.945	0.820	0.820	1.22	4.62	34.8	7.5	1000.2
7	16.023	0.840	0.840	1.25	4.65	36.3	7.8	1000.2
8	16.104	0.860	0.860	1.27	4.68	38.1	8.1	1000.1
9	16.188	0.880	0.880	1.29	4.72	40.3	8.5	1000.1
10	16.273	0.900	0.900	1.32	4.75	42.7	9.0	1000.1
11	16.360	0.920	0.920	1.34	4.78	45.3	9.5	1000.1
12	16.448	0.940	0.940	1.36	4.82	48.3	10.0	1000.0
13	16.536	0.960	0.960	1.38	4.85	51.4	10.6	1000.0
14	16.625	0.980	0.980	1.40	4.89	55.2	11.3	1000.0
15	16.714	1.000	1.000	1.42	4.92	59.0	12.0	1000.0

Appendix B

Photographs of the experimental rigs




Fig. B.1 Photograph of the hot wire Rijke tube.




Fig. B.2 Photograph of the electric heater.




Fig. B.3 Photograph of the ducted flame rig.