Data Assimilation with Laplace's Method in Thermoacoustics

Symposium on Thermoacoustics in Combustion: Industry meets Academia (SoTiC 2021)
Sept. 6 - Sept. 10, 2021
Munich, Germany
Paper No.: 8435
©The Author(s) 2021

Matthew Juniper[1] and Matthew Yoko[1]

Abstract

Several physics-based candidate models of a laboratory thermoacoustic system are developed. The most likely parameters of these models are inferred by assimilating data from around 7000 experimental measurements. The parameter covariance matrix is calculated with Laplace's method using first and second order adjoint methods. This covariance matrix quantifies the uncertainties in each parameter, as well as the joint uncertainty of each pair of parameters. The posterior likelihood is integrated in parameter space to obtain the Marginal Likelihood and Occam Factor for each candidate model. The Marginal Likelihood quantifies the evidence for each model, given the data, and is used to rank the models. We find that a two parameter $n - \tau$ model is the most likely model and that a model containing Lighthill's time delay is the least likely. This method turns qualitatively-accurate models into quantitatively-accurate models with known uncertainty bounds. The method is general and could be applied widely in thermoacoustics.

Keywords

Thermoacoustics, Data Assimilation, Laplace Method, Bayesian Inference

Introduction

Although it is possible to create *qualitatively*-accurate thermoacoustic models, it is difficult to create a *quantitatively*-accurate model of a particular system from quantitatively-accurate models of its components[1]. This is because thermoacoustic systems are often, for physical reasons, exceedingly sensitive to small changes[2]. This sensitivity introduces significant systematic error into models and their parameters. For example, a highly detailed model of thermoacoustic oscillations of an electric heater in a tube[3], despite being carefully tuned to be quantitatively correct at one heater position, was only qualitatively correct at nearby heater positions (Figs. 5-5 to 5-8)[3]. For large devices, accurate prediction of thermoacoustic behaviour is similarly challenging[4].

In this study we propose qualitatively-accurate physics-based thermoacoustic models and then rigorously (i) tune their parameters by assimilating data from experiments; (ii) quantify the uncertainties in each model's parameters; (iii) quantify the evidence (the marginal likelihood) for each model; (iv) compare the models and select the best model.

We use a Bayesian approach with Laplace's method[5,6], which is technically more difficult to implement than methods such as Markov Chain Monte Carlo. For this problem, Laplace's method is several orders of magnitude faster than Markov Chain Monte Carlo. It therefore provides a practical method to compare different models quickly and accurately. The long term goal is to use these quantitatively-accurate model for adjoint-based control and design.

Experimental configuration

The experimental configuration is a vertical Rijke tube containing an electric heater[7]. The stainless steel tube is 1 m long, has internal diameter 47.4 mm and wall thickness 1.7 mm. The heater consists of two concentric annular ceramic discs with inner diameter 31.6 mm and outer diameter 47 mm. Flow passes through the central hole in each disc and over nichrome wire, which is wrapped around the discs. The electric heater is placed at 19 different positions from the bottom end of the tube (Tab. 1). Eight probe microphones record the pressure near the inner surface of the tube from $x_\mathrm{m}/L = 0.25$ to 0.95 in steps of 0.1. Eight thermocouples are placed from $x_\mathrm{t}/L = 0.2$ to 0.9 in steps of 0.1. The ambient temperature is measured with a thermocouple at the inlet of the tube.

The heater power is set to $0, 7, 15, 30, 50, 80, 130, 180$ Watts, for 62.5 minutes at each power. Every 15 seconds, a loudspeaker at the base of the tube forces the system close to its resonant frequency for 6 seconds, referred to in this paper as a 'ping'. The probe microphones measure the response throughout the tube at 10 kHz during the forced and the decaying period. The system is always thermoacoustically stable at these heater powers, although the decay rate depends on the thermoacoustic driving from the heater.

The experiment exhibits a long timescale (typically 1000 seconds), at which the tube heats up and reaches steady state, and a short timescale (typically 1 second), at which the acoustic oscillations decay. At the long timescale we assimilate the thermocouple and sound speed measurements. At the short timescale we assimilate the decay rate measurements, assuming that the long timescale flow is quasi-steady. The quasi-steady assumption is reasonable because of the large difference in the timescales.

[1]Cambridge University Engineering Department, UK

Corresponding author:
Matthew Juniper, Cambridge University Engineering Department, Trumpington Street, Cambridge, CB2 1PZ, UK.
Email: mpj1001@cam.ac.uk

Table 1. Summary of all the experiments assimilated in this paper. Column 1: The three different cold experiments are labelled C1, C2 & C3 and the hot experiments are labelled H. 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. Column 4 is the 'empty tube' case, in which the heater and prongs have been removed from the tube.

			Heater position (cm)																			
	T	Q	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	–	–
H	✓	7	–	–	–	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	–	–
H	✓	30	–	–	–	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	–	–
H	✓	80	–	–	–	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	–	–
H	✓	180	–	–	–	40	40	40	40	40	40	40	40	40	40	40	40	40	40	40	–	–

Measurement of the local speed of sound

The acoustic frequency at the short timescale depends strongly on the local speed of sound in the tube. A previous study[7] inferred the local speed of sound from the local gas temperature measurements. These were assimilated into a model of the long timescale flow using an ensemble Kalman filter. Assimilation of the short timescale data into physics-based models using Markov Chain Monte Carlo implied that the thermoacoustic behaviour of the heater depended linearly on its position in the tube[8]. Although such a model could fit the data well, we could not conceive a physical reason for this dependence and suspected instead a systematic bias in one of the measurements. The linear dependence on heater position was consistent with a dependence on the length of hot gas in the tube, which in turn was consistent with a systematic error in the local temperature. This was also consistent with the gas temperature measurement being affected by radiation from the heater. We therefore decided to measure the local speed of sound directly with the probe microphones.

Before each ping, the loudspeaker sends an impulse down the tube, known in this paper as a 'click'. The 8 probe microphones measure the response along the tube at 62.5 kHz. We find the phase shift that maximises 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. This local sound speed is fed into a model of the long timescale flow via an Ensemble Kalman Filter, described in the next section.

The channels of the acquisition system have slightly different delays. We measure these delays by performing several clicks at ambient sound speed. We then subtract these delays from the measurements, thus removing a significant source of systematic error. We find that this delay depends on the sampling frequency so we perform a second calibration for the ping experiments.

Inference of the long timescale model parameters with an Ensemble Kalman Filter

The long timescale flow is an unsteady plug flow conjugate heat transfer model extended from a previous study[7] to include 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: the inviscid drag coefficient of the heater, the Nusselt number on the outer surface of the tube, the thermal resistance at the tube mounts, and the proportion of supplied power that conducts through the tube wall. Two parameters describe the Nusselt number distribution inside the tube. For a given configuration, the model outputs the velocity of the convective flow, the local temperature in the gas, and the local temperature on the inner and outer walls of the tube.

100 realisations of the long timescale model are iterated as an ensemble. The parameters of each ensemble member are randomly sampled from a uniform prior distribution. At each assimilation step, the ensemble forecast and experimental measurements are supplied to an Ensemble Kalman Filter (EnKF). These measurements consist of the temperature at 17 axial locations along the tube wall, the click arrival time at each microphone, and the measured heater power. The EnKF returns the expected values and variances of the system's state and the model's parameters. As more data become available, the state and parameters converge to a constant values with high certainty. The local sound speed and local convective velocity are then extracted and used in the model of the short timescale flow.

Thermoacoustic Network Model for the short timescale flow

In the ping experiments, acoustic oscillations are forced by the loudspeaker and damped by acoustic radiation from the tube, and by visco-thermal damping in the acoustic boundary layers, around the prongs holding the heater, at the heater, and at the thermocouples. Acoustic oscillations are also driven or damped by heat release rate fluctuations at the hot wire. Acoustic waves are modelled in 40 to 50 acoustic elements within the tube, depending on the configuration[7]. The speed of sound in each element is extracted from the long timescale model. At the interfaces between the acoustic elements, the waves in adjacent elements are matched by jump conditions for the momentum and energy equations. The viscothermal drag and the heat release rate from the wire are all modelled as local feedback mechanisms at these interfaces[9]. Viscous drag is modelled as local feedback from

the velocity field into the momentum equation, thermal drag is modelled as local feedback from the temperature field into the energy equation, and the heat release rate from the wire is modelled as local feedback from the velocity field into the energy equation. Wave reflection at the upstream and downstream ends of the tube are modelled by complex reflection coefficients.

We only assimilate data from pings recorded after the long timescale system has reached steady state. This removes systematic error due to the delay between the experiment and the Ensemble Kalman filter. For each configuration we group all the fixed parameters, variable parameters, and associated experimental measurements. We then obtain the eigenvalue (decay rate and frequency) of each configuration and, using first and second order adjoint methods, obtain the first and second derivatives of the eigenvalue with respect to all the variable parameters. These derivatives are then combined to obtain the derivatives with respect to the viscothermal drag at the prongs holding the heater, at the heater itself, and at the thermocouples, as well as the derivatives with respect to the heat release rate model at the heater. These derivatives are used in the data assimilation process, described next.

Maximum a Posteriori parameter estimation

Each short timescale model, $\mathcal{H}_i$, has a set of variable parameters, $\mathbf{a}$. First we assume that each model is true and infer its parameters from the data D[5]. The posterior probability of the parameters $\mathbf{a}$ is:

$$P(\mathbf{a}|D,\mathcal{H}_i) = \frac{P(D|\mathbf{a},\mathcal{H}_i)P(\mathbf{a}|\mathcal{H}_i)}{P(D,\mathcal{H}_i)} \tag{1}$$

At this level of inference the denominator of (1) is ignored and the *maximum posterior likelihood* of the model parameters, $\mathbf{a}_{MP}$, is found by maximizing the numerator.

The data, D, consists of nine complex numbers for each ping: the eigenvalue, s (whose real part is the growth rate and whose imaginary part is the angular frequency), and the Fourier-decomposed complex pressure, P, at the eight microphone locations. A cost function, $\mathcal{J}$, is defined as the negative log of the posterior likelihood. If all distributions are assumed to be Gaussian then $\mathcal{J}$ is the sum of the squares of the discrepancy between the model parameters and their prior estimates, weighted by the confidence in the prior estimates, added to the sum of the squares of the discrepancies between the model predictions and the experimental measurements, weighted by the confidence in the experimental measurements:

$$\begin{aligned}
\mathcal{J} &= -\log\{P(D|\mathbf{a},\mathcal{H}_i)P(\mathbf{a}|\mathcal{H}_i)\} \\
&= (s_r(\mathbf{a})-z_r)^T C_{sr}^{-1}(s_r(\mathbf{a})-z_r)\ldots \\
&+ (s_i(\mathbf{a})-z_i)^T C_{si}^{-1}(s_i(\mathbf{a})-z_i)\ldots \\
&+ (P_r(\mathbf{a})-Q_r)^T C_{pr}^{-1}(P_r(\mathbf{a})-Q_r)\ldots \\
&+ (P_i(\mathbf{a})-Q_i)^T C_{pi}^{-1}(P_i(\mathbf{a})-Q_i)\ldots \\
&+ (\mathbf{a}-\mathbf{a}_f)^T C_{aa}^{-1}(\mathbf{a}-\mathbf{a}_f)
\end{aligned} \tag{2}$$

where: $\mathbf{a}$ is a column vector containing the parameter values; C_{aa} is the prior covariance matrix of the parameters; $\mathbf{a}_f$ is a column vector containing the prior estimates of the parameter values; $s_r(\mathbf{a})$ is the model's growth rate; C_{sr} is the covariance matrix of the growth rate measurements; z_r is the measured growth rate in the experimental configuration represented by parameters $\mathbf{a}$; $s_i(\mathbf{a})$ is the model's frequency; z_i is the measured frequency; C_{si} is the covariance of the frequency measurements; P_r and P_i are the real and imaginary components of the pressure predictions at each microphone; Q_r and Q_i are the real and imaginary components of the pressure measurements.

The first derivative of the cost function $\mathcal{J}$ with respect to the parameters, $\mathbf{a}$, is derived via a first order Taylor expansion of (2). It is expressed in terms of the first derivatives of s and P with respect to the parameters. The derivatives of s and P are calculated from the network model using first order adjoint methods[2,10]. A BFGS (Broyden-Fletcher-Goldfarb-Shanno) gradient-based optimization algorithm is used to find the parameter values, $\mathbf{a}_{MP}$, that minimize the cost function $\mathcal{J}$, given the prior covariance $C_{\mathbf{aa}}$, and the measurement variances C_{sr}, C_{si}, C_{pr}, C_{pi}. This is known as the *maximum a posteriori estimation* (MAP).

If the measurement uncertainties are unknown, for example if an un-measured or unknown factor is affecting the measurements, then the measurement variances C_{sr}, C_{si}, C_{pr}, C_{pi} can be allowed to float. The derivative of $\mathcal{J}$ with respect to C_{sr}, C_{si}, C_{pr}, C_{pi} is easily calculated from (2). These derivatives can be included in the BFGS optimization algorithm in order to find the values of C_{sr}, C_{si}, C_{pr}, C_{pi} and $\mathbf{a}_{MP}$ at which $\mathcal{J}$ is minimized.

Parameter error estimation with Laplace's Method

The MAP method gives the most likely parameters, $\mathbf{a}$, given a model, $\mathcal{H}_i$, but does not yet give the uncertainties in the parameters. This is found with the Laplace Approximation, which is also known as the Saddle Point Method. In order to estimate these uncertainties we re-use the assumption that $P(\mathbf{a}|D,\mathcal{H}_i)$ is Gaussian around $\mathbf{a}_{MP}$ and we define its inverse covariance matrix around this point as $\mathbf{A}$:

$$-\log\{P(\mathbf{a}|D,\mathcal{H}_i)\} = \frac{1}{2}(\mathbf{a}-\mathbf{a}_{MP})^T\mathbf{A}(\mathbf{a}-\mathbf{a}_{MP}) + \text{const} \tag{3}$$

By inspection, $\mathbf{A}$ is simply the Hessian of $\mathcal{J}$:

$$A_{ij} = \frac{\partial^2\mathcal{J}}{\partial a_i a_j} \tag{4}$$

The second derivative of the cost function $\mathcal{J}$ with respect to the parameters, $\mathbf{a}$, is derived via a second order Taylor expansion of (2). This is expressed in terms of the first and second derivatives of s and P with respect to the parameters. The second derivatives, like the first derivatives, are found with adjoint methods[11,12]. It is worth mentioning that the contribution of the second derivatives to $\mathbf{A}$ is usually an order of magnitude smaller than the contribution of the first derivatives. It is also worth mentioning that the approximate Hessian created during the BFGS optimization algorithm is usually an excellent approximation to $\mathbf{A}$. This means that the second derivatives of s and P, which are expensive to calculate, often do not need to be calculated. In this paper, they are only used to improve accuracy of the final calculation and are not used in interim calculations.

Model Comparison with Laplace's Method

At the second level of inference, we wish to infer which model is most plausible given the data[5]. By Bayes' theorem the posterior probability of each model is:

$$P(\mathcal{H}_i|D) \propto P(D|\mathcal{H}_i)P(\mathcal{H}_i) \tag{5}$$

The second term on the RHS is our prior estimate for the probability of model $\mathcal{H}_i$, while the first term is the denominator of (1), which is known as the *evidence* or *marginal likelihood*. Assuming that each prospective model has equal probability $P(\mathcal{H}_i)$ then the models are ranked by evaluating the evidence:

$$P(D|\mathcal{H}_i) = \int_{\mathbf{a}} P(D|\mathbf{a}, \mathcal{H}_i)P(\mathbf{a}|\mathcal{H}_i)\mathrm{d}\mathbf{a} \tag{6}$$

Assuming furthermore that $P(\mathcal{H}_i|D)$ is Gaussian, this can be approximate by

$$P(D|\mathcal{H}_i) \approx P(D|\mathbf{a}_{MP}, \mathcal{H}_i)P(\mathbf{a}_{MP}|\mathcal{H}_i)\left(\det(\mathbf{A}/2\pi)\right)^{-1/2} \tag{7}$$

where $\mathbf{A}$ is the Hessian calculated in (4). This evidence, which is also known as the Marginal Likelihood, is calculated cheaply with Laplace's Method. The most likely model is that with the largest Marginal Likelihood. The larger the number of parameters, $\mathbf{a}$, in the model, $\mathcal{H}_i$, the larger the Hessian, $\mathbf{A}$. This penalizes over-elaborate models even if they fit the data well.

The Marginal Likelihood, like the MAP point, depends on the measurement uncertainty. Laplace's method also gives the gradient of the Marginal Likelihood with respect to the measurement uncertainties. This gradient can then be used to find the Maximum Marginal Likelihood (MML) of a model, given some data, allowing the measurement noise to float. In this paper we calculate both the MAP parameter values and the MML parameter values, and find that both are similar.

MAP and error estimation from the dataset

The data assimilation process described above is an optimization problem solved with a gradient-based method. It is tempting to assimilate all the parameters simultaneously using the entire experimental dataset. This requires, however, a good initial guess for the parameters and is usually impractical. Indeed there may be many local optima corresponding to unphysical model parameters and these would need to be de-selected after a global optimization procedure. Instead, in this paper, we assimilate the parameters sequentially using *a priori* knowledge that some of the parameters are zero in some of the experiments.

We start by assimilating the reflection coefficients, R_u and R_d, at ambient temperature. We use the decay rates and frequencies from the empty tube experiments, for which we can assume that $R_u = R_d \equiv R$ by symmetry. Using reasonable estimates for the measurement noises we obtain the posterior expected values and covariances for R shown in row 1 of table 2. We use these values as priors for the next assimilation, inflating the covariance of R so that it is not tightly constrained. We then re-assimilate R, together with the viscothermal drag coefficients of the heater prongs, k_{vis_p} and k_{th_p}, using experiments C3 in table 1, in which the heater prongs were moved down the tube without the heater attached. At this stage we also calculate the measurement noises that maximize the MAP. This gives row 2 of table 2. The prongs have a small but non-negligible influence on the system. Having calculated k_{vis_p} and k_{th_p}, we freeze them so that they play no further role in the assimilation.

Next we use the above prior for R, with covariance inflation, and re-assimilate R, together with the viscothermal drag coefficients of the heater, k_{vis_h} and k_{th_h}, using experiments C1, C2, and C3, which contain the prongs and heater but no thermocouples. This gives row 3 of table 3.

Next we use the above priors, with covariance inflation, and re-assimilate R, k_{vis_h}, k_{th_h}, and a factor K, defined such that $k_{vis_t} = K \times k_{vis_h}$ and $k_{th_t} = K \times k_{th_h}$, using the hot experiments at 0 Watts. This gives row 4 of table 3 and, further, that $\mathrm{Re(K)} = 0.0156 \pm 0.0074$ and $\mathrm{Im(K)} = (-8.99 \pm 2.69) \times 10^{-3}$.

For all the cold experiments, Fig. 1 shows the model predictions alongside the experimental values, with uncertainty intervals of two standard deviations. For the experiments with prongs, heater, and thermocouples, each heater position was tested on a different day. The variation in the predicted and measured growth rate and frequency is due to the different temperature and pressure on each day.

The reflection coefficients and the visco-thermal drag of the prongs, heater, and thermocouples are assimilated with the heater turned off so that the thermoacoustic effect is inactive and can be excluded from the assimilation. When the heater is turned on, the air leaving the tube is hot and it cannot be assumed that the downstream reflection coefficient, R_d, will not change[13]. We therefore measure R_d using the multi-microphone method during the hot experiments. This requires us to re-calibrate the probe microphones and their acquisition system at the new acquisition rate. We calibrate the microphones just before each hot run. We predict the pressure mode from the 0 Watt experiments using the visco-thermal drag coefficients assimilated previously. The calibration factor is calculated by dividing the predicted complex pressure by the measured complex pressure. With the calibrated microphones, we use the multi-microphone method to measure R_d during the hot experiments. The thermoacoustic effect is active and as yet unquantified, meaning that only the microphones between the heater and the downstream end of the tube can be used. We find that $\mathrm{Abs}(R_d)$ changes by less than 0.01 and $\mathrm{Angle}(R_d)$ changes by less than 0.04 radians across all the hot experiments. We find later that, when assimilating k_{eu_h}, there is negligible difference between using R_d from the multi-microphone method and R_d from the 0 Watt experiments.

Having carefully obtained priors for the reflection coefficients and the visco-thermal drag coefficients of the prongs, heater, and thermocouples, we now assume that these are fixed across all experiments and assimilate k_{eu_h}, the heat release rate at the heater as a function of the acoustic velocity around the heater. This allows us to inspect the inferred k_{eu_h} and its uncertainty as a function of heater power, Q_h and heater position, X_h. Figure 2 shows the expected values and two standard deviations of $\mathrm{Abs}(k_{eu_h})/Q_h$ and $\mathrm{Angle}(k_{eu_h})$ from each configuration independently. As expected, the uncertainty is large when the heater is around the middle of

Table 2. Parameter values assimilated during assimilations 1 and 2: R is the reflection coefficient at both ends, k_{vis_p} and k_{th_p} are the viscous and thermal drag coefficients of the prongs, $V(s_r)$ is the variance of the growth rate measurements and $V(s_i)$ is the variance of the frequency measurements

	Re(R)	Im(R)	$\mathrm{Re(k_{vis_p})}$	$\mathrm{Im(k_{vis_p})}$	$\mathrm{Re(k_{th_p})}\times10^{7}$	$\mathrm{Im(k_{th_p})}\times10^{6}$	$V(s_r)$	$V(s_i)$
1	-0.9754 ± 0.0008	$+0.0913 \pm 0.0056$	–	–	–	–	0.01	1.00
2	-0.9755 ± 0.0002	$+0.0900 \pm 0.0015$	$+0.102 \pm 0.009$	$+0.887 \pm 0.109$	$+3.18 \pm 0.275$	-2.99 ± 0.329	0.070	0.841

Table 3. As for table 2 but for k_{vis_h} and k_{th_h}, which are the viscous and thermal drag coefficients of the heater.

	Re(R)	Im(R)	$\mathrm{Re(k_{vis_h})}$	$\mathrm{Im(k_{vis_h})}$	$\mathrm{Re(k_{th_h})}\times10^{6}$	$\mathrm{Im(k_{th_h})}\times10^{5}$	$V(s_r)$	$V(s_i)$
3	-0.9755 ± 0.0001	$+0.0900 \pm 0.0005$	$+3.80 \pm 0.0623$	$+22.5 \pm 0.433$	$+1.95 \pm 0.288$	-2.74 ± 0.201	0.116	0.813
4	-0.9756 ± 0.0001	$+0.0902 \pm 0.0005$	$+3.88 \pm 0.103$	$+22.4 \pm 0.428$	$+2.18 \pm 0.490$	-2.82 ± 0.204	0.216	0.888

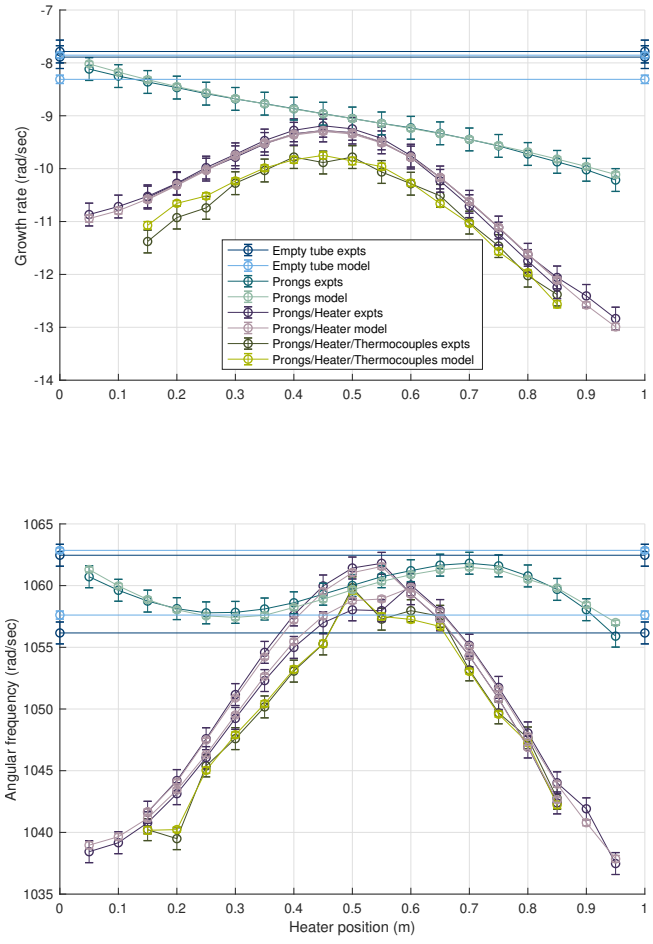

Figure 1. Model predictions and experimental measurements of the growth rate (top) and frequency (bottom) for the C1, C2, C3 and 0 Watt H experiments in table 1 after assimilation of the visco-thermal drag coefficients of the prongs, heater, and thermocouples. Error bars show two standard deviations.

the tube, because it is then located at a velocity node and the heat release rate barely fluctuates. Also as expected, the uncertainty is large when the heater power is low, because the thermoacoustic effect is weak and difficult to observe. In regions in which the thermoacoustic effect is strong, this method successfully assimilates k_{eu_h} with small uncertainty. The values vary smoothly, particularly at high heater powers, even though all assimilations are independent of each other and the experiments were done on different days.

Before assimilating all data into candidate models, it is worth interpreting $\mathrm{Abs}(k_{eu_h})$ and $\mathrm{Angle}(k_{eu_h})$ through a simple time delay, through King's law[14], and through Lighthill's model[15]. If the heat release rate at the heater is proportional to the velocity at the heater some time, τ, earlier such that $q_h' = nu'(t-\tau)$ then we calculate τ with:

$$\tau \equiv \frac{\mathrm{Angle}(k_{eu_h})}{\omega} \tag{8}$$

where $\omega = s_i$ is the angular frequency of the acoustic oscillations. τ is plotted as a function of Q_h and X_h in figure 3. It is nearly constant across all configurations.

King's law predicts that $\mathrm{Abs}(k_{eu_h})/Q_{King}$ should be a constant, where

$$Q_{King} \equiv \frac{1}{(\pi\ \mathrm{Re\ Pr}\ /2)^{-0.5}+2}\,\frac{Q_h}{U_h} \tag{9}$$

Figure 4 shows $\mathrm{Abs}(k_{eu_h})/Q_{King}$ for all configurations. There is excellent consistency as the heater power varies but a significant linear dependence on heater position, which is unlikely to be physical. This could be caused artificially by a systematic error in U_h, which is inferred with the Kalman filter from the long timescale model.

At these Reynolds numbers (Re ~ 20), Lighthill's model[15] predicts that τ/τ_{Light} should be 1, where $\tau_{Light} \equiv 0.2d_{wire}/U_h$. Figure 5 shows τ/τ_{Light} across all the configurations. There is a significant dependence on heater position, which is unlikely to be physical, and significant deviation from 1. This could also be caused by a systematic error in U_h, although such a large deviation is not consistent with our experimental observations. This process highlights the importance of measuring U_h or U_{in} if possible. In our experiment we tried this with a hot wire and with shadowgraphy but could not achieve sufficient accuracy. Future experiments will impose U_{in} with forced convection.

Finally we impose candidate models $\mathcal{H}_i$ and assimilate their parameters for all the configurations. We fix k_{vis_t} and k_{th_t} but re-assimilate k_{vis_h} and k_{th_h} to accommodate possible changes when the heater is hot. The above parameters are less influential than k_{eu_h}, meaning that the hot measurements could be assimilated either by a small adjustment in k_{eu_h} or a large adjustment in k_{vis_h} or k_{th_h}. Having obtained good priors for k_{vis_h} and k_{th_h} with the cold experiments, it is important not to allow them to shift far from their cold values. We achieve this by using their posterior covariances from the cold experiments as the prior covariances of the hot experiment. As before, we re-assimilate the measurement noise.

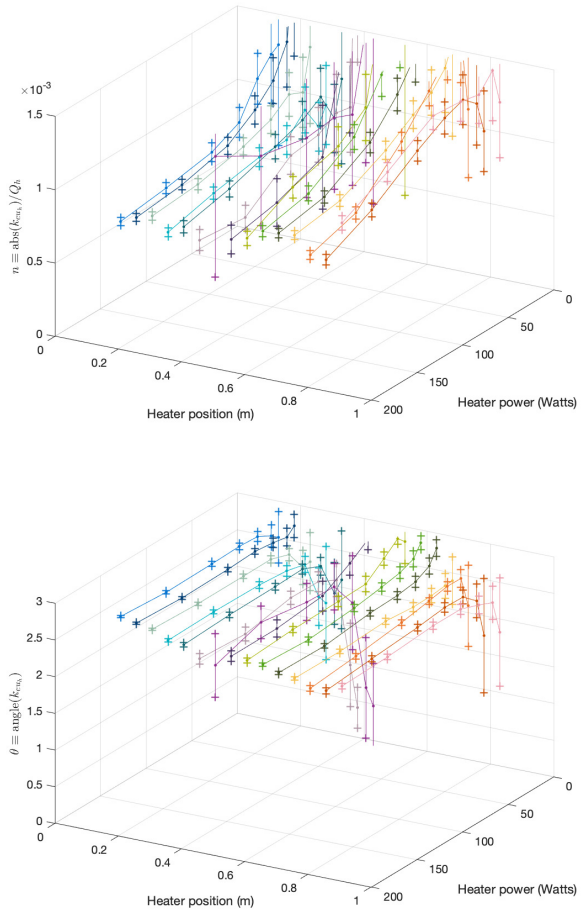

Figure 2. The heat release rate is modelled as feedback from the fluctuating velocity into the energy equation with fluctuating heat release $q'_h = u'_h \times k_{eu_h}$. This figure shows Abs$(k_{eu_h})/Q_h$ and Angle(k_{eu_h}) assimilated for each configuration individually.

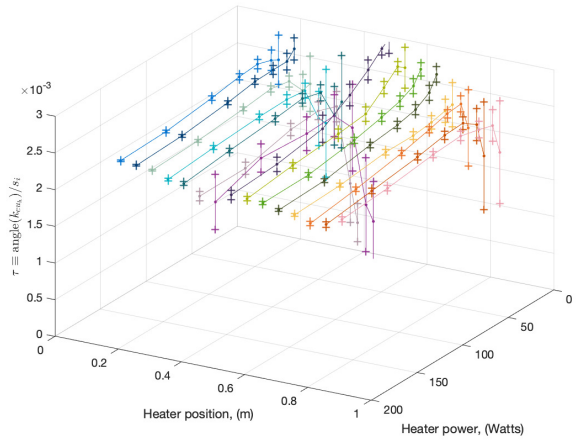

Figure 3. The same data as figure 2(bottom) but with $q'_h = \text{Abs}(k_{eh_h}) \times u'_h(t-\tau)$, where $\tau = \text{Angle}(k_{eu_h})/s_i$

Figure 6 compares the modelled and measured growth rates and frequencies for model A, in which $q'_h = n \times Q_h \times u'(t-\tau)$. The model parameters are n and τ. Figure 7 compares the modelled and measured results for model

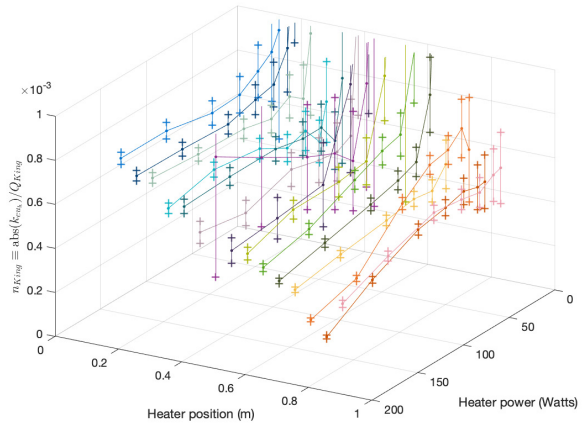

Figure 4. The same data as figure 2(top) but plotting Abs$(k_{eu_h})/Q_{King}$ from equation 9

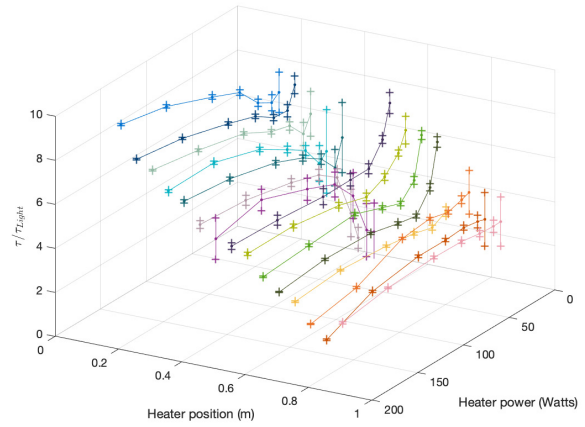

Figure 5. The same data as figure 2(bottom) but with $q'_h = \text{Abs}(k_{eh_h}) \times u'_h(t-\tau_{Light})$.

B, in which $q'_h = n_{King} \times Q_{King} \times u'(t-\tau)$. The model parameters are n_{King} and τ. Figure 8 compares the modelled and measured results for model C, in which $q'_h = n_{King} \times Q_{King} \times u'(t-f\tau_{Light})$. The model parameters are n_{King} and f. Table 4 shows the parameter values and the log marginal likelihood per datapoint for each model. This table includes model D, in which $q'_h = n \times Q_h \times \exp(\mathrm{i}\theta)$ where n and θ are model parameters, and model E, in which $q'_h = n_{King} \times Q_h \times \exp(\mathrm{i}\theta)$ where n_{King} and θ are model parameters. The table includes the log best fit likelihood per datapoint, which measures how well the model fits the data (less negative implies better fit), and the log Occam Factor per datapoint, which measures how much the model has learned from the data (more negative implies better learning). The Occam Factor is the ratio of the posterior accessible volume of $\mathcal{H}_i$'s parameter space to the prior accessible volume, *i.e.* the factor by which $\mathcal{H}_i$'s hypothesis space collapses when the data arrive[5]. As expected, the most suitable models learn most from the data.

Assuming that the models were equally probable before the experiment, the evidence for each model is given by

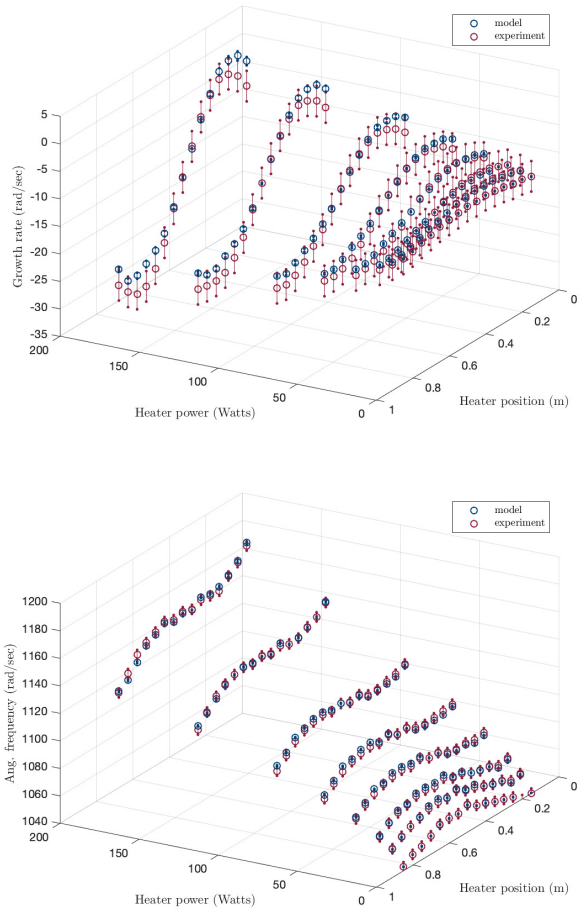

Figure 6. Predictions vs. measurements for model A

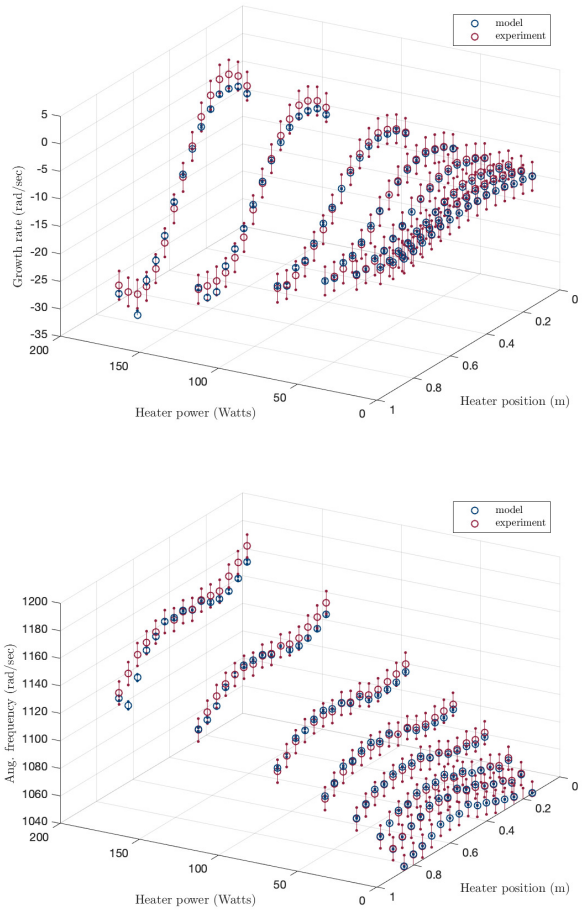

Figure 7. Predictions vs. measurements for model B

the Marginal Likelihood. This process therefore ranks the models in order of suitability as: A, D, B, E, C. Model C can only fit the data with an un-realistically large measurement error in the frequency. This indicates that the simple $n-\tau$ law (A) is best, with a time delay of 2.26 ms. Figure 6 shows that it is quantitatively accurate over the range investigated experimentally. Note that it would be possible to improve the growth rate fit, if desired, by reducing the uncertainty in the growth rate measurement. This would, of course, worsen the frequency fit. King's law with a simple time delay (B) performs reasonably well but, as we saw earlier, could be being let down by systematic error in U_h. The model with Lighthill's time delay (C) performs badly.

Conclusion

In this paper we propose several physics-based candidate models of a thermoacoustic rig. We automate experiments on the rig and obtain thousands of datapoints at around 170 different configurations. We assume that all uncertainties are normally distributed and find an expression for the likelihood of the parameters, given the data, the model, and prior expected values and covariances of the parameters. We assimilate the parameters of each candidate model with first order gradient-based optimization, using first order adjoints. We then obtain the uncertainties of the parameters with Laplace's method, using second order adjoints. We integrate the posterior likelihood in order to obtain the Marginal Likelihood, which quantifies the evidence for each model. We then rank each model and conclude that a simple $n-\tau$ model performs best, with a time delay $\tau = 2.26 \pm 0.02$ ms

The method demonstrated here uses experimetal data to turn qualitatively-accurate models into quantitatively-accurate models with known uncertainty bounds. It provides a rigorous and quantitative method to rank candidate models, based on Bayesian inference. This method can be applied generally. The most challenging procedure is deriving the first and second derivatives of the model predictions with respect to all the parameters. Here, manually-derived first and second order adjoint methods are used. With recent software developments, however, automatic differentiation could also work. In future work we will apply the same method to a flame in a tube.

Acknowledgements

The experimental rig was built by Nick Jamieson, George Rigas, and Francesco Garita. The experiments in this paper were performed by Matthew Yoko. The results were analysed and written up by Matthew Juniper. Matthew Yoko acknowledges financial support from the Cambridge Trusts, the Skye Foundation, and the Oppenheimer Memorial Trust.

Table 4. Characteristics of five models, $\mathcal{H}_i$, for the heat release rate, q'_h, as a function of the acoustic velocity, $u'(t)$, at the heater. The second column contains the models with their assimilated parameter values. During the assimilation, k_{vis_h} and k_{th_h} are allowed to drift from their prior expected values, with tight prior covariances; columns 3 and 4 show their posterior values. $V(s_r)$ is the assimilated measurement uncertainty in the growth rate and $V(s_i)$ is that in the frequency. ML is the marginal likelihood (6), BFL is the best fit likelihood, and OF is the Occam Factor[5].

$\mathcal{H}_i$	$q'_h \times 10^3$	k_{vis_h}	$k_{th_h} \times 10^6$	$V(s_r)$	$V(s_i)$	log(ML)	log(BFL)	log(OF)
A	$0.904 \times Q_h \times u'(t - 0.00226)$	$3.94 + 22.5\mathrm{i}$	$2.39 - 26.7\mathrm{i}$	1.92	2.81	-3.85	-3.68	-0.17
B	$0.422 \times Q_{King} \times u'(t - 0.00223)$	$3.85 + 22.2\mathrm{i}$	$2.09 - 28.7\mathrm{i}$	1.68	17.28	-4.67	-4.52	-0.15
C	$0.292 \times Q_{King} \times u'(t - \tau_{Light} \times 2.06)$	$3.85 + 22.4\mathrm{i}$	$2.04 - 28.2\mathrm{i}$	2.45	128.35	-5.13	-5.00	-0.14
D	$0.901 \times Q_h \times \exp(2.58\mathrm{i})$	$3.93 + 22.4\mathrm{i}$	$2.35 - 26.4\mathrm{i}$	2.64	2.20	-3.89	-3.72	-0.17
E	$0.427 \times Q_{King} \times \exp(2.54\mathrm{i})$	$3.86 + 22.1\mathrm{i}$	$2.11 - 28.7\mathrm{i}$	1.76	17.17	-4.69	-4.54	-0.15

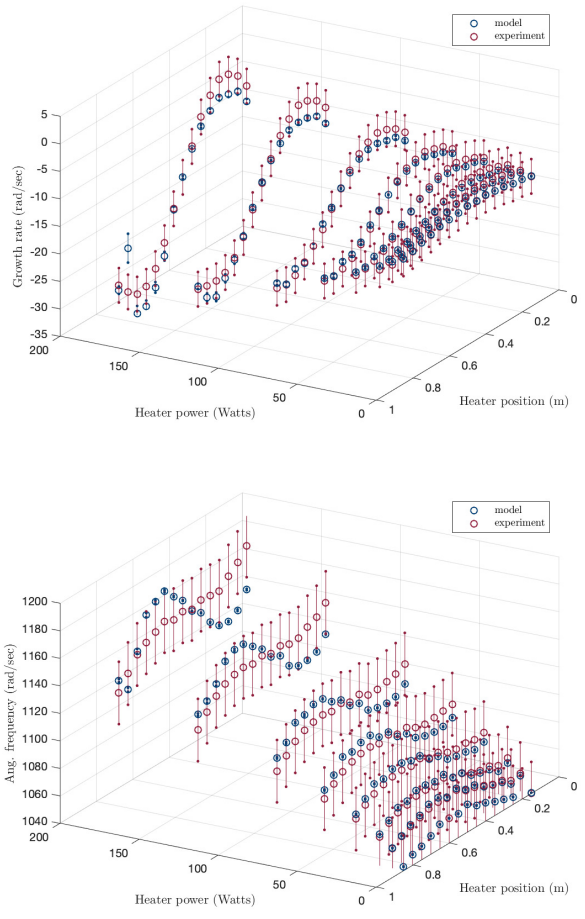

Figure 8. Predictions vs. measurements for model C

References

1. Mongia HC, Held TJ, Hsiao GC et al. Challenges and Progress in Controlling Dynamics in Gas Turbine Combustors Introduction. *Journal of Propulsion and Power* 2003; 19(5): 822–829.
2. Juniper MP and Sujith RI. Sensitivity and Nonlinearity of Thermoacoustic Oscillations. *Annual Review of Fluid Mechanics* 2018; 50: 661–689.
3. Matveev KI. *Thermoacoustic instabilities in the Rijke tube: Experiments and Modeling*. PhD Thesis, California Institute of Technology, 2003.
4. Lieuwen TC and Banaszuk A. Background Noise Effects on Combustor Stability. *Journal of Propulsion and Power* 2005; 21(1): 25–31.
5. MacKay DJC. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
6. Yu H. Inverse Problems in Thermoacoustics. PhD Thesis, University of Cambridge 2020;
7. Garita F, Yu H and Juniper M. Assimilation of Experimental Data to Create a Quantitatively Accurate Reduced-Order Thermoacoustic Model. *Journal of Engineering for Gas Turbines and Power* 2021; 143(2).
8. Garita F Physics-based Statistical Learning in Thermoacoustics. PhD Thesis, University of Cambridge 2021;
9. Juniper MP Sensitivity analysis of thermoacoustic instability with adjoint Helmholtz solvers. *Physical Review Fluids* 2018; 3: 110509
10. Aguilar J, Magri L and Juniper M. Adjoint-based sensitivity analysis of low-order thermoacoustic networks using a wave-based approach. *Journal of Computational Physics* 2017; 341: 163—-181.
11. Tammisola O, Giannetti F, Citro V et al. Second-order perturbation of global modes and implications for spanwise wavy actuation. *Journal of Fluid Mechanics* 2014; 755: 314–335.
12. Magri L, Bauerheim M and Juniper MP. Stability analysis of thermo-acoustic nonlinear eigenproblems in annular combustors. Part I Sensitivity. *Journal of Computational Physics* 2016; 325: 395–410.
13. Cummings A. High temperature effects on the radiation impedance of an unflanged duct exit. *Journal of Sound and Vibration* 1977; 52(2): 299–304.
14. King LV. 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 A* 1914; 214: 373–432.
15. Lighthill MJ. The response of laminar skin friction and heat transfer to fluctuations in the stream velocity. *Proceedings of the Royal Society of London, Series A* 1954; 224: 1156: 1–23.