JASA ARTICLE

Learning the relationship between operating condition and flame response from acoustic data[a)]

Matthew Yoko,[1] Claus Lahiri,[2] André Fischer,[2] and Matthew P. Juniper[1,b)]
[1]*Department of Engineering, University of Cambridge, Cambridge, United Kingdom*
[2]*Rolls-Royce Deutschland, Dahlewitz, Germany*

ABSTRACT:
Predicting how the thermoacoustic response of a combustor changes across operating conditions is a long-standing challenge because small uncertainties in the flame response lead to large uncertainties in the thermoacoustic response. In this paper, we address this challenge for the Rolls-Royce SCARLET test rig using a combination of Bayesian inference, Gaussian process regression, and information-theoretic experiment design. Starting from a physics-based acoustic network model whose flame parameters are inferred at several operating conditions using Bayesian inference (described in a companion paper), we use Gaussian process regression in Part A to learn how the five parameters of a flame model vary with six operating condition parameters. The resulting Gaussian process model predicts the flame response at unseen conditions with quantified uncertainty, and, when coupled into the acoustic network, produces operating maps of the full thermoacoustic response. In Part B, we use metrics from information theory to identify the small number of forcing frequencies that are most informative about the flame parameters. For the SCARLET rig, three optimally chosen frequencies recover the flame transfer function to the same fidelity as the full dataset of around 20 frequencies forced from both upstream and downstream, reducing the data required by up to 90%. © *2026 Acoustical Society of America.*
https://doi.org/10.1121/10.0043841

(Received 19 December 2025; revised 15 April 2026; accepted 19 April 2026; published online 8 June 2026)
[Editor: Ning Xiang] Pages: 5188–5198

I. INTRODUCTION

Thermoacoustic oscillations arise in combustion systems when unsteady heat release rate interacts constructively with acoustic waves in the combustion chamber, leading to self-sustained pressure fluctuations. In practical combustors, such as gas turbines or rocket engines, this feedback can amplify small perturbations, posing risks of excessive noise, vibration, and even structural failure. From the perspective of acoustics, thermoacoustics extends the familiar modelling framework to include a nonlinear source term associated with fluctuating heat release rate. This "flame response" source term describes how the heat release rate of the flame fluctuates in response to acoustic velocity perturbations. To represent this coupling in a tractable way, it is common to approximate the flame response using a linear transfer function, known as the flame transfer function (FTF).

Many of the challenges faced in thermoacoustics are driven by uncertainty. This could be uncertainty in particular features of a physics-based model, or uncertainty in the underlying physics entirely. When a system is sensitive to its inputs, as is the case in thermoacoustics (Juniper and Sujith, 2018), small uncertainties in the inputs drive large uncertainties in the outputs. This makes it difficult to accurately model such systems, unless we have near-perfect precision for all inputs.

A striking example of this is the electrically heated Rijke tube, which is a canonical example of a system that displays thermoacoustic oscillations. The Rijke tube consists of an open-ended tube with an electric heater placed inside. Acoustic velocity perturbations modulate the heat release rate at the heater, amplifying or damping the acoustic oscillations depending on the position of the heater in the tube. Although the system appears simple, it has proven extremely challenging to model accurately. In one example, a detailed physics-based model was constructed, and the model parameters were carefully hand-tuned to match experimental data at one operating condition. The model predictions, however, quickly became inaccurate as the operating condition was changed (Matveev, 2003).

Probability theory offers a natural solution to these challenges by providing tools to quantify and reduce our uncertainties. We have previously demonstrated a Bayesian framework for assimilating data into physics-based models of thermoacoustic oscillations (Juniper and Yoko, 2022; Yoko and Juniper, 2024a). The framework allows us to (1) quantify uncertainty in the model parameters, (2) propagate uncertainty from the parameters to the model predictions, (3) reduce these uncertainties by learning from sparse, noisy data, and (4) calculate the likelihood of a given model, allowing us to quantitatively rank models.

[a)]This paper is a special issue on Bayesian Methods in Acoustical Science and Engineering
[b)]Email: mpj1001@cam.ac.uk

 © 2026 Acoustical Society of America

https://doi.org/10.1121/10.0043841

In a companion paper, we apply this framework to the Rolls-Royce SCARLET thermoacoustic test rig (Yoko *et al.*, 2026). We construct a physics-based model of the rig by (1) considering several candidate sub-models for each rig component, (2) assimilating data into these models to learn the model parameters, (3) using Bayesian model comparison to rank the candidates, and (4) selecting those with the most evidence, given the data. We show that the resulting model precisely reconstructs the measured acoustic response of the rig, while also providing access to unmeasured quantities such as the FTF. Using this framework, we can therefore obtain a calibrated, low-order model for the flame response at any single operating condition.

In this paper, we address the next question: how does the thermoacoustic response change across operating conditions? This is precisely the question that Matveev (2003) attempted to answer for the Rijke tube. We tackle this in two parts. In the first part, we use the companion paper's framework to infer the flame response at 75 operating conditions, producing a dataset of operating conditions and corresponding flame responses. Because we lack a physics-based model for how the flame response depends on operating condition, we use Gaussian process (GP) regression to learn this mapping from the data. The resulting model can predict the flame response at unseen operating conditions, enabling detailed studies into the influence of each operating parameter on the flame and acoustic responses.

Training the GP model requires flame response data at many operating conditions. At each operating condition, the current experimental practice involves around 40 separate forcing tests (20 frequencies in two directions), which is costly for an industrial test rig such as SCaled Acoustic Rig for Low Emission Technology (SCARLET). In the second part of this paper, we investigate an information-based experiment design framework to minimize the number of forcing frequencies required to characterize the flame response at each operating condition.

II. EXPERIMENTAL CONFIGURATION

The Rolls-Royce SCARLET test rig is used to measure the response of an industrially relevant combustion chamber to acoustic forcing. The rig, illustrated in Fig. 1, consists of four key components: the combustion chamber, measurement sections, sirens, and dampers. The combustion chamber is a modular, single can combustor with an easily replaceable fuel spray nozzle. It includes the major features that would be found in a modern aero-engine, such as a bypass annulus and an effusion-cooled liner. On either end of the combustion chamber is a measurement section, which is a circular duct fitted with microphones (high-temperature piezoelectric pressure transducers) at five axial locations. At the end of each measurement section, two sirens are coupled to the rig to provide acoustic forcing. Finally, a damper is fitted on either end of the sirens to minimize the reflection of acoustic waves back into the rig. At the downstream end

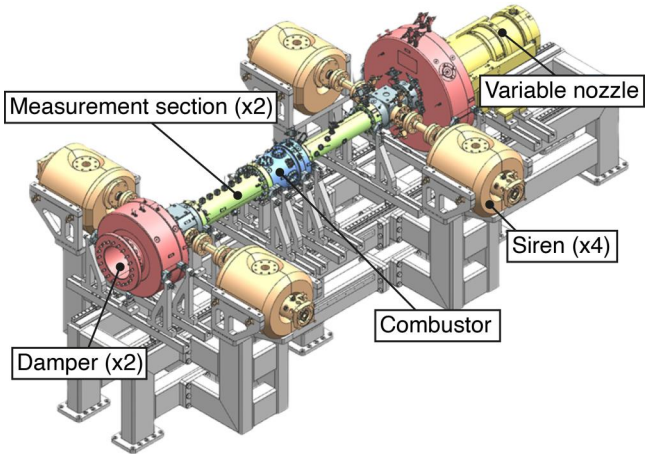

FIG. 1. Diagram of the SCARLET rig highlighting the major components.

of the rig, there is a back-pressure valve, which provides control over the pressure in the rig.

The rig operates at realistic engine conditions, with inlet pressures up to 30 bar and inlet temperatures up to 950 K. The maximum inlet mass flow rate is 4 kg/s, and the combustor pressure drop can be varied between 2% and 5% using the back-pressure valve. At any given operating condition, the rig is forced sequentially in the upstream and downstream directions using the sirens. The sirens produce sinusoidal excitation over a range of amplitudes for frequencies up to 1.2 kHz, which is the theoretical limit for plane wave propagation in the measurement sections at ambient conditions. The excitation amplitude can be varied by shifting the phase relation between the two sirens of each pair. For the experiments relevant to this paper, the velocity perturbations range from 5% to 12% of the mean flow velocity.

The rig provides control over six operating parameters: (1) the inlet mass flow rate, (2) the inlet temperature, (3) the inlet pressure, (4) the pressure drop across the combustor, (5) the global air-fuel ratio (AFR), and (6) the fuel distribution between the pilot and main flames.

III. MODELLING FRAMEWORK

In the companion paper (Yoko *et al.*, 2026), we construct a physics-based model of the SCARLET rig using an acoustic network approach (Dowling and Stow, 2003; Morgans and Yang, 2025; Schuermans *et al.*, 2003). The rig is decomposed into a series of ducts connected by jump conditions that describe the area changes, branch splits and junctions, perforated liner, swirl injector, and flame. This is shown schematically in Fig. 2. Many of the jump conditions introduce model parameters that are not known to sufficient precision *a priori*. Rather than describing these parameters as point estimates, as is standard practice, we describe them as Gaussian probability distributions centered at the most probable value, with a width that quantifies our uncertainty. When data become available, we perform a Bayesian update on these parameter distributions, reducing the uncertainty and improving the precision of our model. Where multiple

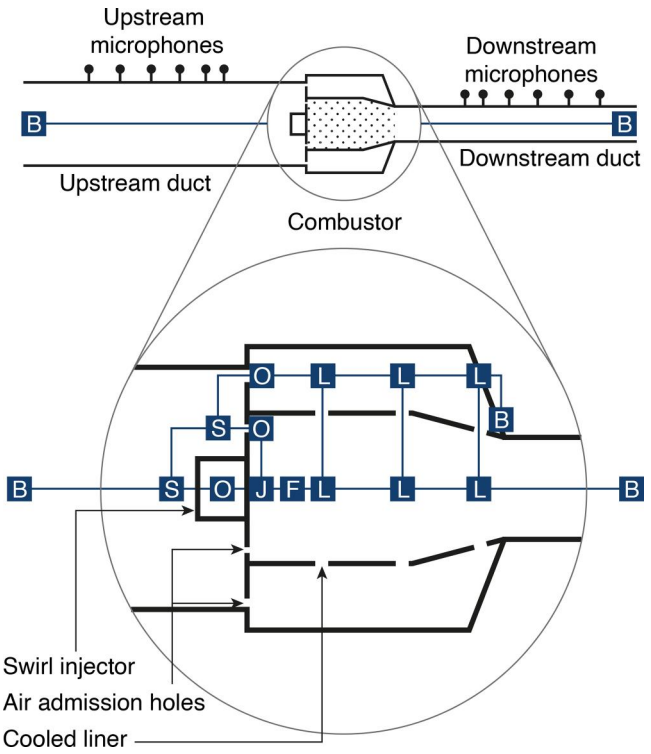

FIG. 2. Schematic of the acoustic network model of the SCARLET rig. The network consists of several ducts (lines) connected by boundary conditions (B) and jump conditions including branch splits (S) and junctions (J), orifices (O), which are just area decreases followed by increases, and special elements for the injector (I), liner (L), and flame (F).

candidate sub-models are available for a given rig component, we use Bayesian model comparison to rank and select the best-supported candidate.

The model parameters can be categorized into acoustic and flame parameters. We infer the acoustic parameters from microphone measurements taken under non-reacting conditions, and show that the resulting acoustic model generalizes to a wide range of operating conditions, allowing us to hold these parameters fixed. Consequently, the only parameters that change with operating condition are those of the flame model, which is described in detail in Sec. III A. At each new operating condition, we infer these flame model parameters from reacting data. Further detail is provided in the companion paper, but for the purpose of this paper, the framework can be viewed as a tool that takes in microphone measurements at a single operating condition, and returns the flame model parameters with quantified uncertainty.

A. Flame model

The flame response to acoustic forcing is represented by a FTF, $\mathcal{F}$, which we model as a sum of discrete time delays (Polifke, 2020) as follows:

$$\mathcal{F} = \frac{q'/\overline{q}}{u'/\overline{u}} = \sum_{i=1}^{N} n_i e^{-s\tau_i}, \tag{1}$$

where q is the heat release rate, u is the velocity, primes denote perturbations, and overbars denote mean quantities. Each term in the sum on the right-hand side contains two unknown parameters: the interaction index, n_i, and the time delay, τ_i, which is multiplied by the complex frequency, s. Additionally, the number of terms, N, is also unknown.

In the companion paper, we use Bayesian model comparison to rank eight candidate flame models: discrete and Gaussian distributed time delays with one to four delays each. These models have been widely used in thermoacoustics (AEsøy *et al.*, 2020; Blondé *et al.*, 2023; Ghani and Albayrak, 2023; Schuermans *et al.*, 2004), where the time delays are typically fitted to data using least squares or similar methods, and the number of delays is typically presumed. We find that, for the burner considered here, a discrete time delay model with three delays ($N = 3$) has the most evidence in the data (Yoko *et al.*, 2026, their Sec. V B 1). We further show that three inferred time delays are physically meaningful: the first is consistent with the convective time scale between the injector dump plane and the flame, the second with vortex shedding from structures in the injector, and the third with circulation fluctuations (Yoko *et al.*, 2026, their Sec. V B 2). Finally, we find that the data support a low-frequency gain approaching zero (Yoko *et al.*, 2026, their Sec. V B 2), which is expected for this type of fuel injector where the fuel mass flow rate is insensitive to acoustic perturbations (Polifke and Lawn, 2007, their Sec. 2.2). In the current paper, we exploit this by constraining the low-frequency limit to zero ($\sum n_i = 0$), which fixes one of the interaction indices and leaves five unknown flame parameters per operating condition: two interaction indices (n_2, n_3) and three time delays (τ_1, τ_2, τ_3). For a detailed discussion of the physical interpretation and validation of the inferred flame responses, the reader is referred to the companion paper.

B. From individual operating conditions to operating maps

Applying the inference procedure described above over a range of operating conditions produces a dataset with six operating condition parameters as inputs and five flame model parameters as outputs. Our goal is to learn the mapping between these, so that we can predict the flame response at operating conditions at which we have not collected data. We lack a physics-based model for how the flame response depends on operating condition, so in Part A of this paper we use GP regression to learn this mapping from the data. A separate but related question is how to minimize the experimental cost of populating this dataset. In Part B, we use information theory to identify the small number of forcing frequencies that contain the most information about the flame parameters, allowing us to characterize each operating condition far more cheaply.

IV. PART A: LEARNING THE MAPPING FROM OPERATING CONDITION TO FLAME RESPONSE

We collect experimental data from a single fuel spray nozzle at 75 different operating conditions and infer the

https://doi.org/10.1121/10.0043841

flame model parameters at each condition to generate the dataset that we use to train the GP model. The 75 operating conditions were accumulated across several experimental campaigns and are therefore irregularly distributed in operating parameter space. This is well suited to GP regression, which does not require data on a regular grid. The range of operating conditions covered by the test matrix is summarized in Table I.

A. GP regression

GP regression is a probabilistic method for learning unknown functions from data (Rasmussen and Williams, 2006). Unlike conventional regression, which assumes a specific functional form (e.g., linear or polynomial), a GP places a probability distribution directly over functions. Given a set of observations, the GP returns a prediction at any new input together with a measure of uncertainty, which narrows near observed data and grows where data are sparse. This makes GPs well suited to our problem, where we have a limited, irregularly sampled dataset, and we want predictions with uncertainty estimates.

The behavior of a GP is governed by two choices: the mean function and the kernel function. The mean function captures any broad trend that we expect in the data, so that predictions far from observed data revert to the broad trend rather than to zero. The kernel function encodes how strongly we expect the function value at one input to be related to its value at another. In this work, we use the squared exponential kernel, which favors smoothly varying functions, and we select a linear mean function to capture first-order trends in the data. Other kernels are available to model periodicity, roughness, or other behaviors, but the squared exponential is a natural first choice for the physical quantities considered here, which are expected to vary smoothly with operating conditions.

The kernel and mean function each introduce a small number of hyperparameters (e.g., length scales, signal variance). As in the Bayesian framework described in the companion paper, we select these by maximizing the marginal likelihood of the model, which balances data fit against model complexity. This allows us to objectively compare different modelling choices and select the simplest configuration that describes the data.

Previous work in thermoacoustics has applied GPs directly to acoustic data, either to infer the flame response (Guo *et al.*, 2021a), as a surrogate model for uncertainty quantification (Guo *et al.*, 2021b; Kulkarni *et al.*, 2021), or for optimization and control (Dharmaputra *et al.*, 2024; Reumschüssel *et al.*, 2024). Unlike previous work, we do not use the GP for flame response identification or uncertainty estimation, as these are handled by the physics-based inference framework (Yoko *et al.*, 2026). Instead, we use the GP to bridge a gap in our physical understanding: how the flame model parameters depend on operating condition. Specifically, we train the GP to predict the five flame parameters from the six operating condition parameters.

TABLE I. Summary of the range of operating condition parameters in the test matrix. The parameters are inlet mass flow rate ($\dot{m}$), temperature (T), AFR, pilot fraction (PF), pressure (P), and pressure drop (ΔP). Note that in this table we present the thermal power (P_{th}) as a surrogate for AFR.

	$\dot{m}$ (kg/s)	T (K)	P_{th} (MW)	PF (%)	P (bar)	ΔP (%)
Minimum	1.6	770	1.7	0	10	3
Maximum	4.0	852	5.6	40	30	4.5

B. Training, validation, and testing

When training machine learning models, the dataset is typically split into training, validation, and test sets. The validation set is used to tune model hyperparameters, while the test set is withheld from both training and validation, and used to provide an unbiased measure of generalization. Because our dataset is relatively small, we adopt leave-one-out cross-validation, where each training point is omitted in turn, the model is retrained on the remaining points, and the prediction error at the omitted point is recorded. The hyperparameters are then tuned to minimize the prediction error at the omitted points. This approach allows efficient use of limited data for both training and validation. We therefore allocate 65 points for training/validation, and reserve 10 points as a test set. The allocation of data to the two groups is done randomly.

Figure 3 shows a parity plot of the trained GP model, which compares the GP prediction (y axis) to the target parameter values that we are trying to reconstruct (x axis). A perfect GP model would place all points along the 45° parity line. We see that the training and validation data (blue) fall consistently close to the parity line. The majority of the test data (orange) fall within a ±5% deviation of the parity line, meaning that the model generalizes well. Note that the interaction index n_1 does not need to be predicted by the GP because it used to enforce the constraint on the low-frequency limit.

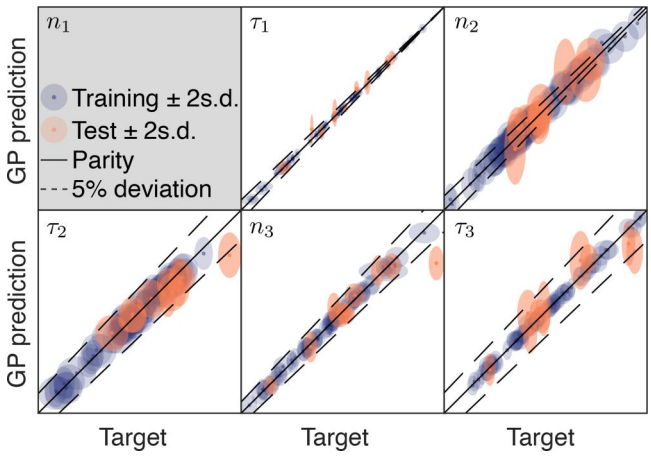

FIG. 3. Parity plot showing the reconstruction performance of the GP model. The target values, which are the parameters we infer from the experimental data, are shown on the x axis. The GP prediction is plotted on the y axis. A perfect GP would place all ellipses along the 45° parity line.

C. Results

Once the GP model is trained on data from a wide range of operating conditions, it can be used to interpolate or extrapolate the FTF at unseen operating conditions, with quantified uncertainties. The closer the operating condition is to previously observed data, the narrower are the uncertainty bounds. This model can then be used to (1) perform detailed studies to determine the influence of each operating variable on the flame response, or (2) interpolate between observed data to generate detailed operating maps of the thermoacoustic system.

To investigate the influence of different operating conditions on the flame response, we select a baseline condition and then sequentially perturb the operating parameters. The baseline condition has a mass flow rate of 3.9 kg/s, an inlet temperature of 825 K, a thermal power of 4.3 MW, a pilot fraction of 20%, an inlet pressure of 29 bar, and a pressure drop over the combustor of 4%. Figure 4 shows how the flame response changes as we sweep the four most influential operating parameters over their full range, while the remaining parameters are held fixed at the baseline condition. We note that the FTFs do not exhibit the low-pass filter behavior that is common to many flames in the literature. This is because, unlike the simple laboratory-scale burners that are typically studied, the SCARLET burner is characterized by high injection velocities and a compact flame, resulting in short convective time scales. Consequently, even relatively high forcing frequencies correspond to low Strouhal numbers based on the convective time scale, and so the low-pass roll-off should be expected to occur at frequencies beyond the limit of plane wave propagation, which is the frequency threshold of these experiments. Similar behavior has been independently reported on SCARLET (Eder *et al.*, 2023, 2024), as well as other industrially relevant flames (Latour *et al.*, 2024; Soundararajan *et al.*, 2022, 2023).

Figure 4(a) shows the effect of sweeping the inlet mass flow rate while holding all other operating parameters constant, including the mixture ratio of the flame. This primarily increases the injection velocity, thereby altering the thermoacoustic time delays. As one would expect, this causes a significant change in the gain, as well as an increase in the rate of phase decay with frequency.

Figure 4(b) shows the effect of sweeping the inlet temperature while holding all other operating parameters constant. Raising the inlet temperature increases the injection velocity slightly, increases the flame speed, and potentially accelerates droplet evaporation. These effects compete, with one acting to lengthen, and the others to shorten the flame. The net result is a steady increase in gain with inlet temperature, and negligible impact on the phase.

Figures 4(c) and 4(d) show the effect of altering the fuel delivery, either through [Fig. 4(c)] the AFR, or [Fig. 4(d)] the fuel distribution between the pilot and main flames. We see in Fig. 4(c) that an increase in global AFR acts to decrease the gain, accompanied by a slight change in the phase. We see a strikingly similar trend in Fig. 4(d), showing that increasing the fraction of fuel delivered to the pilot flame steadily reduces the gain. Increasing the pilot fraction at a fixed global AFR results in an increased local AFR in the mains flame. With the majority of the fuel being supplied to the mains flame, we expect it to dominate the thermoacoustic response, and therefore expect a similarity in these two sweeps.

D. Predicting the thermoacoustic response across operating conditions

So far, the GP model has been used in isolation to predict how the FTF changes with operating condition. We now couple the GP flame model with the physics-based acoustic network model of the rig described in Sec. III. This combined model takes an operating condition as input, uses

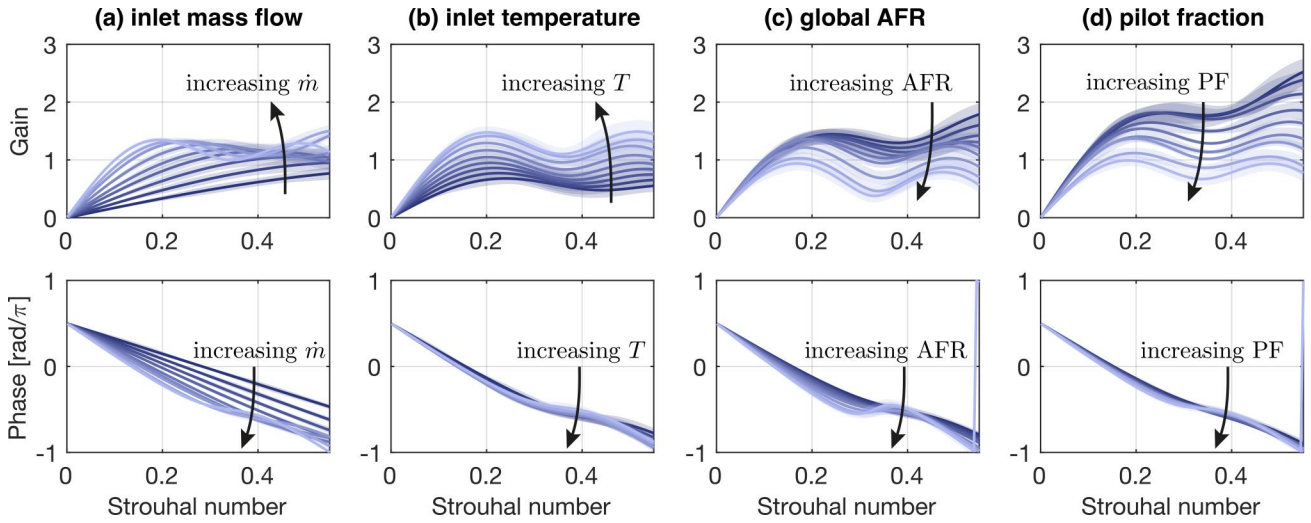

FIG. 4. FTFs predicted by the GP model while sweeping through four test parameters: (a) inlet mass flow rate, (b) inlet temperature, (c) AFR, and (d) pilot fraction. In each sub-plot, the labelled test parameter is swept from the minimum (darkest shade) to maximum (lightest shade) values, while the remaining parameters are held fixed at a baseline value. The parameter ranges are summarized in Table I.

https://doi.org/10.1121/10.0043841

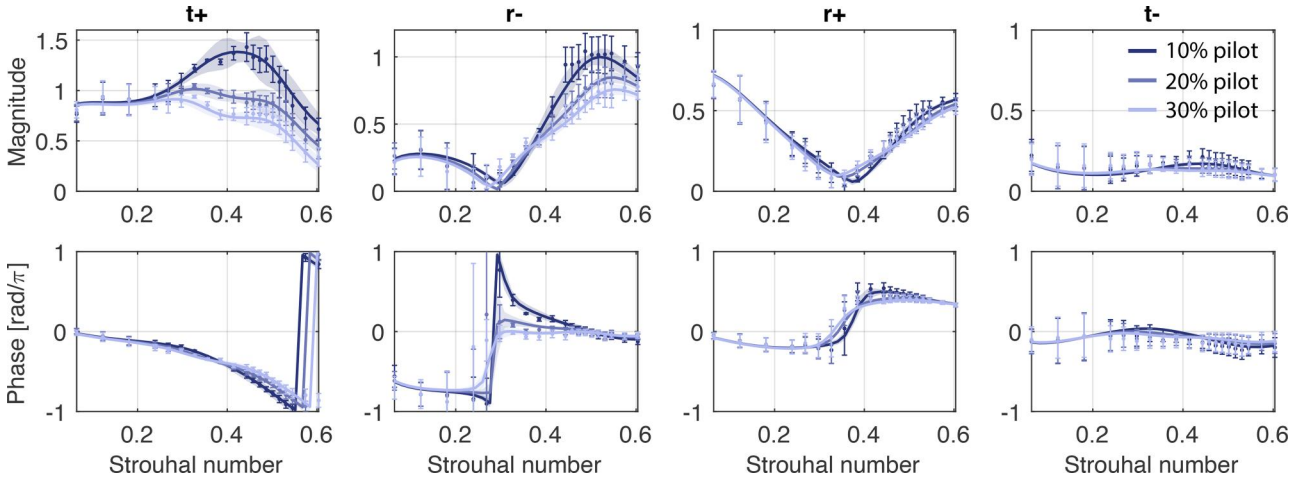

FIG. 5. Scattering matrices of the combustion chamber when the pilot fraction is swept from 10% pilot (darkest shade) to 30% (lightest shade). Experimental measurements are shown as error bars denoting ±2 SD, while the model predictions are shown as solid lines with shaded regions denoting ±2 SD of uncertainty. The elements of the scattering matrix are denoted as r^-, t^-, t^+, and r^+, where r and t denote reflection and transmission coefficients, respectively, and + and – denote the forward and backward travelling waves, respectively.

the GP to predict the flame model parameters, inserts these into the acoustic network, and returns the full thermoacoustic response of the rig. This allows us to (1) make a direct comparison against experimental measurements and (2) explore in detail how the thermoacoustic response of the rig changes with operating condition.

In this section, we use the combustor scattering matrix to represent the acoustic response of the rig. This matrix describes how incoming acoustic waves are reflected and transmitted by the combustor, mapping the forward and backward travelling waves in the upstream and downstream measurement sections onto one another. This includes all elements enclosed by the circle in Fig. 2.

We focus on a sweep of pilot fraction because an equivalent sweep was carried out experimentally, allowing direct comparison with the data. Figure 5 shows the magnitude and phase of all four elements of the combustor scattering matrix at three pilot fractions where experimental data are available (10%, 20%, and 30%). At all three conditions, the model recovers the experimental measurements over the full frequency range. This is expected, since the GP is trained on flame parameters that were themselves inferred from these data, but it confirms that the GP introduces no significant loss of fidelity in the process.

Importantly, the combined model is not restricted to conditions at which data were collected. In Fig. 6, we use this model to produce an operating map of the forward wave transmission coefficient magnitude, $|t^+|$, as the pilot fraction is swept continuously from 0% to 40%. The forward wave transmission coefficient describes how incoming acoustic waves are transmitted through the combustor. When $|t^+| > 1$, the combustor locally amplifies forward-travelling waves. This indicates net conversion of thermal energy into acoustic energy, presenting a potential for thermoacoustic instability if sufficient acoustic energy is fed back with an appropriate phase by the boundaries (note that this will not

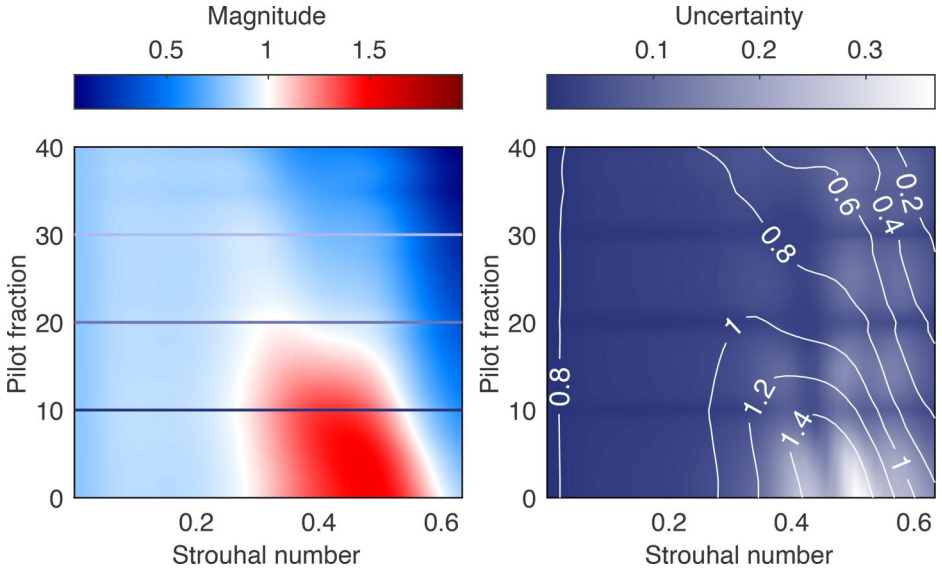

FIG. 6. Operating map of the forward wave transmission coefficient magnitude, $|t^+|$, as the pilot fraction is swept from 0% to 40%. Left: Heat map of the mean predicted value of $|t^+|$. The three pilot fractions where data are available are shown as horizontal lines, with colors corresponding to those in Fig. 5. Right: Heat map of uncertainty (in standard deviations) in the predicted value. To assist with orientation to the magnitude plot, contours of the magnitude from the left panel are show in white on the right panel.

be the case in SCARLET, which is terminated with dampers to avoid global instability, but may be the case in an engine where the combustor is placed between choked boundaries). We see from the left half of Fig. 6 that, as the pilot fraction reduces below about 20%, $|t^+|$ exceeds unity in a narrow band of frequencies centred near $\mathrm{St} \approx 0.45$, indicating elevated risk of instability in this range.

The right half of Fig. 6 shows the uncertainty in the transmission coefficient due to the uncertainty in the GP flame model prediction. Three distinct horizontal bands of low uncertainty are visible, corresponding to the three pilot fractions where experimental data were available, and the flame model confidence is therefore highest. The uncertainty is largest where the model is most sensitive to the flame parameters, which coincides with the regions in which the wave amplification is strongest. Together, the mean prediction and the uncertainty map allow a designer to assess instability risk across the operating envelope. Regions where the model predicts wave amplification or attenuation with high confidence can guide design or operating decisions. Regions where the model is uncertain signal that caution is required until further data are collected to reduce the uncertainty, allowing informed decisions to be made.

E. Summary

We have demonstrated the use of a GP regression model to learn the mapping between the operating condition in a test rig and the parameters of a flame response model. The resulting GP model can predict the flame response at unseen operating conditions, with quantified uncertainty, allowing us to explore the influence of operating conditions at a level of detail that would be too costly to perform experimentally.

Applying this process to a rig such as SCARLET has potential to provide valuable insight into designers early in the design phase. The challenge is that, being a physics-agnostic method, the GP model requires training data at many operating conditions. At each operating condition, the flame response must be characterized experimentally, which is costly for an industrial test rig like SCARLET. Any reduction in the experimental cost of characterizing the flame response at each operating condition would allow more operating conditions to be tested within a given budget, which is the subject of the following section.

V. PART B: INFORMATION-BASED OPTIMAL EXPERIMENT DESIGN IN THERMOACOUSTICS

We now return to the problem addressed in the companion paper: learning the flame model parameters at a single operating condition by assimilating microphone measurements into the physics-based model of the rig. In the companion paper, the experimental data at each operating condition were collected by forcing the rig at around 20 frequencies from both the upstream and downstream directions, producing densely sampled FTFs that can be interpreted directly by a human. However, since the physics-based model relates all of these measurements to just five flame parameters, much of these data are redundant for the purpose of parameter learning.

If we replace the requirement for human-interpretable data with the goal of learning the flame parameters as efficiently as possible, we can be far more selective about which experiments to conduct. To do this, we use tools from information theory to identify the small number of forcing frequencies that contain the most information about the unknown flame parameters, allowing us to characterize each operating condition far more cheaply. The cost savings can then be reinvested in testing more operating conditions, providing the volume of training data required for the GP model demonstrated in Part A.

In this part of the paper, any reference to "model" refers to the physics-based model of the rig, as described in Sec. III in brief and the companion paper in detail (Yoko *et al.*, 2026, their Appendix A). A reference to "unknown parameters" refers to the flame model parameters described in Sec. III A, which we are trying to infer from the data. Finally, a reference to "experiment" refers to a test on the SCARLET rig at a particular forcing frequency and direction, which produces a set of microphone measurements.

A. Information content of an experiment

Not all experiments are equally informative. An experiment whose outcome we can already predict with confidence tells us little, while an experiment whose outcome is uncertain has the potential to teach us a great deal. This intuition can be made precise using Shannon entropy (Shannon, 1948), which quantifies the total uncertainty in a probability distribution. Consider the probability distribution that describes our state of knowledge about the model parameters, $\boldsymbol{a}$, after we have assimilated a dataset containing i observations, $\boldsymbol{D}_i$. We denote this posterior distribution $p(\boldsymbol{a}|\boldsymbol{D}_i)$. The Shannon entropy of this distribution is given by the following:

$$S_i = -\int_{\boldsymbol{A}} p(\boldsymbol{a}|\boldsymbol{D}_i) \log_2 p(\boldsymbol{a}|\boldsymbol{D}_i)\, \mathrm{d}\boldsymbol{a}, \tag{2}$$

where $\boldsymbol{A}$ is the set of values that the parameters can take with non-zero probability and the base-two logarithm gives information in bits, where one bit halves the uncertainty. Note that if no data have been collected yet, $\boldsymbol{D}_i$ is empty and $p(\boldsymbol{a}|\boldsymbol{D}_i)$ is equal to the prior distribution, $p(\boldsymbol{a})$. The information gained from a subsequent observation $\boldsymbol{d}_{i+1}$ is the reduction in entropy as follows:

$$\mathcal{I} = S_i - S_{i+1}. \tag{3}$$

In our inference framework, we use a Laplace approximation, so the posterior is approximated as a multivariate Gaussian with covariance matrix $\boldsymbol{C}_{a|D_i}$ (see Yoko *et al.* 2026, their Sec. IV for further detail). The Shannon entropy of a multivariate Gaussian is

https://doi.org/10.1121/10.0043841

$$S_i = \frac{1}{2}\log_2|\boldsymbol{C}_{a|D_i}| + \text{const.}, \tag{4}$$

where the value of const. does not change with subsequent observations, such that it cancels in the information gain. When applying the Laplace approximation, a subsequent observation at frequency f updates the inverse covariance matrix as follows:

$$\boldsymbol{C}^{-1}_{a|D_{i+1}} = \boldsymbol{C}^{-1}_{a|D_i} + \boldsymbol{J}_f^\top \boldsymbol{C}_f^{-1}\boldsymbol{J}_f, \tag{5}$$

where $\boldsymbol{J}_f$ is the Jacobian of the model output with respect to the model parameters, evaluated at frequency f, and $\boldsymbol{C}_f$ is the measurement noise covariance. The information gain from this observation is, then,

$$\mathcal{I}_f = \frac{1}{2}\log_2\frac{|\boldsymbol{C}_{a|D_i}|}{|\boldsymbol{C}_{a|D_{i+1}}|} = \frac{1}{2}\log_2|\boldsymbol{C}_f^{-1}||\boldsymbol{C}_f + \boldsymbol{J}_f\boldsymbol{C}_{a|D_i}\boldsymbol{J}_f^\top|. \tag{6}$$

The crucial feature of this result is that the information gain depends on the current uncertainty in the parameters, $\boldsymbol{C}_{a|D_i}$, the model sensitivity, $\boldsymbol{J}_f$, and the measurement noise, $\boldsymbol{C}_f$, but not on the actual measurement outcome. This means that we can evaluate the utility of any candidate experiment before conducting it (Lindley, 1956; MacKay, 1992). We therefore plan experiments by recursively (1) computing the model sensitivity, $\boldsymbol{J}_f$, at each candidate forcing frequency using adjoint methods or automatic differentiation, (2) using Eq. (6) to compute the expected information gain from an experiment at each candidate frequency, and (3) selecting the frequency with the highest expected information gain to conduct the next experiment. A more detailed discussion of this method in the context of thermoacoustics is provided in Yoko and Juniper (2024b).

B. Results

To demonstrate the optimal experiment design framework, we take an existing dataset in which the frequency was densely sampled, and we sequentially identify and assimilate the most informative forcing frequencies. We repeat this until our model uncertainty falls below experimental uncertainty throughout the frequency range. The results are shown in Fig. 7.

The first four rows in Fig. 7 show the magnitude of the four scattering matrix elements. The phase is omitted for space, but the results are identical. The fifth row plots the amount of information we expect to gain by forcing the rig at any frequency. The four columns represent four states of knowledge. The first column represents the prior state of knowledge, i.e., we have not yet assimilated any data. The second column represents our state of knowledge after assimilating data from a single frequency observation, and columns three and four follow.

Considering Fig. 7 (column a), we see from the scattering matrix elements that our prior knowledge is poor (the model does not fit the data well), and our uncertainty is large. The expected information gain plot shows that the information content steadily increases with forcing frequency up until $\mathrm{St} \approx 0.5$ when it begins to fall again.

There are three interesting features of the expected information gain plot in Fig. 7 (column a, row v). First, we see that the expected information gain falls to exactly zero at the low-frequency limit. This is because we have enforced the low-frequency limit in the flame model, so we have left no uncertainty in the low-frequency behavior, and can therefore learn nothing about it from the data.

Second, the peak of expected information corresponds to the most thermoacoustically active frequency predicted by the model, given the current state of knowledge. This is where we are likely to observe the most extreme deviations between the model and the data, giving us the highest chance of being surprised by the data.

Finally, we see that forcing from the upstream end of the rig provides more information than forcing from the downstream end. This is because of the extremely low transmission coefficient for the backward-travelling wave (t-) for this combustor, which means that little of the downstream excitation penetrates into the upstream measurement section, resulting in an experiment with poor information content. While the difference may seem small, we recall that a single bit halves our uncertainty in the parameters. The difference between upstream and downstream forcing is about three bits, meaning we will become eight times more confident in the parameter values if we force from upstream rather than downstream.

Moving to column Fig. 7 (column b, rows i–iv), we see that, after assimilating the data collected at the first optimal frequency, the model becomes locally accurate and confident. This accuracy and confidence decay quickly with frequency as we move away from the observed data.

From Fig. 7 (column b, row v), we see that the expected information at the observed frequency drops to almost zero, i.e., we do not expect to learn much more by repeating the same experiment. Notably, the information content of the downstream forcing experiment at the same frequency has also dropped to almost zero, even though we only assimilated data from the upstream forcing case. Following an observation from a single forcing direction, the physics-based model can already predict, with high confidence, how the rig will respond to forcing in the other direction. There is therefore little additional information to be gained by forcing in both directions when we assimilate the data into a physics-based model.

Similarly, in Fig. 7 (column c, rows i–iv), we see that, after the second optimal observation, the model becomes locally accurate and confident around the new frequency. We do, however, see a more global reduction in uncertainty at this stage, with even the predictions at lower frequencies becoming more confident.

This is confirmed by the expected information plot in Fig. 7 (column c, row v), which shows a global reduction in expected information right down to the low frequencies. We see, however, that the expected information at untested frequencies is still higher than that at the two observed

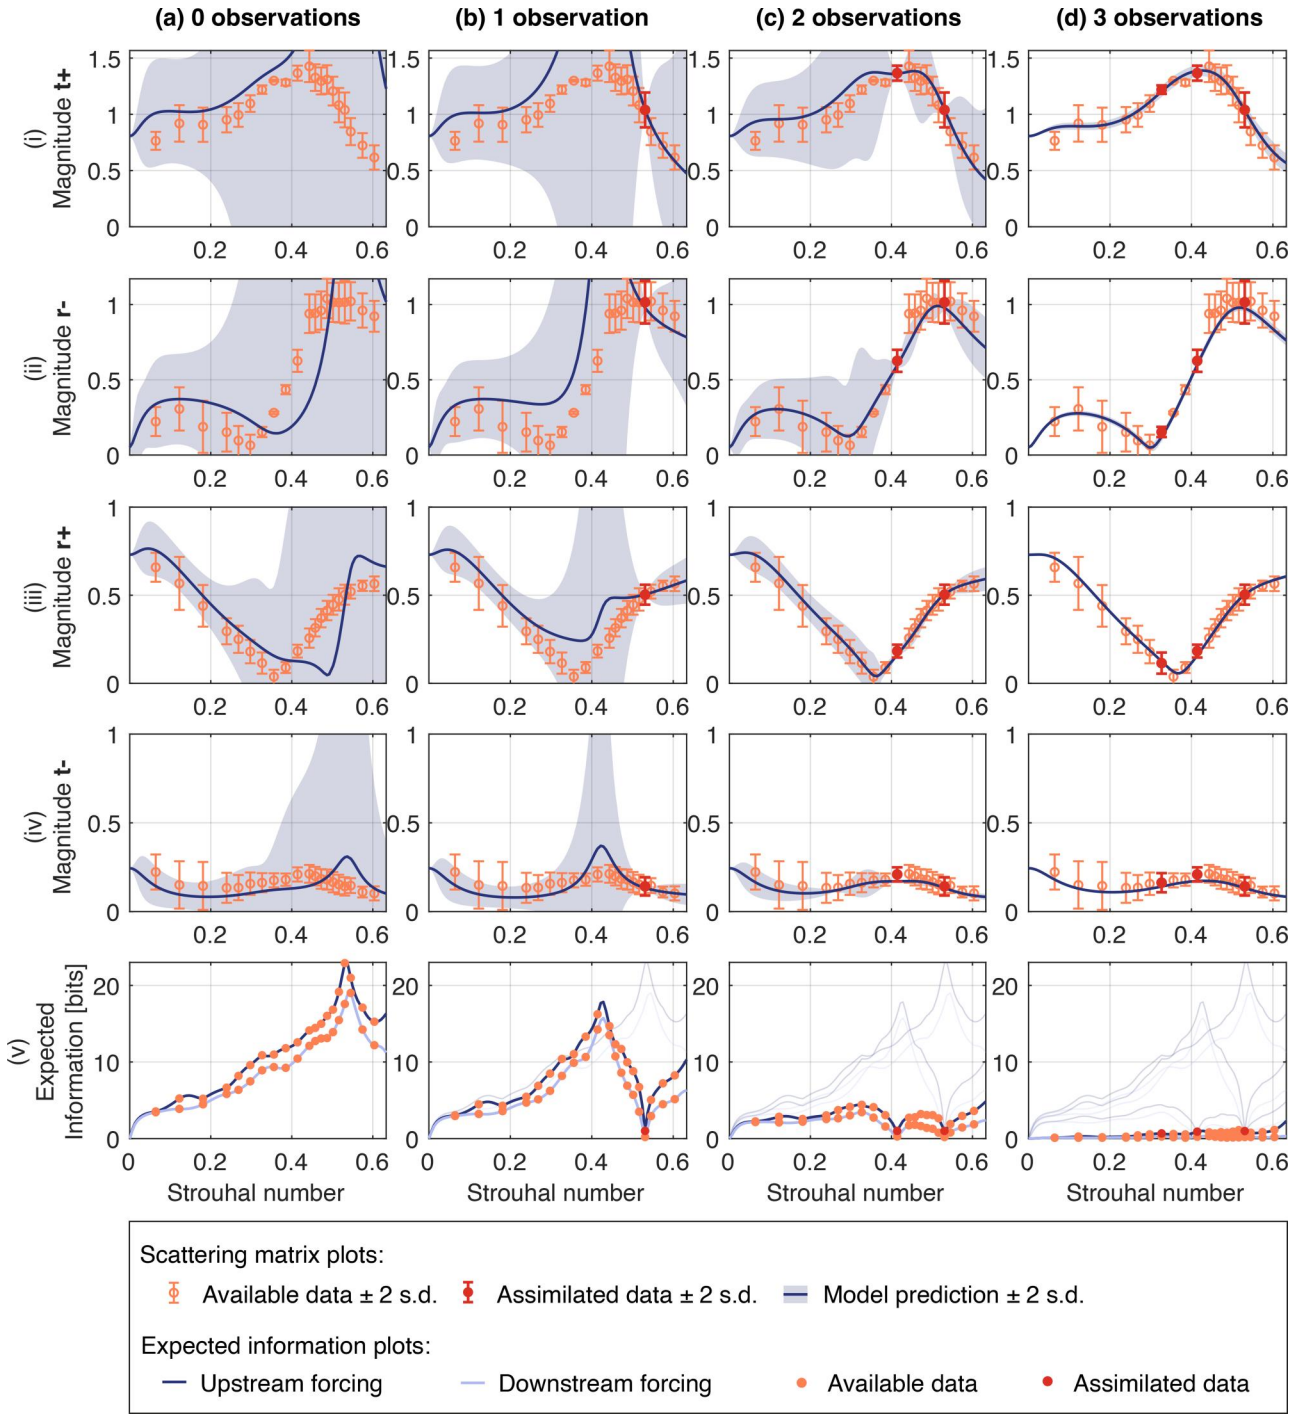

FIG. 7. Using information metrics to select the optimal forcing frequencies. The first four rows plot the magnitude of the scattering matrix elements. The fifth row plots the expected information content of a forcing experiment at any frequency, in either direction. The four columns show the state of knowledge and expected information gain after zero, one, two, and three experimental data points have been observed.

frequencies, so there is still valuable information that we can gain from additional data.

Finally, Fig. 7 (column d, rows i–iv) shows that, after assimilating the data from three optimal frequencies, the model becomes accurate and confident over the entire frequency range. Similarly, in Fig. 7 (column d, row v), we see that the expected information content of future experiments is close to zero, indicating that we can stop collecting data. Additional experiments will only serve to average out noise in the measurements.

In Fig. 8, we show how the estimated FTF changes as we assimilate the three optimal experiments. For reference, the FTF inferred from the full dataset is shown in gray. As with the scattering matrices in Fig. 7, we see that our prior (column a) was poor, significantly overpredicting the gain. As we assimilate more data, our estimate of the FTF

https://doi.org/10.1121/10.0043841

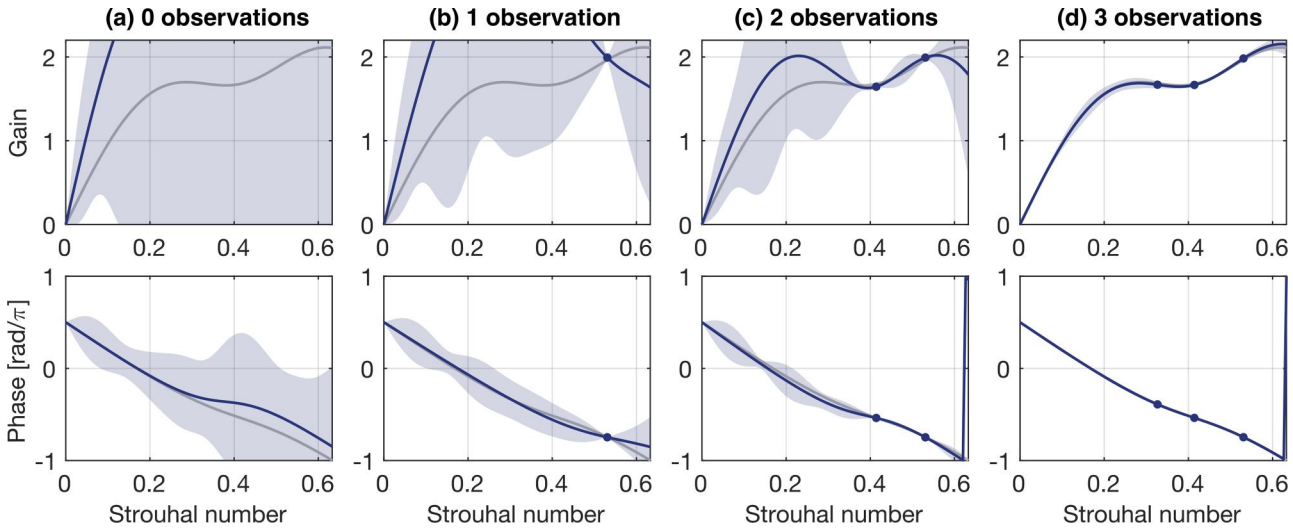

FIG. 8. Comparing the FTF inferred after sequentially assimilating the three optimal observations (blue) to the FTF inferred using all the data (gray).

becomes locally accurate. After assimilating the three optimal experiments, we achieve global accuracy, and the FTF is almost identical to that inferred when all the data are assimilated.

It is not surprising that we require only three experimental observations to obtain a good approximation of the FTF. Our model for the flame response is a distributed time delay model with three time delays, which has six unknown parameters (five if we constrain the low-frequency limit). In each forcing experiment, we measure the magnitude and phase of the acoustic pressure. These data contain information about both the gain and the phase of the FTF at the tested frequency, effectively placing two constraints on the values of our unknown parameters. We therefore only need three of these observations to become fully confident in the FTF model parameters.

This shows that more elaborate models (those with more parameters) require more data. In the companion paper, we use the full frequency range at one operating condition to show that the three time delay model has the highest probability given the data. In this paper, we then use three frequency observations at another operating condition to infer the parameters of the flame model. We therefore needed a large amount of data from one operating condition to choose the model, and then little data to train it on subsequent operating conditions, assuming that the model remains valid.

By using the physics-based model of the rig to identify maximally informative experiments, we have learned the unknown parameters of the model using less than 10% of the data that is typically collected. The same concept can be applied in the GP model that we demonstrated in Part A of this paper. This allows us to identify the operating conditions that would provide maximum information for training the GP model to map between operating condition and FTF parameters. Indeed, when applied to a GP model, our information-theoretic sampling simplifies to the popular maximum variance sampling, which is commonly used for active learning in GPs (Rasmussen and Williams, 2006; Sengupta and Juniper, 2022).

VI. CONCLUSION

In this paper, along with its companion paper (Yoko *et al.*, 2026), we demonstrate how we can combine physics-based modelling, data assimilation, and probabilistic machine learning to efficiently build precise models of complex thermoacoustic systems. In the companion paper, we demonstrate a Bayesian framework that infers the parameters of a flame model at a single operating condition from microphone measurements. In this paper, we extend this to predicting the flame response across operating conditions, as well as reducing the experimental cost of data collection.

In Part A, we use a GP to learn how the five flame model parameters vary across a six-dimensional operating space. The GP model can predict the flame response at unseen conditions with quantified uncertainty, enabling isolated parameter sweeps and detailed operating maps that would be too costly to produce experimentally. By coupling the GP predictions back into the physics-based acoustic network, we produce operating maps of the entire thermoacoustic system that reconstruct the experimental data where they are available, and make predictions with quantified uncertainty where they are not.

In Part B, we show how to reduce the experimental cost of populating the dataset required to train the GP model. Using a closed-form expression for the expected information gain of a candidate experiment, we identify the small number of forcing frequencies that are most informative about the flame parameters at each operating condition. For the SCARLET rig, three optimally chosen frequencies recover the FTF to the same fidelity as the full dataset of around 20 frequencies in two directions, a reduction of 90%. In practice, it would be pragmatic to collect some additional data for redundancy, but even with 100% redundancy, the data required are still reduced by 80%. This reduction makes it

practical to characterize many more operating conditions within a fixed experimental budget, supporting the data requirements of the GP model in Part A.

In future work, we will apply this method to a live test campaign rather than down-sampling an existing dataset. In particular, we will use optimal experiment design to select both the forcing frequencies at each operating condition and the operating conditions themselves, closing the loop between data collection and model training. We will also use the insight gained from the GP model to inform the derivation of a physics-based model for how the flame parameters vary with operating condition, which will further reduce the data requirements.

ACKNOWLEDGMENTS

This work was supported by the European Union, under Grant Agreement No. 101102004 within the HEAVEN project. The test rig was funded with financial support by German Federal Ministry for Economic Affairs and Climate Action as part of the AG-Turbo program "COORETEC–AG Turbo 2020," acknowledged as per resolution of the German Federal Parliament under Grant No. 03ET2012B. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or Clean Aviation Joint Undertaking. Neither the European Union nor Clean Aviation JU can be held responsible for them.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts of interest to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available from Rolls-Royce. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors upon reasonable request and with the permission of Rolls-Royce.

AEsøy, E., Aguilar, J. G., Wiseman, S., Bothien, M. R., Worth, N. A., and Dawson, J. R. (**2020**). "Scaling and prediction of transfer functions in lean premixed," Combust. Flame **215**, 269–282.

Blondé, A., Schuermans, B., Pandey, K., and Noiray, N. (**2023**). "Effect of hydrogen enrichment on transfer matrices of fully and technically premixed swirled flames," J. Eng. Gas Turbines Power **145**(12), 1–11.

Dharmaputra, B., Reckinger, P., Schuermans, B., and Noiray, N. (**2024**). "BOATS: Bayesian optimization for active control of thermoacoustics," J. Sound Vib. **582**, 118415.

Dowling, A. P., and Stow, S. R. (**2003**). "Acoustic analysis of gas turbine combustors," J. Propul. Power **19**(5), 751–764.

Eder, A. J., Fischer, A., Lahiri, C., Merk, M., Staufer, M., Eggels, R., Silva, C. F., and Polifke, W. (**2023**). "Identification of the dynamics of a turbulent spray flame at high pressure," in *Symposium on Thermoacoustics in Combustion*, Zurich, Switzerland.

Eder, A. J., Merk, M., Hollweck, T., Fischer, A., Lahiri, C., Silva, C. F., and Polifke, W. (**2024**). "Model-based inference of flame transfer matrices from acoustic measurements in an aero-engine test rig," J. Eng. Gas Turbines Power **147**, 1–31.

Ghani, A., and Albayrak, A. (**2023**). "From pressure time series data to flame transfer functions: A framework for perfectly premixed swirling flames," J. Eng. Gas Turbines Power **145**(1), 1–9.

Guo, S., Silva, C. F., and Polifke, W. (**2021a**). "Robust identification of flame frequency response via multi-fidelity Gaussian process approach," J. Sound Vib. **502**, 116083.

Guo, S., Silva, C. F., Yong, K. J., and Polifke, W. (**2021b**). "A Gaussian-process-based framework for high-dimensional uncertainty quantification analysis in thermoacoustic instability predictions," Proc. Combust. Inst. **38**(4), 6251–6259.

Juniper, M. P., and Sujith, R. (**2018**). "Sensitivity and nonlinearity of thermoacoustic oscillations," Annu. Rev. Fluid Mech. **50**, 661–689.

Juniper, M. P., and Yoko, M. (**2022**). "Generating a physics-based quantitatively-accurate model of an electrically-heated Rijke tube with Bayesian inference," J. Sound Vib. **535**, 117096.

Kulkarni, S., Guo, S., Silva, C. F., and Polifke, W. (**2021**). "Confidence in flame impulse response estimation from large eddy simulation with uncertain thermal boundary conditions," J. Eng. Gas Turbines Power **143**(12), 121002.

Latour, V., Durox, D., Renaud, A., and Candel, S. (**2024**). "Data-driven model-based instability prediction in an annular combustor relying on a flame response mapping of the operating domain," Combust. Flame **270**, 113782.

Lindley, D. V. (**1956**). "On a measure of the information provided by an experiment," Ann. Math. Statist. **27**(4), 986–1005.

MacKay, D. J. C. (**1992**). "Information-based objective functions for active data selection," Neural Comput. **4**(4), 590–604.

Matveev, K. (**2003**). "Thermoacoustic instabilities in the Rijke tube: Experiments and modeling," Ph.D. thesis, California Institute of Technology, Pasadena, CA.

Morgans, A. S., and Yang, D. (**2025**). "Thermoacoustic instability in combustors," Annu. Rev. Fluid Mech. **57**, 9–33.

Polifke, W. (**2020**). "Modeling and analysis of premixed flame dynamics by means of distributed time delays," Prog. Energy Combust. Sci. **79**, 100845.

Polifke, W., and Lawn, C. (**2007**). "On the low-frequency limit of flame transfer functions," Combust. Flame **151**(3), 437–451.

Rasmussen, C. E., and Williams, C. K. (**2006**). *Gaussian Processes for Machine Learning* (Massachusetts Institute of Technology, Cambridge, MA).

Reumschüssel, J. M., von Saldern, J. G., Ćosić, B., and Paschereit, C. O. (**2024**). "Data-driven optimization of a gas turbine combustor: A Bayesian approach addressing NO_x emissions, lean extinction limits, and thermoacoustic stability," Data-Centric Eng. **5**, e32.

Schuermans, B., Bellucci, V., Guethe, F., Meili, F., Flohr, P., and Paschereit, C. O. (**2004**). "A detailed analysis of thermoacoustic interaction mechanisms in a turbulent premixed flame," in *Proceedings of the ASME Turbo Expo*, pp. 539–551.

Schuermans, B., Bellucci, V., and Paschereit, C. O. (**2003**). "Thermoacoustic modeling and control of multi burner combustion systems," in *Proceedings of the ASME Turbo Expo*, Atlanta, Georgia.

Sengupta, U., and Juniper, M. P. (**2022**). "Thermoacoustic stabilization of combustors with gradient-augmented Bayesian optimization and adjoint models," Int. J. Spray Combust. Dyn. **14**(3), 266–272.

Shannon, C. E. (**1948**). "A mathematical theory of communication," Bell Syst. Tech. J. **27**(4), 623–656.

Soundararajan, P. R., Durox, D., Renaud, A., and Candel, S. (**2023**). "Impact of spray dynamics on combustion instabilities investigated by changing the atomizer recess in a swirl combustor," Combust. Flame **252**, 112757.

Soundararajan, P. R., Durox, D., Renaud, A., and Vignat, G. (**2022**). "Swirler effects on combustion instabilities analyzed with measured FDFs, injector impedances and damping rates," Combust. Flame **238**, 111947.

Yoko, M., and Juniper, M. P. (**2024a**). "Adjoint-accelerated Bayesian inference applied to the thermoacoustic behaviour of a ducted conical flame," J. Fluid Mech. **985**, A38.

Yoko, M., and Juniper, M. P. (**2024b**). "Optimal experiment design with adjoint-accelerated Bayesian inference," Data-Centric Eng. **5**, e17.

Yoko, M., Lahiri, C., Fischer, A., and Juniper, M. P. (**2026**). "In-situ inference of the thermoacoustic properties of an industrial combustion system," J. Acoust. Soc. Am. **159**, 5150–5169.