



# Bayesian inference of physics-based models of acoustically-forced laminar premixed conical flames

Alessandro Giannotta <sup>a,b</sup> , Matthew Yoko <sup>b</sup> , Stefania Cherubini <sup>a</sup>,  
Pietro De Palma <sup>a</sup>, Matthew P. Juniper <sup>b</sup> ,\*

<sup>a</sup> Department of Mechanics, Mathematics and Management, Politecnico di Bari, via Re David, 200, Bari, 70125, Italy

<sup>b</sup> Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, UK

## ARTICLE INFO

Dataset link: <http://doi.org/10.17863/CAM.106960>, [https://github.com/aleGiannotta/Bayesian\\_Inference\\_Lam\\_Flame.git](https://github.com/aleGiannotta/Bayesian_Inference_Lam_Flame.git)

### Keywords:

Thermoacoustics  
Data assimilation  
Laplace's method  
Bayesian inference

## ABSTRACT

We perform twenty experiments on an acoustically-forced laminar premixed Bunsen flame and assimilate high-speed footage of the natural emission into a physics-based model containing seven parameters. The experimental rig is a ducted Bunsen flame supplied by a mixture of methane and ethylene. A high-speed camera captures the natural emission of the flame, from which we extract the position of the flame front. We use Bayesian inference to combine this experimental data with our prior knowledge of this flame's behaviour. This prior knowledge is expressed through (i) a model of the kinematics of a flame front moving through a model of the perturbed velocity field, and (ii) *a priori* estimates of the parameters of the above model with quantified uncertainties. We find the most probable *a posteriori* model parameters using Bayesian parameter inference, and quantify their uncertainties using Laplace's method combined with first-order adjoint methods. This is substantially cheaper than other common Bayesian inference frameworks, such as Markov Chain Monte Carlo. This process results in a quantitatively-accurate physics-based reduced-order model of the acoustically forced Bunsen flame for injection velocities ranging from 1.75 m/s to 2.99 m/s and equivalence ratio values ranging from 1.26 to 1.47, using seven parameters. We use this model to evaluate the heat release rate between experimental snapshots, to extrapolate to different experimental conditions, and to calculate the flame transfer function and its uncertainty for all the flames. Since the proposed model relies on only seven parameters, it can be trained with little data and successfully extrapolates beyond the training dataset. Matlab code is provided so that the reader can apply it to assimilate further flame images into the model.

**Novelty and Significance Statement** Thermoacoustic systems tend to be extremely sensitive to small parameter changes, which makes them difficult to model *a priori* from existing models in the literature. This means, however, that thermoacoustic models tend to be easy to train using data-driven methods because, with well-chosen experiments, their parameters can be easily observed from experimental data. This paper presents a novel use of Bayesian inference to combine experimental measurements, numerical simulations, and prior knowledge about flame behaviour. We outline our methodology and demonstrate its effectiveness using a laminar premixed Bunsen flame. Our approach yields a quantitatively-accurate physics-based model that predicts the expected value and uncertainty bounds of the flame transfer function between velocity and heat release rate perturbations. The proposed model contains only seven physical parameters, which is fewer parameters than non-physics-based models, and can therefore be trained on relatively little data. We also illustrate how the trained model effectively extrapolates beyond the training dataset. Our numerical code and experimental data are open access.

## 1. Introduction

The strength of the mechanism driving thermoacoustic oscillations depends strongly on the phase difference between the heat release rate (h.r.r) and pressure oscillations. This phenomenon was described

qualitatively by Rayleigh [1], and expressed in terms of the rate of acoustic energy increase by Chu [2], where the influence of the phase dependence mentioned above can be deduced from the first term on the right hand side of Eq. (19b). In turn, this phase difference usually depends strongly on the flame's dynamics [3] and on the acoustics of

\* Corresponding author.

E-mail address: [mpj1001@cam.ac.uk](mailto:mpj1001@cam.ac.uk) (M.P. Juniper).

a combustion chamber. This dependence is particularly strong if time delays,  $\tau$ , in the flame are similar to or greater than the period,  $T$ , of acoustic chamber modes [4] because small changes to either  $\tau$  or  $T$  lead to large changes in the phase difference between pressure and h.r.r.. As a result of these strong, coupled dependencies, a combustor's thermoacoustic behaviour tends to be extremely sensitive to changes in the geometry or flame behaviour [5]. For the same reason, the outputs of faithful thermoacoustic models are sensitive to small changes in models and their parameters. On the negative side, this makes thermoacoustic systems difficult to model *a priori* from existing models in the literature. On the positive side, this makes thermoacoustic models easy to train from experimental data because their parameters tend to be easily observable from experimental data. With well-chosen experiments, we can therefore (i) tune the parameters of candidate models and (ii) compare candidate models against each other and select the one with most evidence, given the experimental data [6,7]. Also on the positive side, this extreme sensitivity means that thermoacoustic systems can often be stabilized by making small changes, which is attractive in industrial settings [5]. The challenge is to model systems accurately and to determine stabilizing modifications as quickly and accurately as possible.

The most influential component of a thermoacoustic system is the mechanism linking acoustic velocity and pressure fluctuations at the base of the flame to subsequent h.r.r. fluctuations in the body of the flame. (In most settings, only the acoustic velocity fluctuations are influential.) This mechanism is difficult to model or simulate [8], so researchers often rely on experimental measurements of the h.r.r. as a function of the acoustic velocity. While in some premixed flames, unfiltered light emission has been shown to reasonably approximate the heat release rate [9–12], this is not always the case. In general, fluctuating natural emission from the flame is not a consistently reliable measurement of the fluctuating heat release rate (h.r.r.) [13]. Alternatives, such as PLIF to identify reaction zones [14] are more accurate but are technically difficult and, in large systems, impractical. This paper seeks to circumvent the above problems by combining experimental measurements with numerical simulations of a flame model. The flame model is physics-based and qualitatively-accurate. It gives, amongst other things, the flame front dynamics under harmonic forcing and the heat release rate as a function of time. The physics-based model has several physical parameters, such as the laminar flame speed and the Markstein length, whose expected values and uncertainties are inferred from the flame position, as found from experimental natural emission measurements. Therefore, we first use the natural emission data to estimate the flame front position over time. Next, we tune the model parameters by minimizing the discrepancy between the flame front dynamics predicted by the model and that obtained from the experiment. Finally, with the calibrated parameters, we use the model to compute the heat release rate (h.r.r.) and its uncertainty, as a function of the acoustic velocity perturbations.

In previous work, Juniper & Yoko [6] applied this process to a hot wire Rijke tube. From several candidate models in the literature, they selected those with the most evidence, given the data, and created a quantitatively accurate physics-based model of this thermoacoustic system. In subsequent work on the same system, they used Bayesian optimal experimental design [15] to evaluate the most informative experiments. Following this, they extended the method to infer the flame response of laminar and weakly turbulent flames from pressure measurements [16,17]. These study revealed that, in some cases, the uncertainty in the inferred flame responses could be large when the source of information was pressure measurements alone. In this paper we revisit the flames studied in [16] in order to assess an alternative source of information about the flame response: the dynamics of the flame front position. We use a high-speed camera to record the natural emission of the flames in both steady conditions and during harmonic acoustic forcing. We assimilate this data into a physics-based model of a conical Bunsen flame from a previous study [18]. This model contains

sufficient parameters to describe the forced experimental conical flame and is differentiable with respect to its parameters. We can therefore obtain the first derivatives of the model outputs with respect to model parameters and use these gradients to (i) find the most probable parameter values, given some data, and (ii) quantify our uncertainty in those parameters and the resulting uncertainty in the model predictions. This process creates a quantitatively-accurate physics-based model of the h.r.r. fluctuations as a function of the velocity perturbation. Unlike the output of many other data-driven approaches, the resulting model is interpretable, trustworthy and extrapolatable.

The paper is structured as follows: in section 2 we explain the experiments and the image processing; in section 3 we report the key feature of the model and how we generate the gradients of the model output with respect to its parameters; in section 4 we explain the Bayesian inference methodology and Laplace's approximation; in section 5.1 and 5.2 we show how we assimilate the experimental data to create the quantitatively-accurate physics-based model of the flame; in section 5.3 we show how we employ the model to predict the flame transfer function and its uncertainty; in section 5.4 we show how the model is capable of extrapolating outside the training dataset; in section 6 we discuss the conclusions and outline some future work.

## 2. Experiments

The experimental observations consist of high-speed footage of a conical flame under steady and acoustically-forced conditions. The footage is processed to isolate the position of the flame front.

### 2.1. Data collection

The experimental configuration is a laminar premixed conical flame inserted into a vertical duct, as illustrated in Fig. 1. The lower end of the duct is fixed to a plenum chamber, through which co-flow air is supplied. The upper end is open to the atmosphere. The duct is a 0.8 m long section of quartz tube with an internal diameter of 75 mm. Quartz is used to allow for optical access to the flame.

The burner is a 0.85 m long section of brass tubing with an internal diameter of 14 mm. The outlet of the burner is fitted with a nozzle that is chosen such that the system could become thermoacoustically unstable. At the injection plane, the nozzle diameter is 9.35 mm. The burner is fuelled by a mixture of methane and ethylene over a wide range of equivalence ratios and fuel mass flow rates, which allow us to study flames with a wide range of thermoacoustic responses. The reactants are metred using a set of mass flow controllers, from which they flow through a choke plate, decoupling the supply lines from the acoustic fluctuations in the rig.

We study 20 rich premixed flames, which we select to produce a wide range of flame responses without encountering blow-off. The flames are fuelled by a blend of methane and ethylene. The ethylene serves to increase the luminosity of the flame, improving the signal-to-noise ratio of the flame images.

We parameterize the flames based on (i) the convective time delay,  $\tau_c = L_f/\bar{u}$ , which is the time taken for a perturbation travelling at the bulk velocity in the burner tube,  $\bar{u}$ , to traverse the length of the flame,  $L_f$ , and (ii) the mean heat release rate of the inner cone,  $\bar{q}$ . We focus on the heat release rate of the inner cone because we have observed that, in our rig, the outer diffusion flame front does not respond to acoustic forcing, and therefore does not contribute to the thermoacoustic effect.

We group the 20 flames into five sets of four flames each, and select the composition of each flame such that the time delay is constant within a group and the mean heat release rate of the inner cone varies. The convective time delays range from 9.5 ms to 16 ms, and the mean heat release rates range from 375 W to 600 W. These flames produce a thermoacoustic behaviour ranging from strongly damped, to neutral, to strongly driven. The flame properties are summarized in Table A.1 in Appendix A, and illustrated in Fig. 2.

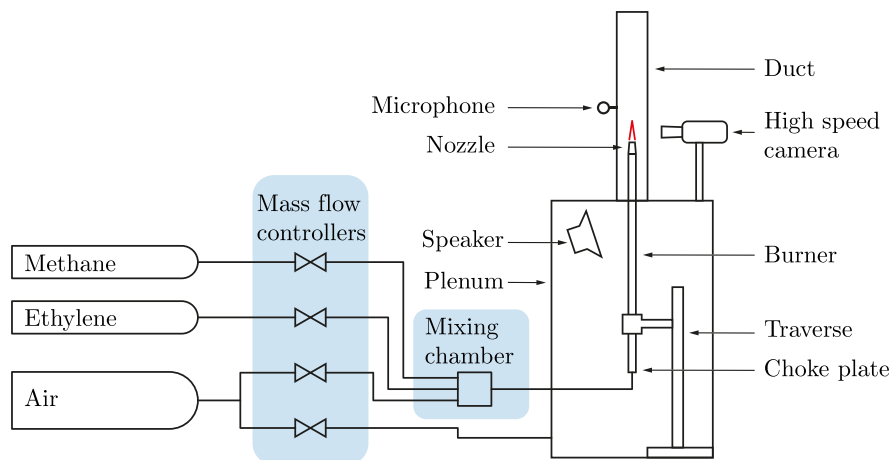


Fig. 1. Diagram of the experimental rig.

Depending on the flame properties, the system will either be linearly stable or self-excited. Our method requires footage of the flame under both stable and acoustically forced conditions. If the system is linearly stable, we begin by recording footage of the stable flame. We then use a loudspeaker to harmonically force the flames near the fundamental frequency of the system, and record high speed footage of the flame during forcing. If the system is self-excited, we begin by recording footage of the flame during limit-cycle oscillation. To record footage of the flame in the stable condition, we use active feedback control to stabilize the system. We record the emitted sound with a microphone and route this signal to a phase-shift amplifier. The phase-shift amplifier inverts the phase of the microphone signal and we play this through the loudspeaker, cancelling the acoustic oscillations and stabilizing the system.

## 2.2. Data processing

The raw footage is processed in order to (i) isolate the flame front position and (ii) produce a Euclidian distance field with the flame front as the zeroth level set. The images are first averaged to improve the signal-to-noise ratio. For the steady flames, this involves averaging over 200 frames. For the unsteady flames, it involves phase-averaging over 20 frames captured at 10 phase angles. The averaged images are straightened and centred so that the symmetry axis of the flame is vertical and in the centre of the frame. This is necessary to undo any misalignment in the camera setup. To do this, we exploit the symmetry of the flame and find the rigid transformation that maps the image to its mirror. From this, we can determine the transformation that moves the symmetry axis of the flame to be vertical and centred in the frame. Once the flame is centred, we perform an Abel deconvolution to undo the line-of-sight integration of the natural luminosity of the flame. We use an implementation of the ‘onion-peeling’ deconvolution [19]. This process produces a clearly defined flame front but accumulates noise towards the centre line of the image. We perform image segmentation on the Abel deconvolved image using Chan–Vese segmentation [20]. This separates the image signal from the background and allows us to interrogate the connectivity map of each cluster of pixels. We discard weakly connected clusters to denoise the image, leaving just the flame front as the dominant image structure. Finally, we perform a Euclidian distance transform on the binary image of the isolated flame front. This produces a matrix of the same dimensions as the original image, where each element contains the normal distance between the corresponding pixel and the flame front. We use this in our data assimilation process to quantify the discrepancy between the predicted and observed flame front positions.

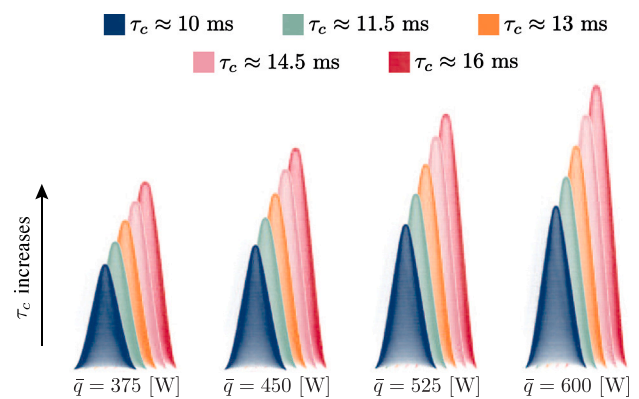


Fig. 2. Processed steady flame images from the 20 flames. Images are artificially coloured according to the approximate convective time delay. Flames with equal mean heat release rates are overlaid. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

## 3. Reduced-order model of a laminar premixed Bunsen flame

### 3.1. Model description

The flame model is derived from the G-equation and assumes that the flame front is infinitesimally thin and axisymmetric. The flame speed incorporates stretch effects through the Markstein length. The velocity field is modelled as the sum of a steady flow and a perturbation, which takes the form of a wave that originates at the burner rim and propagates longitudinally at a constant speed [21–23]. Although heat loss at the burner rim is not modelled directly, the flame base position in steady conditions and its trajectory under harmonic forcing are imposed as boundary conditions and assimilated using experimental data. All model details can be found in Ref. [18]. The model can be expressed compactly as follows:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), t; \mathbf{a}), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (1)$$

where  $\mathbf{x}(t) = (\mathbf{r}(t), z(t))$  is the state vector defining the flame front position in terms of radial coordinates,  $\mathbf{r}$ , and longitudinal coordinates,  $z$ , at a given time,  $t$ , and  $\mathbf{x}_0$  is the initial state. The vector  $\mathbf{a}$  contains the physics-based parameters, while  $\mathbf{f}$  is the physics-based nonlinear operator encapsulating the flame front dynamics. The main feature of the model is that it has been designed to be differentiable with respect to the state and the parameters. The heat release rate  $q(t)$  can also be expressed in compact form through:

$$q(t) = g(\mathbf{x}(t); \mathbf{a}), \quad (2)$$

where  $g$  is a nonlinear operator that links the heat release rate  $q(t)$  to the flame front position and shape defined by  $\mathbf{x}(t)$ . The model parameters  $\mathbf{a}$  are: (i) the flame aspect ratio  $\beta = \beta(\bar{u}, S_L)$ , which depends on the bulk velocity in the burner tube  $\bar{u}$  and the unstretched laminar flame speed  $S_L$ ; (ii) the nondimensional Markstein length  $\mathcal{M}$ , which depends on the thermal expansion parameter, the effective Lewis number of the mixture, the Zel'dovich number and the thermal conductivity of the mixture [24]; (iii) the shape parameter  $\mu$ , which linearly combines a uniform and a parabolic mean velocity profile [25]; (iv) the wavelength of the harmonic perturbation velocity field  $K$ ; (v) the amplitude of the acoustic forcing  $\varepsilon_V$ ; (vi) the amplitude of the flame base oscillations  $\lambda$ ; (vii) the initial phase of the flame base oscillations  $\phi_0$ . These seven parameters are sufficient to describe the flame front dynamics qualitatively. All lengths, including the Markstein length, are normalized by the nozzle radius  $R = 4.675$  mm.

### 3.2. Obtaining periodic solutions and their sensitivities using adjoint methods

All flames are perturbed by a harmonic acoustic forcing with period  $T$ . We consider the following vector-valued function:

$$\mathbf{F}(\mathbf{x}_0; \mathbf{a}) = \mathbf{x}(T) - \mathbf{x}_0, \quad \text{where } \mathbf{x}(T) = \int_0^T \mathbf{f}(\mathbf{x}(t), t; \mathbf{a}) dt + \mathbf{x}_0. \quad (3)$$

For a given parameter set  $\mathbf{a}$ , a solution  $\mathbf{x}(t)$  of Eq. (1) with an initial state  $\mathbf{x}_0$  is periodic with period  $T$  if:

$$\mathbf{F}(\mathbf{x}_0; \mathbf{a}) = \mathbf{0}. \quad (4)$$

We solve Eq. (4) using the trust-region dogleg algorithm and by providing the partial derivative of  $\mathbf{F}$  with respect to the initial state  $\mathbf{x}_0$

$$\frac{\partial \mathbf{F}}{\partial \mathbf{x}_0} = \frac{\partial \mathbf{x}(T)}{\partial \mathbf{x}_0} - \mathbf{I}, \quad (5)$$

where  $\mathbf{I}$  is the identity matrix, and  $\partial \mathbf{x}(T)/\partial \mathbf{x}_0$  is the Jacobian matrix containing the partial derivatives of the state at time  $T$  with respect to changes in the initial state  $\mathbf{x}_0$ . We compute these derivatives using adjoint methods, as discussed in Refs. [26–28]. Further details are provided in Appendix B. We can also obtain the gradient of the periodic solution  $\mathbf{x}(t)$  with respect to the parameters  $\mathbf{a}$ , by differentiating Eq. (3) with respect to  $\mathbf{a}$ :

$$\frac{d\mathbf{F}}{d\mathbf{a}} = \frac{\partial \mathbf{x}(T)}{\partial \mathbf{a}} + \frac{\partial \mathbf{x}(T)}{\partial \mathbf{x}_0} \frac{d\mathbf{x}_0}{d\mathbf{a}} - \frac{d\mathbf{x}_0}{d\mathbf{a}}. \quad (6)$$

If Eq. (4) holds true for any variation in the parameters (i.e., the solution remains periodic when the parameters change), we can set Eq. (6) to zero. Consequently, we have:

$$\frac{d\mathbf{F}}{d\mathbf{a}} = \mathbf{0} \implies \left( \frac{\partial \mathbf{x}(T)}{\partial \mathbf{x}_0} - \mathbf{I} \right) \frac{d\mathbf{x}_0}{d\mathbf{a}} = - \frac{\partial \mathbf{x}(T)}{\partial \mathbf{a}}, \quad (7)$$

which can be solved for the sensitivity matrix  $d\mathbf{x}_0/d\mathbf{a}$ . Next, we can obtain the sensitivity of the periodic solution at any arbitrary time by differentiating the state  $\mathbf{x}(t)$  with respect to the model parameters. The state depends directly on the parameters and the initial condition. Thus, we obtain:

$$\frac{d\mathbf{x}(t)}{d\mathbf{a}} = \frac{\partial \mathbf{x}(t)}{\partial \mathbf{a}} + \frac{\partial \mathbf{x}(t)}{\partial \mathbf{x}_0} \frac{d\mathbf{x}_0}{d\mathbf{a}}. \quad (8)$$

In Appendix B, we describe how to compute the matrices  $\partial \mathbf{x}(t)/\partial \mathbf{a}$  and  $\partial \mathbf{x}(t)/\partial \mathbf{x}_0$ . Finally, we compute the gradients of the heat release rate (h.r.r.) with respect to the parameters by differentiating Eq. (2):

$$\frac{dq(t)}{d\mathbf{a}} = \frac{\partial g(\mathbf{x}(t); \mathbf{a})}{\partial \mathbf{a}} + \frac{\partial g(\mathbf{x}(t); \mathbf{a})}{\partial \mathbf{x}(t)} \frac{d\mathbf{x}(t)}{d\mathbf{a}}. \quad (9)$$

This quantity is essential for the rapid calculation of uncertainties in the heat release rate and, consequently, in the flame transfer function.

## 4. Bayesian data assimilation

Bayesian data assimilation provides a set of tools that allow the researcher to (i) infer the most probable parameters of a physics-based model, given some data, (ii) quantify the uncertainty in the model parameters, and estimate the systematic uncertainty in the model and data, (iii) rank several candidate models to select the best model, given the data, and (iv) identify optimal experiment designs and sensor placements. In this paper, we use points (i) and (ii), so these are described in detail. A demonstration of model ranking can be found in [16], and a demonstration of optimal experiment design can be found in [15].

In this section, we introduce the following notation to outline the Bayesian inference methodology: the data  $\mathbf{D}$  contains several experimental observations of the flame position,  $\mathbf{x}^e$ , in steady and/or unsteady configurations; the model  $\mathcal{H}$  encodes a physics-based reduced-order model, such that for a given set of parameters,  $\mathbf{a}$ , the model  $\mathcal{H}(\mathbf{a})$  returns the predicted flame front position,  $\mathbf{x}$ .

### 4.1. Parameter inference

When performing parameter inference, we assume that the model is structurally correct and we infer its most probable parameters,  $\mathbf{a}_{\text{MP}}$ . The model is rarely free of structural error, however, and we will revisit this assumption later. We encode our level of uncertainty in the parameter values through a probability distribution, which we denote  $P(\bullet)$ . Using any prior knowledge we have about the unknown parameters (which may be none at all), we propose a prior probability distribution over the parameter values,  $P(\mathbf{a}|\mathcal{H})$ . We then assimilate the data,  $\mathbf{D}$ , by performing a Bayesian update on the parameter values:

$$P(\mathbf{a}|\mathbf{D}, \mathcal{H}) = \frac{P(\mathbf{D}|\mathbf{a}, \mathcal{H})P(\mathbf{a}|\mathcal{H})}{P(\mathbf{D}|\mathcal{H})} \quad (10)$$

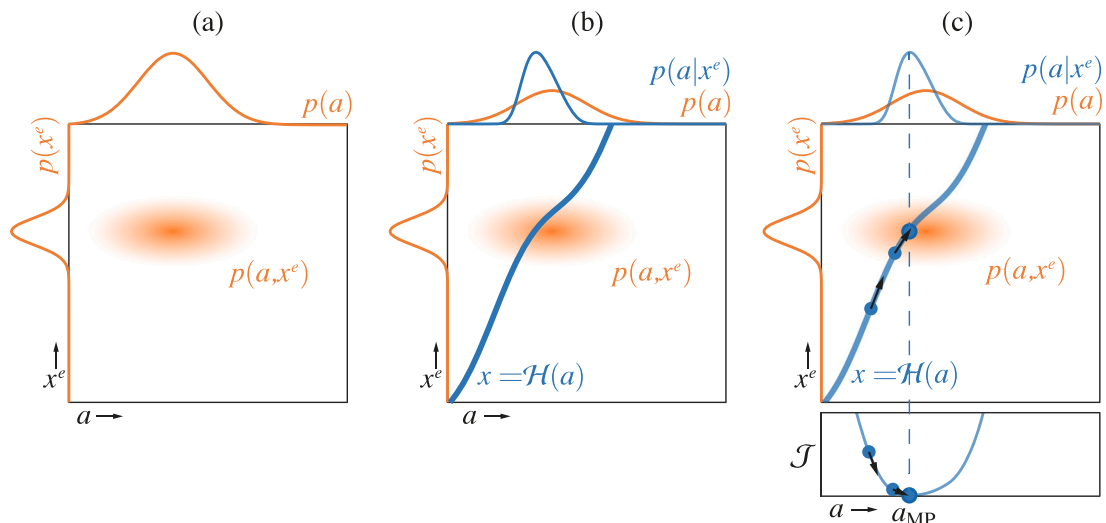
The quantity on the left-hand side of Eq. (10) is the posterior probability of the parameters, given the data. It is generally computationally intractable to calculate the full posterior, because this requires integration over a large parameter space. The integral typically cannot be evaluated analytically, and requires thousands of model evaluations to compute numerically. At the parameter inference stage, however, we are only interested in finding the most probable parameters, which are those that maximize the posterior. We therefore use an optimization algorithm to find the peak of the posterior without evaluating the full distribution.

The right-hand side of Eq. (10) does not depend on the parameters, so we can find the most probable parameters by maximizing the product  $P(\mathbf{D}|\mathbf{a}, \mathcal{H})P(\mathbf{a}|\mathcal{H})$ . This process is made computationally efficient by (i) assuming that the experimental uncertainty is Gaussian distributed, such that  $P(\mathbf{D}|\mathbf{a}, \mathcal{H})$  is a Gaussian distribution over the data for a given set of parameters, and (ii) choosing the prior parameter distribution,  $P(\mathbf{a}|\mathcal{H})$ , to be Gaussian. Assumption (i) is reasonable for well-designed experiments in which the uncertainty is dominated by random error, which is typically Gaussian distributed. For assumption (ii) we note that the choice of prior is often the prerogative of the researcher, and we are free to exploit the mathematical convenience offered by the Gaussian distribution. Under the assumption of Gaussian distributions, we can transform the optimization problem into a quadratic minimization problem by defining the cost function,  $\mathcal{J}$ , to be the negative log of the product  $P(\mathbf{D}|\mathbf{a}, \mathcal{H})P(\mathbf{a}|\mathcal{H})$ :

$$\begin{aligned} \mathcal{J}(\mathbf{a}) &= -\log \{ P(\mathbf{D}|\mathbf{a}, \mathcal{H})P(\mathbf{a}|\mathcal{H}) \} \\ &= \sum_{i=1}^{N_e} \frac{1}{2} (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e)^T \mathbf{C}_{ee}^{-1} (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e) \\ &\quad + \frac{1}{2} (\mathbf{a} - \mathbf{a}_p)^T \mathbf{C}_{aa}^{-1} (\mathbf{a} - \mathbf{a}_p) + \text{const.} \end{aligned} \quad (11)$$

where: the summation is over the  $N_e$  frames of the footage;  $\mathbf{x}(\mathbf{a}, t_i)$  and  $\mathbf{x}_i^e$  are column vectors of the predicted and measured flame front





**Fig. 3.** Illustration of parameter inference on a simple univariate system. (a) the marginal probability distributions of the prior and data,  $p(a)$  and  $p(x^e)$ , as well as their joint distribution,  $p(a, x^e)$  are plotted on axes of parameter value,  $a$ , vs observation outcome,  $x^e$ . (b) the model,  $\mathcal{H}$ , imposes a functional relationship between the parameters,  $a$ , and the predictions,  $x$ . Marginalizing along the model predictions yields the true posterior,  $p(a|x^e)$ . This is computationally intractable for even moderately large parameter spaces. (c) instead of evaluating the full posterior, we use gradient-based optimization to find its peak. This yields the most probable parameters,  $a_{\text{MP}}$ .

positions at the time of the  $i$ th snapshot;  $\mathbf{C}_{ee}$  is the covariance matrix describing the experimental uncertainty;  $\mathbf{a}$  and  $\mathbf{a}_p$  are column vectors of the current and prior parameter values respectively;  $\mathbf{C}_{aa}$  is the covariance matrix describing the uncertainty in the prior, and  $\text{const.}$  is a constant that arises from the Gaussian pre-exponential factors, which does not impact the maximum posterior parameter estimation  $\mathbf{a}_{\text{MP}}$ . In this study, we evaluate the discrepancy between the model predictions and the experimental observations by projecting the predicted flame front position onto the Euclidian distance field produced from the corresponding frame of the high speed footage. This produces a column vector in which each entry is the normal distance between the predicted and observed flame front position at each point in the model. We minimize Eq. (11) over the parameters using the BFGS gradient-based optimization algorithm, where the gradient information is provided using adjoint methods.

The parameter inference process is illustrated in Fig. 3 for a simple system with a single parameter,  $a$ , and a single observable variable,  $x^e$ . In (a) we show the marginal probability distributions of the prior,  $p(a)$ , and the data,  $p(x)$ . The prior and data are independent, so we construct the joint distribution,  $p(a, x^e)$  by multiplying the two marginals. In (b), we overlay the model predictions,  $x$ , for various values of  $a$ . Marginalizing along the model predictions yields the true posterior,  $p(a|x^e)$ . This is possible for a cheap model with a single parameter, but exact marginalization quickly becomes intractable as the number of parameters increases. In (c) we plot the cost function,  $\mathcal{J}$ , which is the negative log of the unnormalized posterior. We show the three steps of gradient-based optimization that are required to find the local minimum, which corresponds to the most probable parameters,  $a_{\text{MP}}$ .

#### 4.2. Uncertainty quantification

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

##### 4.2.1. Parametric uncertainty

Once we have found the most probable parameters,  $\mathbf{a}_{\text{MP}}$ , we estimate the uncertainty in these parameter values using Laplace's method [6,29,30]. This method approximates the posterior as a Gaussian distribution with a mean of  $\mathbf{a}_{\text{MP}}$ , and a covariance given by the Hessian of

the cost function:

$$\begin{aligned} \mathbf{C}_{aa}^{\text{MP-1}} &\approx \frac{\partial^2 \mathcal{J}}{\partial a_i \partial a_j} \\ &\approx \mathbf{C}_{aa}^{-1} + \mathbf{J}^T \mathbf{C}_{ee}^{-1} \mathbf{J} \end{aligned} \quad (12)$$

where  $\mathbf{J}$  is the Jacobian matrix of the model predictions with respect to the parameters  $\mathbf{a}$ , and the higher-order terms have been neglected. It was claimed in Section 1 that the sensitivity of thermoacoustics can be beneficial because it makes the uncertain parameters easy to infer from experimental data. This is clear from Eq. (12), which shows that the precision (inverse variance) of the posterior parameter estimate increases with parameter sensitivity.

The accuracy of the Laplace approximation depends on the functional dependence between the model predictions and the parameters. This is shown graphically in Fig. 4 for three univariate systems. In (a), the model is linear in the parameters. Marginalizing a Gaussian joint distribution along any intersecting line produces a Gaussian posterior, so Laplace's method is exact. In (b), the model is weakly nonlinear in the parameters. The true posterior is skewed, but the Gaussian approximation is still reasonable. This panel also shows a geometric interpretation of Laplace's method: the approximate posterior is given by linearizing the model around  $\mathbf{a}_{\text{MP}}$ , and marginalizing the joint distribution along the linearized model. In (c), the model is strongly nonlinear in the parameters, so the true posterior is multi-modal and the main peak is highly skewed. In this case, the gradient-based optimization algorithm will only find a single local minimum, which will depend on the choice of initial condition for the optimization. Furthermore, we see that the posterior estimated using Laplace's method is a poor approximation of the true posterior. This problem can be avoided by reducing the extent of the nonlinearity captured by the joint distribution by (i) shrinking the joint distribution by providing more precise prior information or more precise experimental data, or (ii) re-parameterizing the model to reduce the strength of the nonlinearity [30, Chapter 27].

In many cases, the use of this approximate inference framework will be justifiable, given the substantial reduction in computational cost compared to sampling methods, which are the only viable alternative for constructing the posterior. In previous work, we compared the computational cost of our framework to two sampling approaches [16]. The comparison was done on a computationally cheap thermoacoustic network model. Applying our framework to this model, we can compute the posterior probability of five unknown parameters in under 5 s on a single core on a laptop. The same inference problem takes 35

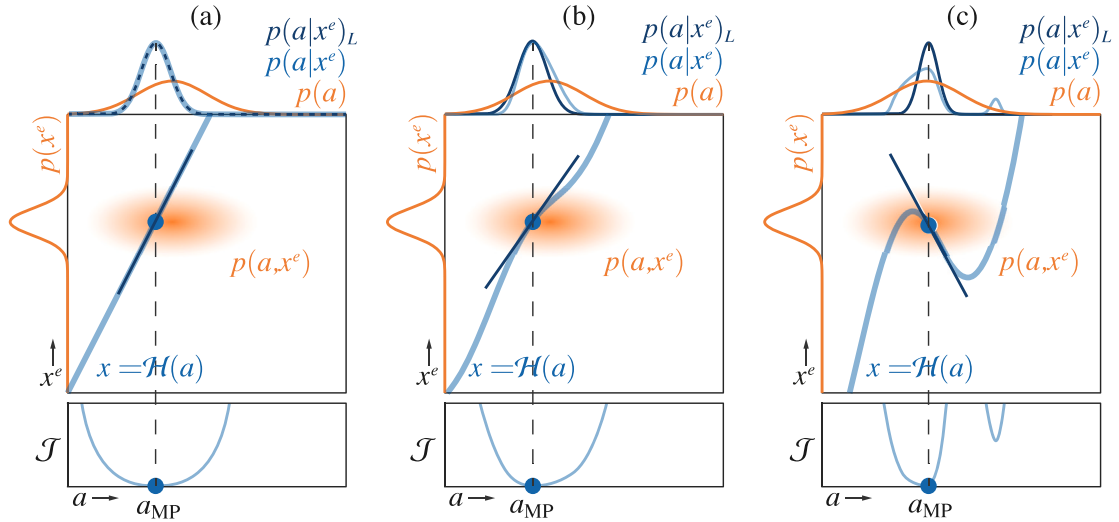


Fig. 4. Illustration of uncertainty quantification for three univariate systems, comparing the true posterior,  $p(a|z)$  to the approximate posterior from Laplace's method,  $p(a|z)_L$ . (a) the model is linear in the parameters, so the true posterior is Gaussian and Laplace's method is exact. (b) the model is weakly nonlinear in the parameters, the true posterior is slightly skewed, but Laplace's method yields a reasonable approximation. (c) the model is strongly nonlinear in the parameters, the posterior is multi-modal and Laplace's method underestimates the uncertainty.

CPU hours running on a workstation when solved with Markov Chain Monte Carlo, and 22 CPU hours when solved with Hamiltonian Monte Carlo.

#### 4.2.2. Uncertainty propagation

To quantify the uncertainty in the model predictions due to the parametric uncertainty, we linearize the model around  $\mathbf{a}_{MP}$  and propagate the parameter uncertainties through the linearized model [31]. The uncertainty in the model prediction is given by:

$$\mathbf{C}_{xx} = \mathbf{J}^T \mathbf{C}_{aa} \mathbf{J} \quad (13)$$

where  $\mathbf{C}_{xx}$  represents the covariance in the model predictions, and  $\mathbf{J}$  denotes the Jacobian matrix of the model predictions with respect to the parameters  $\mathbf{a}$ . The square of the marginal uncertainty in each model prediction,  $(\sigma_{x_i})^2$ , is given by the diagonal entries of  $\mathbf{C}_{xx}$ .

#### 4.2.3. Systematic uncertainty

In most cases, experimental data will contain some systematic uncertainty, and models will contain some structural uncertainty. These uncertainty sources cannot be quantified *a-priori*, and are often referred to as “unknown unknowns”. We can, however, construct a total covariance matrix,  $\mathbf{C}_{tt}$ , which encodes the total uncertainty due to (i) the known experimental uncertainty, (ii) the unknown systematic experimental uncertainty, and (iii) the unknown structural model uncertainty. We can then estimate this total covariance from the posterior discrepancy between the model and the data. This must be done simultaneously with parameter inference, because the posterior parameter distribution depends on the total uncertainty in the model and data. We therefore replace  $\mathbf{C}_{ee}$  with  $\mathbf{C}_{tt}$  in Eq. (11), and estimate the total uncertainty by simultaneously minimizing  $\mathcal{J}$  with respect to  $\mathbf{a}$  and  $\mathbf{C}_{tt}^{-1}$ .

We begin by calculating the derivative of  $\mathcal{J}$  with respect to  $\mathbf{C}_{tt}^{-1}$ , assuming that the observed variables are uncorrelated, and keeping in mind that the normalizing constant,  $K$ , depends on  $\mathbf{C}_{tt}$ :

$$\mathcal{J} = \sum_{i=1}^{N_e} \frac{1}{2} (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e)^T \mathbf{C}_{tt}^{-1} (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e) + N_e \log \left( \sqrt{(2\pi)^k |\mathbf{C}_{tt}|} \right) + \frac{1}{2} (\mathbf{a} - \mathbf{a}_p)^T \mathbf{C}_{aa}^{-1} (\mathbf{a} - \mathbf{a}_p) + \log \left( \sqrt{(2\pi)^k |\mathbf{C}_{aa}|} \right) \quad (14)$$

$$\frac{\partial \mathcal{J}}{\partial \mathbf{C}_{tt}^{-1}} = \sum_{i=1}^{N_e} \frac{1}{2} (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e) (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e)^T \circ \mathbf{I} - \frac{1}{2} N_e \mathbf{C}_{tt} \quad (15)$$

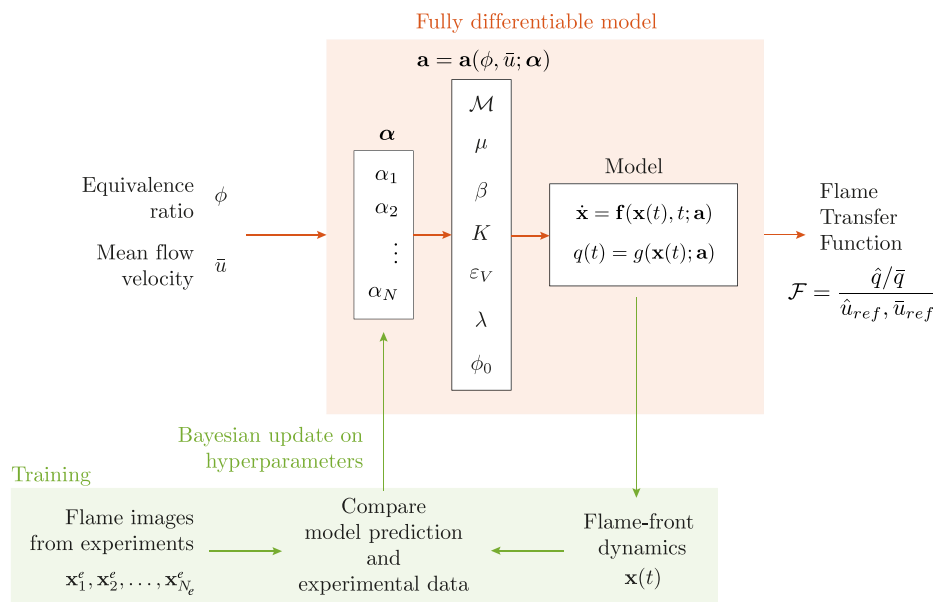
where  $\mathbf{I}$  is the identity matrix, and  $\circ$  denotes the Hadamard product. For a given set of parameters, the most probable  $\mathbf{C}_{tt}$  sets equation (15) to zero. This gives the estimate:

$$\mathbf{C}_{tt} = \frac{1}{N_e} \sum_{i=1}^{N_e} (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e) (\mathbf{x}(\mathbf{a}, t_i) - \mathbf{x}_i^e)^T \circ \mathbf{I} \quad (16)$$

which is the expected result that the total variance in the model and data is the mean square of the discrepancy between the model predictions and the data. Although we cannot directly identify the source of the unknown uncertainty because the experimental and model uncertainties cannot be disentangled, the inferred total uncertainty can assist the researcher with identifying potential error sources. For example, if the unknown error in a single sensor is unexpectedly large, this could indicate a faulty sensor or bad installation. If the unknown error at a certain experimental operating condition is large, this could prompt the researcher to repeat that experiment. If the unknown error grows with one of the input variables, the researcher might investigate the model to see if any important physical phenomena has been neglected.

## 5. Creating the quantitatively-accurate physics-based model

In this study, we perform Bayesian data assimilation in two steps. In the first step, we assimilate the flame front position data of each flame individually. Therefore, for each experimental configuration (i.e. for each combination of mean flow velocity  $\bar{u}$  and equivalence ratio  $\phi$ ), we determine the most probable set of model parameters  $\mathbf{a} = (\mathcal{M}, \mu, \beta, K, \varepsilon_V, \lambda, \phi_0)$ , which have been introduced in Section 3. This process does not create a general model that can predict the behaviour of an arbitrary flame, but is useful for assessing the influence of the model parameters  $\mathbf{a}$  on the model predictions. After completion of this step, we need to introduce a set of hyperparameters  $\alpha$ , such that  $\mathbf{a} = \mathbf{a}(\phi, \bar{u}; \alpha)$ . We then simultaneously assimilate the flame front position data of all flames. This allows us to determine the most probable set of hyperparameters  $\alpha$  so that, for each combination of mean flow velocity  $\bar{u}$  and equivalence ratio  $\phi$ , the model  $\mathbf{a} = \mathbf{a}(\phi, \bar{u}; \alpha)$  provides the best prediction of the physics-based model parameters  $\mathbf{a}$  and, consequently, of the flame front dynamics. In Fig. 5, we show a diagram describing the methodology followed in this study. The end goal is to create a general model that provides the flame transfer function, along with quantified uncertainty bounds, for any given pair of equivalence ratio and bulk velocity in the burner tube.



**Fig. 5.** Illustration showing the methodology for constructing a quantitatively-accurate reduced-order model. The general model depends on the hyperparameters  $\alpha$  and comprises two components: the physics-based reduced order model, which depends on the physical parameters  $\mathbf{a}$ , and a metamodel  $\mathbf{a} = \mathbf{a}(\phi, \bar{u}; \alpha)$  that establishes the relationship between these physical parameters  $\mathbf{a}$  and the hyperparameters  $\alpha$ . The model is fully differentiable and provides the flame front dynamics and the flame transfer function. By comparing the model's predictions of the flame front with experimental images, we refine our hyperparameters  $\alpha$  through Bayesian updates. The resulting model predicts the flame transfer function, with confidence bounds, for any combination of equivalence ratio  $\phi$  and the mean flow velocity  $\bar{u}$ . This model is both interpretable and capable of extrapolation.

### 5.1. Step 1: assimilation of each flame individually

First, we assimilate the steady-state snapshots in order to infer  $\mathcal{M}$ ,  $\beta$  and  $\mu$ . Second, we use these inferred values as priors for the second step, in which we assimilate the unsteady flame images and infer all parameters simultaneously:  $\mathcal{M}$ ,  $\beta$ ,  $\mu$ ,  $K$ ,  $\varepsilon_V$ ,  $\lambda$ ,  $\phi_0$ . Here we show the process for flame 17 in Table A.1, which has an equivalence ratio  $\phi = 1.34$ , a flow rate of  $\bar{u} = 2.86$  m/s and an equal volume fraction of methane ( $\text{CH}_4$ ) and ethylene ( $\text{C}_2\text{H}_4$ ).

#### 5.1.1. Steady-state data assimilation

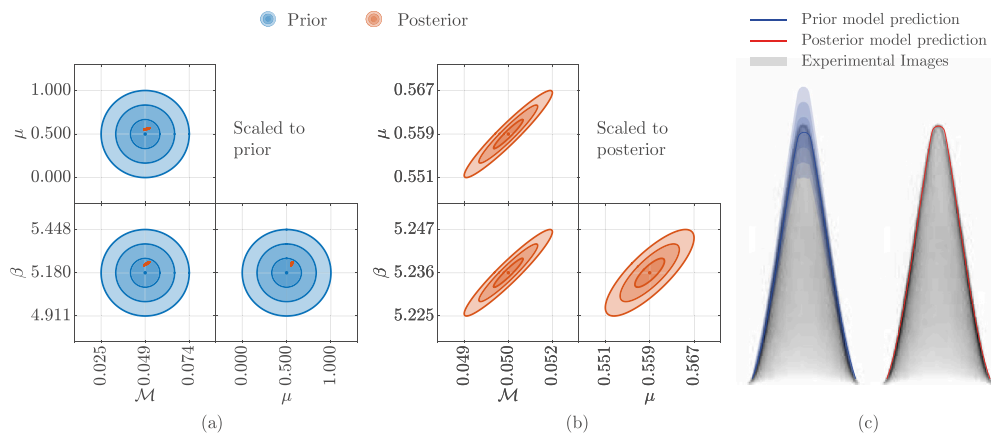
The steady flame front position depends on  $\beta$ ,  $\mathcal{M}$  and  $\mu$ . We can include prior knowledge about the parameter  $\beta = \beta(\bar{u}, S_L)$  because it depends on the unstretched flame speed  $S_L$ , which can be estimated using the open source chemical kinetics code Cantera [32], and the mean flow velocity  $\bar{u}$ , which can be measured during the experiments. Furthermore, we can estimate the value of the Markstein length by using Matalon's formulation [24]. Following this process for flame 17, we assign a prior estimate to the unstretched laminar flame speed of  $S_L = 0.4036 \text{ ms}^{-1}$  and a dimensional Markstein length of  $\tilde{\mathcal{M}} = 0.2579 \text{ mm}$ , corresponding to  $\mathcal{M} = \tilde{\mathcal{M}}/R = 0.0552$ . To construct the parameter covariance matrix, we must assign confidence bounds to these values, which is naturally a subjective process. In this case, we assign an uncertainty of  $6\sigma_{S_L} = 10\%S_L$  and  $6\sigma_{\mathcal{M}} = 100\%\mathcal{M}$  to indicate that we are relatively more confident in our ability to predict laminar flame speed than Markstein length. The only prior knowledge available for the velocity shape parameter  $\mu$  is that it can vary between 0 and 1, so we set a prior with large uncertainty for this parameter. From the steady-state flame snapshot, we can infer the *maximum a posteriori* probability estimate of the aspect ratio,  $\beta$ , the nondimensional Markstein length,  $\mathcal{M}$ , and the velocity shape parameter,  $\mu$ . Fig. 6(a–b) shows the prior and posterior probability distributions of the model parameters. The contours show 1, 2 and 3 standard deviations from the prior  $\mathbf{a}_p$  (blue) and posterior  $\mathbf{a}_{MP}$  (red). We see that the experimental observations have reduced the uncertainties in the model parameters. Fig. 6(c) shows the model prediction  $\pm 3$  standard deviations overlaid on the steady flame image. The largest uncertainties are found at the flame tip, which is the region that is most sensitive to the parameters. The model prediction uncertainty is found by solving Eq. (13).

#### 5.1.2. Unsteady-state (acoustically forced)

We repeat the process with the images of the acoustically forced flame. In this case, we set the values of the three steady parameters using the information gained from the previous step, and we set large uncertainties in the prior values of  $K$ ,  $\varepsilon_V$ ,  $\lambda$  and  $\phi_0$ . Fig. 7(a–b) shows the prior and posterior distribution of the parameters before and after assimilating the unsteady flame images. The results show that the uncertainty in the values of  $K$ ,  $\varepsilon_V$ ,  $\lambda$  and  $\phi_0$  has been reduced significantly through the assimilation of this data, and the values of  $\mathcal{M}$  and  $\mu$  have shifted slightly from the values found during the assimilation of the steady data, without exceeding the 3-standard deviation bounds. Fig. 7(c) shows the model prediction  $\pm 3$  standard deviations, plotted against the unsteady flame images before (blue) and after (red) assimilating the experimental video footage. As in the steady-state case, the largest uncertainties are found at the flame tip. We see that the model matches the data well at each frame. With this, we have produced a digital twin of the flame, which we can use to estimate additional quantities of interest such as the unstretched laminar flame speed  $S_L$ , the Markstein length  $\mathcal{M}$ , the convective speed of velocity perturbations  $K$  and the heat release rate, which were not directly measured in the experiments. Furthermore, because the model is probabilistic, we can also quantify our uncertainty in these estimates.

For the example of flame 17, the unstretched laminar flame speed and the Markstein length are estimated to be  $S_L = 0.4191$  m/s and  $\mathcal{M} = 0.2430$  mm compared to the value of  $S_L^c = 0.4036$  m/s and  $\mathcal{M}^c = 0.2579$  mm obtained from Cantera and the value of the convective speed is estimated to be  $K = 0.8874$  compared to  $K = 1$  in Schuller et al. [21],  $K = 0.9$  in Kashinath et al. [33] and  $K = 0.83$  in Orchini & Juniper [34]. In this case, we were able to provide reasonable prior information for  $\mathcal{M}$  and  $S_L$ . Similar values are, however, inferred even if the prior is inaccurate and uncertain. This methodology could therefore be used to cheaply and easily estimate combustion properties of fuels for which data are not available.

The proposed model contains few parameters, allowing it to be trained on a relatively small amount of data. Furthermore, the data required to train the model is easy to collect from a simple experiment because only snapshots of the natural emission of the flame are required.



**Fig. 6.** (a-b) Probability distributions of parameters in the steady case: prior (blue) and posterior (red). Contours denote 1, 2, and 3 standard deviations from the prior  $\mathbf{a}_p$  and posterior  $\mathbf{a}_{MP}$  estimations. Parameters include the flame aspect ratio  $\beta$ , which is a function of the unstretched laminar flame speed  $S_L$  and the bulk velocity in the burner tube  $\bar{u}$ , the shape parameter for the velocity field  $\mu$ , and the nondimensional Markstein length  $\mathcal{M}$ . Both panels represent the same quantities but panel (a) is scaled to the prior distributions, while panel (b) is scaled to the posterior distributions. (c) Comparison between model predictions before (red line) and after (blue line) assimilating the experimental steady-state flame image. Confidence intervals, spanning 3 standard deviations, are illustrated by the red and blue shadings overlaying the steady flame image. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Once the most probable parameter values are inferred and the flame front position has been predicted, we can estimate the heat release rate using Eq. (2). The resulting heat release rate estimation is shown for flame 8 in Fig. 8. Panel (a) shows the model prediction (red lines) with a confidence of 3 standard deviations (red shading) overlaid on the experimental observations (greyscale). We also plot the predictions between two observations (blue lines and shading). Panel (b) shows the h.r.r. perturbation  $q'$  normalized by its mean value  $\bar{q}$ , computed using the model. The black dots correspond to the time frames observed in the experiments. Fig. 8 highlights the advantages of the Bayesian inference approach: using experimental flame images we know the flame front position at ten different time points during a period of oscillation; we use this data to infer the seven parameters of a physics-based reduced order model, which gives the flame front position and the heat release rate at any point in time with their uncertainties.

## 5.2. Step 2: assimilation of all flames into a general model

At this stage, we introduce a metamodel  $\mathbf{a} = \mathbf{a}(\phi, \bar{u}; \alpha)$ , designed to determine the model parameters  $\mathbf{a}$  based on the equivalence ratio  $\phi$  and bulk flow rate in the burner tube  $\bar{u}$ , given a set of hyperparameters  $\alpha$ . Fig. 9 shows the *maximum a posteriori* probability distribution of the model parameters  $\mathbf{a}_{MP}$  assimilated in the first step, denoted by blue markers with error bars representing 3 standard deviations. Panel (a) shows the assimilated unstretched laminar flame speed  $S_L$  against the equivalence ratio  $\phi$ . Panel (b) shows the assimilated Markstein length  $\mathcal{M}$  against the equivalence ratio  $\phi$ . Panel (c) shows the assimilated shape parameter  $\mu$  against the bulk velocity  $\bar{u}$  in the burner tube. The parameter  $\mu$  decreases linearly with increasing  $\bar{u}$ . Panel (d) shows the assimilated convective speed ratio  $K$  against the acoustic-forcing Strouhal number  $St$ . The plot shows that  $K$  falls within the range of 0.8 to 1, with no discernible pattern with respect to  $St$ . The remaining parameters  $\varepsilon_V$ ,  $\lambda$ , and  $\phi_0$  (not shown) depend on experimental conditions and vary according to external acoustic forcing, both of which are independent of the flame properties.

Based on these observations, the model  $\mathbf{a} = \mathbf{a}(\phi, \bar{u}; \alpha)$  is formulated as follows. (i) The laminar flame speed  $S_L(\phi)$  is expressed as  $S_L(\phi) = S_L^c(\phi) + \alpha_1\phi + \alpha_2$ , indicating that it is equal to the laminar flame speed computed using Cantera,  $S_L^c(\phi)$ , plus a linear correction that depends on the equivalence ratio. In this study, we have data over a narrow range of equivalence ratios, and a simple linear correction for the laminar flame speed is sufficient. However, if we were to study a wider range of equivalence ratios, we would need a more elaborate correction. (ii) The Markstein length  $\mathcal{M}(\phi)$  is given by  $\mathcal{M}(\phi) =$

$\alpha_3\mathcal{M}^c(\phi)$ , indicating that it is equal to the Markstein length computed using Cantera,  $\mathcal{M}^c(\phi)$ , multiplied by the parameter  $\alpha_3$ . (iii) The shape parameter  $\mu = \alpha_4\bar{u} + \alpha_5$  is assumed to vary linearly with the bulk flow velocity and therefore with the Reynolds number of the bulk flow in the burner tube. (iv) The value  $K = \alpha_6$  is assumed to be constant across all flame cases. We keep  $\varepsilon_V$ ,  $\lambda$  and  $\phi_0$  equal to the values obtained in the previous assimilation step because the amplitudes of the acoustic forcing and of the flame base oscillations depend on the specific experimental conditions and are not general properties of the flame.

We then assimilate all the data simultaneously to infer the most probable hyperparameter set, denoted as  $\alpha_{MP}$ . The results of this step are shown in red in Fig. 9. The red lines display the estimates of the physical parameters  $\mathbf{a}$  as functions of the equivalence ratio  $\phi$ , the mean flow velocity  $\bar{u}$ , and the Strouhal number  $St$ , after the assimilation of the hyperparameters  $\alpha$ . The three red contours show 1, 2, and 3 standard deviations from the posterior values.

We now compare the model predictions to experimental data for all 20 flames, while calculating the parameters of the physics-based model using our metamodel. Fig. 10 displays the predicted flame front dynamics for all 20 flames in the library at a single timestep. The experimental observations are shown as greyscale images, and the model predictions are overlaid as red lines with confidence intervals of 1, 2, and 3 standard deviations (depicted as red shading). The numbers at the base of the flames indicate the corresponding flame numbers referenced in Table A.1. In the supplementary material, we provide the animated version of this figure, showing the predictions for all timesteps of the experimental video footage.

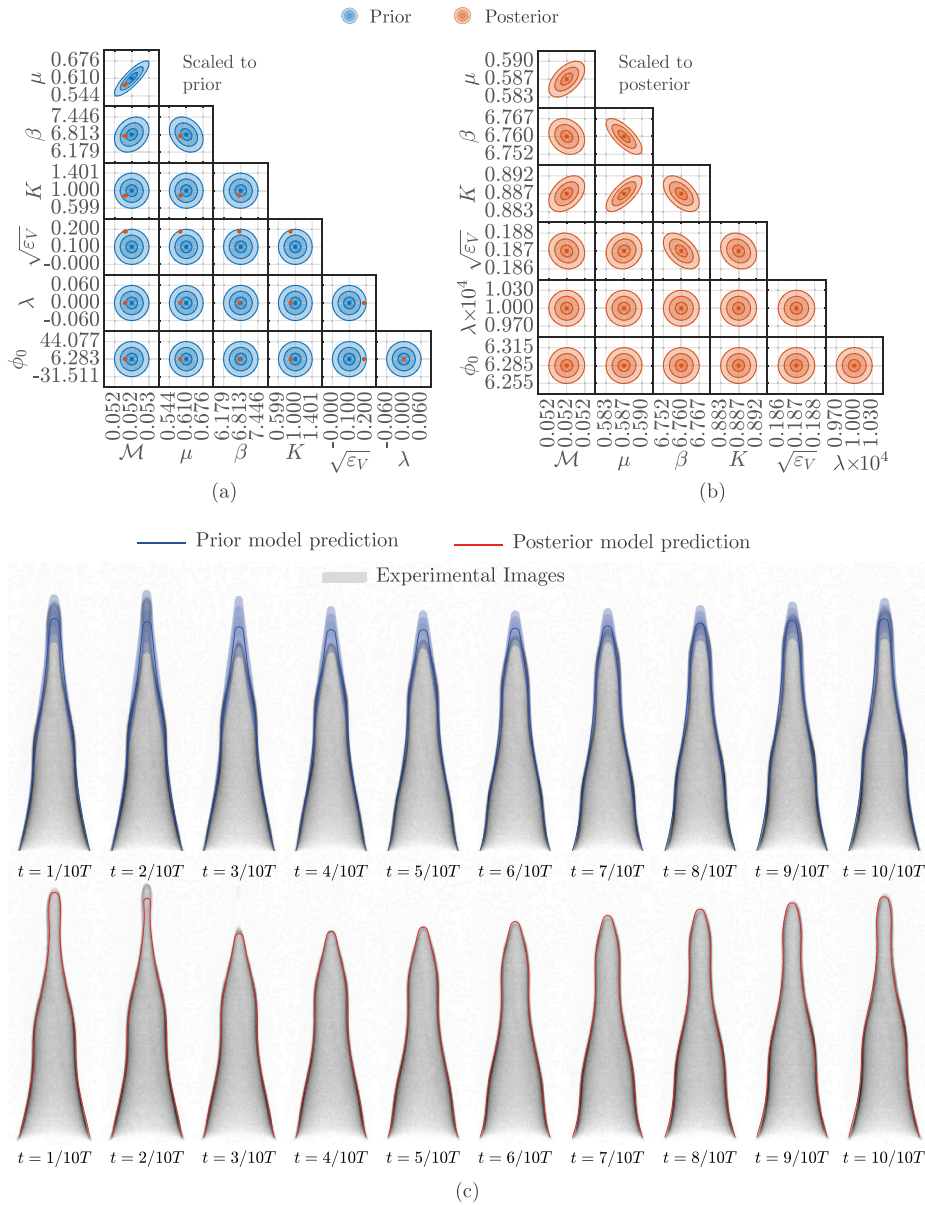
## 5.3. Flame transfer function predictions

With the model parameters inferred from the available data, we use the calibrated model to predict the flame transfer functions for each of the 20 flames with quantified uncertainty bounds. We introduce the flame transfer function  $\mathcal{F}$  as:

$$\mathcal{F}(St; \phi, \bar{u}, \alpha) = \frac{\hat{q}/\bar{q}}{\hat{u}_{ref}/\bar{u}}, \quad (17)$$

where  $\hat{\star}$  denotes the Fourier component of  $\star$  at the Strouhal number  $St$ ;  $q$  denotes the fluctuating heat release rate and  $u_{ref}$  denotes the velocity fluctuations at a reference position. The heat release rate  $q$  is normalized by the mean heat release rate, while the velocity  $\hat{u}_{ref}$  is normalized by the mean flow velocity in the burner tube  $\bar{u}$ . In this study, we define the reference velocity  $\hat{u}_{ref}$  as the component of the





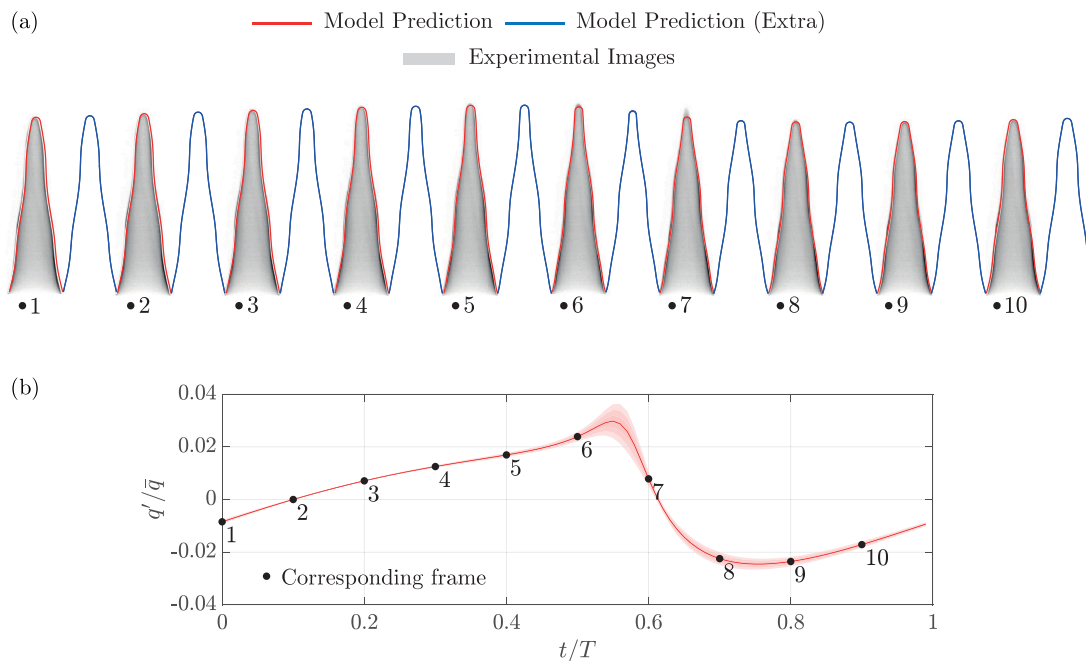
**Fig. 7.** (a-b) Probability distributions of parameters for the acoustically forced case; prior (blue) and posterior (red). Contours depict 1, 2, and 3 standard deviations from the prior  $\mathbf{a}_p$  and posterior  $\mathbf{a}_{MP}$  estimations. Both panels represent the same quantities; however, panel (a) is scaled to the prior distributions, while panel (b) is scaled to the posterior distributions. Parameters include the flame aspect ratio  $\beta$ , which is a function of the unstretched laminar flame speed  $S_L$  and the bulk velocity in the burner tube  $\bar{u}$ , the shape parameter for the velocity field  $\mu$ , the nondimensional Markstein length  $\mathcal{M}$ , the nondimensional velocity perturbation convective speed  $K$ , and the amplitude of velocity perturbations  $\varepsilon_V$ . (c) Comparison of model predictions before (red line) and after (blue line) assimilating the experimental video footage. The confidence interval, spanning 3 standard deviations, is indicated by red and blue shadings overlaid on the experimental flame images. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

velocity field that is perpendicular to the flame front at its base, scaled by  $1/\sin(\alpha_b)$ , where  $\alpha_b$  is the angle between the flame base and the axial direction. This definition aligns with the formulations provided by Orchini & Juniper [35] and Preetham et al. [36]. This scaling is necessary because, in the velocity model employed in our study and theirs, the radial velocity component vanishes as the Strouhal number approaches zero, resulting in purely axial flow. In this limit, the scaling ensures that the magnitude of the reference velocity corresponds to that of the axial velocity, allowing the gain of the flame transfer function to approach 1. While this simple velocity model has clear limitations, it allows us to demonstrate our Bayesian framework in a simple setting. In future work we will refine the velocity model.

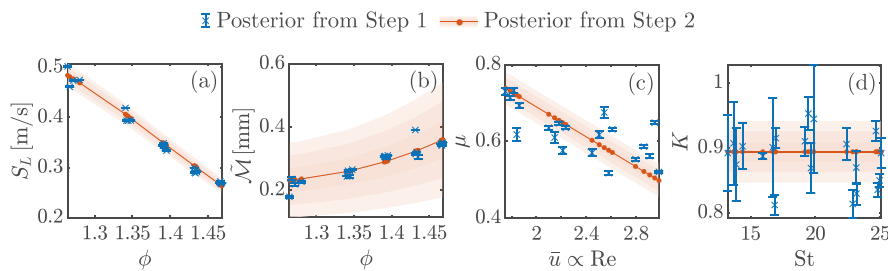
Using Eq. (17) and our calibrated model, we calculate a flame transfer function for each of the 20 flames. The polar plot in Fig. 11(a) shows the predicted flame responses, and their uncertainties, evaluated at the

experimentally observed frequency. Each flame response is coloured according to its corresponding flame group as in Fig. 2. It is worth noting that the flames within the same group, which share the same convective time delay  $\tau$ , exhibit similar phase delays on the polar plot, as one would expect.

Fig. 11(b) shows the predicted flame transfer functions for each of the 20 flames, along with their uncertainties, over a wide range of Strouhal numbers (St). For reference, the Strouhal numbers at which data were collected are highlighted with a dot. We see that the parametric uncertainty (shown with a shaded patch) is low throughout the frequency range for each of the 20 flame transfer functions. This indicates that experimental data at a single forcing frequency contains enough information about the unknown parameters to collapse our parametric uncertainty over a wide range of frequencies. We note, however, that without observations at other forcing frequencies, we cannot



**Fig. 8.** Diagram showing the concept of using Bayesian inference to create a quantitatively accurate model capable of providing physical information not directly available from experiments. The experimental data consists of ten frames during one period of oscillation. Panel (a) shows the model prediction (red line) with a confidence interval of 3 standard deviations (red shading) plotted on top of the experimental flame images, as in Fig. 7. We also use the model to infer the flame position (blue lines) and its uncertainty (blue shading) at moments between observations. Panel (b) shows the h.r.r. perturbation normalized by its mean value, computed using the model. The black dots correspond to the timeframes observed in the experiments. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



**Fig. 9.** Results of the second step of the data assimilation process. From the first step of the process, we obtain the trends of the physical parameters with respect to the equivalence ratio  $\phi$  and the bulk velocity  $\bar{u}$ , which is proportional to the Reynolds number,  $Re$ , inside the burner tube. Then we propose a general model for the parameters and assimilate all the flame conditions together. The blue markers represent the posterior values from the first step of the data assimilation process and the error bars correspond to 3 standard deviations. The red dotted lines correspond to the posterior values from the second step and the shaded areas represent 1, 2 and 3 standard deviations. Panel (a) shows the assimilated values of the unstretched laminar flame speed  $S_L$  plotted against the equivalence ratio  $\phi$ . Panel (b) shows the assimilated values of the dimensional Markstein length  $\bar{M}$  plotted against the equivalence ratio  $\phi$ . Panel (c) shows the assimilated values of the velocity shape parameter  $\mu$  plotted against the bulk velocity in the burner tube  $\bar{u}$ . Panel (d) shows the assimilated values of the convective speed ratio  $K$  plotted against the acoustic perturbation Strouhal number  $St$ .

identify other sources of uncertainty such as systematic experimental uncertainty or structural model uncertainty.

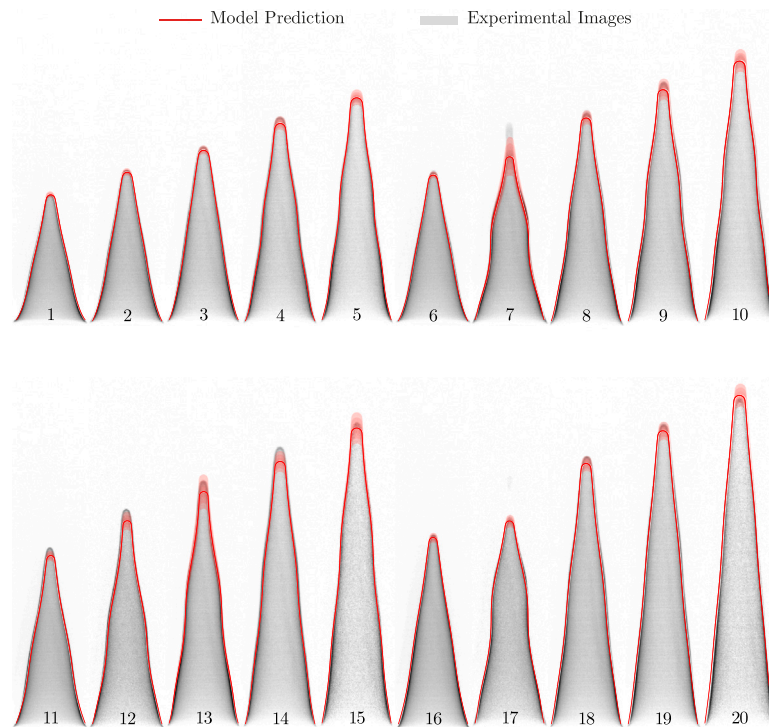
#### 5.4. Extrapolation

The model in this paper contains a small number of parameters because it is physics-based. Consequently, the training process requires only a small amount of data. In this section, we show the model's ability to predict the flame front dynamics of the 20 flames in our library using data from just four flames. We train the model's hyperparameters on flames 1, 6, 11, and 16 of Table A.1 and predict the remaining flames with this model. These four flames have the same convective time delay, but different powers. We therefore test the ability of the model to extrapolate outside the range of trained convective time delays, which is more challenging than scaling with power.

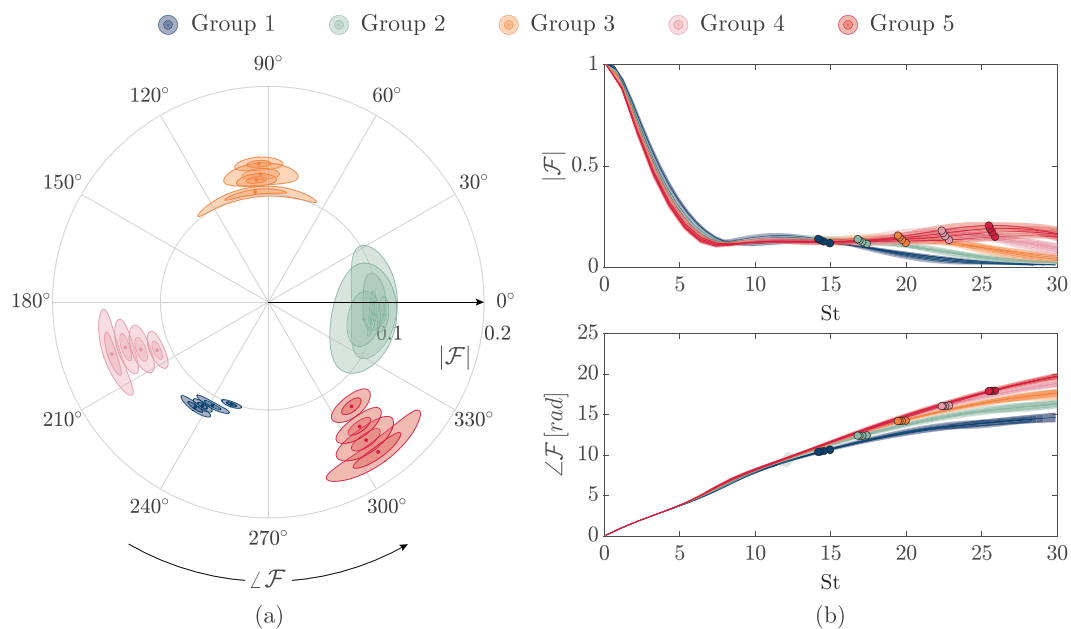
Fig. 12 shows the predicted flame front dynamics for all flames at a single timestep. Red lines show the flames that were used for training and blue lines show the flames used for testing. The red

and blue shading denotes 1, 2, and 3 standard deviations from the predicted flame front position. All model predictions are overlaid onto the experimental flame images for comparison. The number at the base of each flames indicates the flame numbers in Table A.1. In the supplementary material, we provide the animated version of this figure, showing the predictions for all timesteps of the experimental video footage.

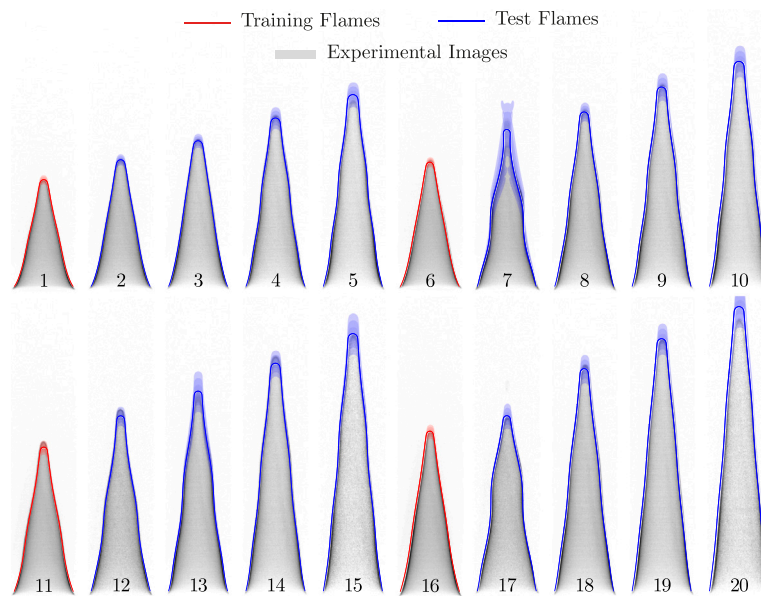
We now compare Figs. 10 and 12. When only one-fifth of the available data has been used, the predicted flame front positions are less accurate and have larger uncertainty. This is particularly clear for flame 7, which has a large acoustic amplitude and is close to pinch off in these figures. Indeed, the model shows significant sensitivity to the model parameters as the acoustic amplitude increases. Similarly, the prediction for flame 20 is less accurate because the flame length is much larger than the length of the four flames used for model training. Even with such a small amount of training data, however, the data still falls within the uncertainty bounds of the model predictions, even when the data is well outside the training range. We further note that Fig. 12 shows the timestep with the largest model misfit and uncertainties. In



**Fig. 10.** The general model predictions for all flames in the library are represented by the red line, with confidence intervals of 3 standard deviations (depicted as red shading), overlaid on the experimental flame images. The numbers at the base of the flames indicate the corresponding flame numbers referenced in Table A.1. Additionally, in the supplementary material, we provide an animated version of this figure, showing predictions at all time steps of the experimental video footage. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



**Fig. 11.** (a) Predicted flame transfer functions at the experimentally observed frequency, presented in a polar plot. Each curve is coloured according to the flame group as shown in Fig. 2. Flames within the same group, sharing the same convective time delay  $\tau$ , exhibit similar phase delays, as expected. (b) Predicted flame transfer functions, with associated uncertainties, over a broad range of Strouhal numbers ( $St$ ). Strouhal numbers where experimental data were available are marked with dots. The shaded regions represent the parametric uncertainty, which remains low across the frequency range for all 20 flame transfer functions. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



**Fig. 12.** Predicted flame front dynamics for all flames in the library at a single timestep. Red lines show flames used for training, while blue lines show flames used for testing. The shaded regions in red and blue indicate 1, 2, and 3 standard deviations from the optimal flame front estimation. Additionally, all model predictions are overlaid onto the experimental flame images for comparison. In the supplementary material, we provide an animated version of this figure, presenting predictions for all time steps of the experimental video footage. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the supplementary material, we provide an animated version of this figure showing the predictions for all timesteps of the video footage.

## 6. Conclusions

In this study, we perform experiments on an acoustically forced, laminar premixed conical flame in a duct. We use a high-speed camera to record snapshots of the natural emission of the flame while steady and forced. We propose a physics-based reduced-order model that provides, amongst other things, the flame front dynamics under harmonic forcing. We combine the model output with the experimental images of the flame to infer the most probable model parameters. This process (i) turns a qualitatively-accurate model into a quantitatively-accurate model, and (ii) quantifies the uncertainty in the inferred model parameters and the model predictions.

The inference process produces a digital twin of the flame, which provides access to quantities that were not directly measured in the experiments. From observations of the steady flame, we can estimate combustion characteristics such as the laminar flame speed and the Markstein length. This can be used to characterize the combustion properties of a fuel for which data are not available. From observations of the perturbed flame, we can infer the fluctuating heat release rate as a response to the velocity perturbation, which is then used to calculate the flame's thermoacoustic response through, for example, the flame transfer function. The proposed model contains few parameters and can therefore be trained on relatively little data. To demonstrate this we infer the model's parameters using one-fifth of the available data and successfully predict the remaining flames in the dataset. The model is also differentiable with respect to the parameters meaning that it can be used as a design optimization tool.

In future work, the flame model will be included within a larger model of the thermoacoustic behaviour of the system, whose other parameters are tuned in the same way. The final result will be a quantitatively-accurate physics-based model of the flame and the thermoacoustic system that is interpretable, trustworthy, and extrapolatable. The results of the current work are limited to simple flames due to the simplicity of the physics-based model. In future work we will increase the fidelity of the model so that the inference framework can be applied to more interesting flames.

The Matlab code developed in this study is made available to the reader, who is invited to modify and apply it to assimilate other flame image data into the flame model.

## CRedit authorship contribution statement

**Alessandro Giannotta:** Writing – original draft, Visualization, Software, Investigation, Formal analysis, Data curation. **Matthew Yoko:** Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Stefania Cherubini:** Writing – review & editing, Formal analysis, Conceptualization. **Pietro De Palma:** Writing – review & editing, Formal analysis, Conceptualization. **Matthew P. Juniper:** Writing – review & editing, Supervision, Methodology, Investigation, Formal analysis, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

This research was partly supported under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.4 - Call for “National Centres” from research to business, funded by the European Union – NextGenerationEU. Project code CN00000023, Concession Decree No. 1033 adopted by Ministero dell'Università e della Ricerca (MUR), CUP - D93C22000410001, Project title “MOST - National Center for Sustainable Mobility”. This research was partly carried out within the NEST - Network 4 Energy Sustainable Transition (D.D. 1243 02/08/2022, PE00000021) and received funding under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.3, funded from the European Union - NextGenerationEU. This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them. Matthew Yoko acknowledges funding for his PhD from The Cambridge Trust, The Skye Foundation and the Oppenheimer Memorial Trust.



**Table A.1**

Summary of the properties of the 20 flames studied. We show the average measured flow rates of air, methane (CH<sub>4</sub>) and ethylene (C<sub>2</sub>H<sub>4</sub>), the equivalence ratio ( $\phi$ ), the bulk velocity in the burner tube ( $\bar{u}$ ), the measured flame length ( $L_f$ ), and the mean heat release rate ( $\bar{Q}$ ).

Flame No.	Air [ln/min]	CH <sub>4</sub> [ln/min]	C <sub>2</sub> H <sub>4</sub> [ln/min]	$\phi$ [-]	$\bar{u}$ [m/s]	$L_f$ meas. [mm]	$\bar{Q}$ [W]
1	6.049	0.325	0.325	1.28	1.75	16.6	375
2	6.147	0.348	0.348	1.35	1.79	20.0	375
3	6.219	0.364	0.364	1.40	1.82	23.1	375
4	6.283	0.379	0.379	1.44	1.84	26.4	375
5	6.338	0.391	0.391	1.47	1.86	28.9	375
6	7.246	0.387	0.387	1.27	2.10	19.8	450
7	7.369	0.416	0.416	1.34	2.15	23.8	450
8	7.459	0.436	0.436	1.39	2.18	27.4	450
9	7.537	0.454	0.454	1.43	2.21	30.9	450
10	7.603	0.468	0.468	1.47	2.24	34.2	450
11	8.444	0.449	0.449	1.27	2.45	23.1	525
12	8.594	0.484	0.484	1.34	2.51	27.6	525
13	8.699	0.508	0.508	1.39	2.55	31.7	525
14	8.790	0.529	0.529	1.43	2.58	36.4	525
15	8.868	0.546	0.546	1.47	2.61	39.2	525
16	9.644	0.512	0.512	1.26	2.80	25.9	600
17	9.818	0.553	0.553	1.34	2.86	30.1	600
18	9.939	0.580	0.580	1.39	2.91	34.8	600
19	10.045	0.604	0.604	1.43	2.95	39.9	600
20	10.134	0.624	0.624	1.47	2.99	43.6	600

## Appendix A. Properties of the flames studied in this paper

See Table A.1.

## Appendix B. Adjoint methods for computing the sensitivity of a dynamical system to variations in initial conditions and model parameters

Consider the general dynamical system:

$$\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x}(t), t; \mathbf{a}) = 0, \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (\text{B.1})$$

where  $\mathbf{x}(t)$  is the state vector, at a given time  $t$ ,  $\mathbf{x}_0$  is the initial state and  $\mathbf{a}$  represents parameters. Any small variation in the initial condition  $\delta\mathbf{x}(0)$  and in the system parameters  $\delta\mathbf{a}$  will result in variations in the state  $\delta\mathbf{x}(t)$  and  $\delta\mathbf{x}(T)$ , where  $\mathbf{x}(T)$  is the state vector at a final time  $T$ . Assuming the system dynamics expressed by Eq. (B.1) hold at all times, the variation of the state is governed by:

$$\frac{d\delta\mathbf{x}}{dt} - \frac{\partial\mathbf{f}}{\partial\mathbf{x}}\delta\mathbf{x} - \frac{\partial\mathbf{f}}{\partial\mathbf{a}}\delta\mathbf{a} = 0 \quad (\text{B.2})$$

To derive the adjoint equation, we multiply this equation by a vector  $\mathbf{w}^T$  having the same size of vector  $\mathbf{x}$  and integrate over the time interval  $[0, T]$ :

$$\int_0^T \mathbf{w}^T \left( \frac{d\delta\mathbf{x}}{dt} - \frac{\partial\mathbf{f}}{\partial\mathbf{x}}\delta\mathbf{x} - \frac{\partial\mathbf{f}}{\partial\mathbf{a}}\delta\mathbf{a} \right) dt = 0. \quad (\text{B.3})$$

The vector  $\mathbf{w}$  is called the *adjoint vector*. Integration by parts yields:

$$\begin{aligned} \mathbf{w}^T(T)\delta\mathbf{x}(T) - \mathbf{w}^T(0)\delta\mathbf{x}(0) - \int_0^T \left( \frac{d\mathbf{w}^T}{dt} + \mathbf{w}^T \frac{\partial\mathbf{f}}{\partial\mathbf{x}} \right) \delta\mathbf{x} dt \\ - \int_0^T \left( \mathbf{w}^T \frac{\partial\mathbf{f}}{\partial\mathbf{a}} \right) \delta\mathbf{a} dt = 0. \end{aligned} \quad (\text{B.4})$$

To ensure the integrand is zero for arbitrary  $\delta\mathbf{x}$ , we impose the following equation:

$$\frac{d\mathbf{w}^T}{dt} + \mathbf{w}^T \frac{\partial\mathbf{f}}{\partial\mathbf{x}} = 0, \quad (\text{B.5})$$

which is solved backward in time from  $T$  to 0. Eq. (B.5) is called the *adjoint equation*. With the adjoint equation satisfied, we are left with:

$$\mathbf{w}^T(T)\delta\mathbf{x}(T) - \mathbf{w}^T(0)\delta\mathbf{x}(0) - \int_0^T \left( \mathbf{w}^T \frac{\partial\mathbf{f}}{\partial\mathbf{a}} \right) \delta\mathbf{a} dt = 0. \quad (\text{B.6})$$

This equation relates variations in the final state to variations in the initial conditions. To compute the sensitivity of the  $i$ th component of the state at time  $T$  to the initial condition, we set:

$$\mathbf{w}(T) = [0, \dots, 0, 1_{\text{ith}}, 0, \dots, 0]^T \quad (\text{B.7})$$

and solve the adjoint Eq. (B.5). Eq. (B.6) can be then expressed by:

$$\delta x_i(T) = \mathbf{w}^T(0)\delta\mathbf{x}(0) + \int_0^T \left( \mathbf{w}^T \frac{\partial\mathbf{f}}{\partial\mathbf{a}} \right) \delta\mathbf{a} dt. \quad (\text{B.8})$$

Finally, the  $i$ th row of the Jacobian matrix with respect to the initial condition is given by:

$$\frac{\partial x_i(T)}{\partial \mathbf{x}(0)} = \mathbf{w}^T(0). \quad (\text{B.9})$$

Analogously, the  $i$ th row of the Jacobian matrix with respect to the model parameters is given by:

$$\frac{\partial x_i(T)}{\partial \mathbf{a}} = \int_0^T \left( \mathbf{w}^T \frac{\partial\mathbf{f}}{\partial\mathbf{a}} \right) dt. \quad (\text{B.10})$$

## Appendix C. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.combustflame.2025.114011>.

## Data availability

The experimental data is openly available at <http://doi.org/10.17863/CAM.106960>. The Matlab code is available at [https://github.com/aleGiannotta/Bayesian\\_Inference\\_Lam\\_Flame.git](https://github.com/aleGiannotta/Bayesian_Inference_Lam_Flame.git).

## References

- [1] J.W.S.B. Rayleigh, The explanation of certain acoustical phenomena, *Nature* 18 (1878) 319–321.
- [2] B.T. Chu, On the energy transfer to small disturbances in fluid flow (Part I), *Acta Mech.* 1 (3) (1965) 215–234, <http://dx.doi.org/10.1007/BF01387235>.
- [3] T.C. Lieuwen, V. Yang, *Combustion Instabilities In Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modelling*, vol. 210, American Institute of Aeronautics and Astronautics, 2005.
- [4] M.P. Juniper, R.I. Sujith, Sensitivity and nonlinearity of thermoacoustic oscillations, *Annu. Rev. Fluid Mech.* 50 (2018) 661–689.
- [5] H.C. Mongia, T. Held, G. Hsiao, R. Pandalai, Challenges and progress in controlling dynamics in gas turbine combustors, *J. Propuls. Power* 19 (2003) 822–829.
- [6] M.P. Juniper, M. Yoko, Generating a physics-based quantitatively-accurate model of an electrically-heated Rijke tube with Bayesian inference, *J. Sound Vib.* 535 (2022) 117096, <http://dx.doi.org/10.1016/j.jsv.2022.117096>.
- [7] M. Yoko, M.P. Juniper, Minimizing the data required to train a physics-based thermoacoustic model, in: 29th International Congress on Sound and Vibration, 2023.
- [8] T. Poinsot, Prediction and control of combustion instabilities in real engines, *Proc. Combust. Inst.* 36 (2017) 1–28.
- [9] P. Palies, D. Durox, T. Schuller, S. Candel, The combined dynamics of swirler and turbulent premixed swirling flames, *Combust. Flame* 157 (9) (2010) 1698–1717.
- [10] J.P. Moeck, J.-F. Bourgoin, D. Durox, T. Schuller, S. Candel, Nonlinear interaction between a precessing vortex core and acoustic oscillations in a turbulent swirling flame, *Combust. Flame* 159 (8) (2012) 2650–2668.
- [11] J.P. Moeck, J.-F. Bourgoin, D. Durox, T. Schuller, S. Candel, Tomographic reconstruction of heat release rate perturbations induced by helical modes in turbulent swirl flames, *Exp. Fluids* 54 (2013) 1–17.
- [12] J. Li, D. Durox, F. Richecoeur, T. Schuller, Analysis of chemiluminescence, density and heat release rate fluctuations in acoustically perturbed laminar premixed flames, *Combust. Flame* 162 (10) (2015) 3934–3945.
- [13] Z. Han, S. Balusamy, S. Hochgreb, Spatial analysis on forced heat release response of turbulent stratified flames, *J. Eng. Gas Turbines Power* 137 (2015).
- [14] R. Yuan, J. Kariuki, A. Dowlut, R. Balachandran, E. Mastorakos, Reaction zone visualisation in swirling spray n-heptane flames, *Proc. Combust. Inst.* 35 (2015) 1649–1656.
- [15] M. Yoko, M.P. Juniper, Optimal experiment design with adjoint-accelerated Bayesian inference, *Data-Centric Eng.* 5 (2024) 1–28.
- [16] M. Yoko, M.P. Juniper, Adjoint-accelerated Bayesian inference applied to the thermoacoustic behaviour of a ducted conical flame, *J. Fluid Mech.* 985 (2024) A38, <http://dx.doi.org/10.1017/jfm.2024.276>.

- [17] M. Yoko, M.P. Juniper, Inferring flame transfer functions of turbulent conical flames from pressure measurements, *J. Eng. Gas Turbines Power* (2024) <http://dx.doi.org/10.1115/1.4066616>.
- [18] A. Giannotta, S. Cherubini, P. De Palma, M.P. Juniper, The effect of flame curvature and flame base movement on the frequency response of a conical Bunsen flame, *Combust. Flame* 259 (2024) 113179.
- [19] C.J. Dasch, One-dimensional tomography: a comparison of Abel, onion-peeling, and filtered backprojection methods, *Appl. Opt.* 31 (1992) 1146, <http://dx.doi.org/10.1364/ao.31.001146>.
- [20] T. Chan, L. Vese, An active contour model without edges, in: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 1682, 1999, pp. 141–151, <http://dx.doi.org/10.1007/3-540-48236>.
- [21] T. Schuller, S. Ducruix, D. Durox, S. Candel, Modelling tools for the prediction of premixed flame transfer functions, *Proc. Combust. Inst.* 29 (2002) 107–113.
- [22] K. Kashinath, S. Hemchandra, M.P. Juniper, Nonlinear thermoacoustics of ducted premixed flames: The influence of perturbation convection speed, *Combust. Flame* 160 (12) (2013) 2856–2865.
- [23] A. Orchini, M.P. Juniper, Linear stability and adjoint sensitivity analysis of thermoacoustic networks with premixed flames, *Combust. Flame* 165 (2016) 97–108.
- [24] M. Matalon, Intrinsic flame instabilities in premixed and nonpremixed combustion, *Annu. Rev. Fluid Mech.* 39 (2007) 163–191.
- [25] H. Yu, M.P. Juniper, L. Magri, A data-driven kinematic model of a ducted premixed flame, *Proc. Combust. Inst.* 38 (4) (2021) 6231–6239.
- [26] P. Luchini, A. Bottaro, Adjoint equations in stability analysis, *Annu. Rev. Fluid Mech.* 46 (1) (2014) 493–517.
- [27] M.P. Juniper, Triggering in the horizontal rijke tube: non-normality, transient growth and bypass transition, *J. Fluid Mech.* 667 (2011) 272–308.
- [28] A. Giannotta, S. Cherubini, P. De Palma, Minimal energy thresholds for triggering in the Rijke tube: the effect of the time delay, *J. Fluid Mech.* 938 (2022) A23, <http://dx.doi.org/10.1017/jfm.2022.149>.
- [29] H. Jeffreys, *Scientific Inference*, third ed., Cambridge University Press, 1973.
- [30] D.J. MacKay, Information-based objective functions for active data selection, *Neural Comput.* 4 (1992) 590–604.
- [31] M.M. Putko, A.C. Taylor, P.A. Newman, L.L. Green, Approach for input uncertainty propagation and robust design in CFD using sensitivity derivatives, *J. Fluids Eng. Trans. ASME* 124 (2002) 60–69, <http://dx.doi.org/10.1115/1.1446068>.
- [32] D.G. Goodwin, H.K. Moffat, R.L. Speth, *Cantera: An Object-Oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes*, version, 2018.
- [33] K. Kashinath, I.C. Waugh, M.P. Juniper, Nonlinear self-excited thermoacoustic oscillations of a ducted premixed flame: bifurcations and routes to chaos, *J. Fluid Mech.* 761 (2014) 399–430.
- [34] A. Orchini, M.P. Juniper, Linear stability and adjoint sensitivity analysis of thermoacoustic networks with premixed flames, *Combust. Flame* 165 (2015) 97–108, <http://dx.doi.org/10.1016/j.combustflame.2015.10.011>.
- [35] A. Orchini, M.P. Juniper, Flame double input describing function analysis, *Combust. Flame* 171 (2016) 87–102.
- [36] Preetham, H. Santosh, T. Lieuwen, Dynamics of laminar premixed flames forced by harmonic velocity disturbances, *J. Propuls. Power* 24 (6) (2008) 1390–1402.