Bayesian parameter inference of a vortically perturbed flame model for the prediction of thermoacoustic instability

Anonymous Author(s)
Affiliation
Address
email

Abstract

Thermoacoustic instabilities can be highly detrimental to the operation of aircraft gas turbine combustors within design conditions, and hence their prediction and suppression are crucial. This work uses a Bayesian machine learning method to infer the parameters of a bluff-body stabilised, physics-informed flame model in real-time. The flame front is modelled using the G-equation, a level-set method which segments the flow into regions of reactants and products. The flow past the bluff-body is modelled with a discrete vortex method (DVM) to account for vortical perturbations on the flame front. Using the physics-informed model with the learned parameters from both the G-equation and the DVM, a flame transfer function (FTF) is obtained, from which the growth rates of instability in the system can be calculated. A heteroscedastic Bayesian neural network ensemble (BayNNE) is trained on a library of flame front simulations with known target parameters in both models. The trained BayNNE is a surrogate model for a Bayesian posterior of the target parameters given the input flame front coordinates. The ensemble predicts some parameters of the DVM with more certainty than others, showing which are more influential in affecting the flame front. Using the learned posterior, the flame fronts are re-simulated, to extrapolate the flame beyond the experimental window where it was observed. Flame results are also extrapolated in parameter space. These extrapolated flame shapes are then used to calculate thermoacoustic frequencies and growth rates of the system. We observe that the growth rates and frequencies do not show a strong dependency on the amplitude of forcing, which is one of the inferred parameters of the physics-informed model. This important result suggests that a FTF derived at high amplitude, when it is observable, is also valid at low amplitude, when it is not observable.

1 Introduction

1.1 Thermoacoustics

The prediction and control of thermoacoustic instability is a persistent challenge in jet and rocket engine design [1]. The drive towards lower $\mathrm{NO_x}$ emissions in gas turbines has led to the use of 'lean premixed pre-vaporised' (LPP) combustion [2]. Although LPP combustion considerably reduces $\mathrm{NO_x}$, it is particularly susceptible to unfavourable thermoacoustic instabilities [3], which are defined as follows. In wall-confined ducted environments such as aircraft gas turbine combustors or rocket

Submitted to NeurIPS 2022 AI for Science Workshop.

motors, acoustic standing waves can be sustained. When the heat release rate perturbations during combustion (q') are in phase with the duct acoustic pressure (p'), a positive feedback loop is generated, driving p' and q' to a thermoacoustically unstable state, with high amplitudes [4] that can cause structural damage and flame blowoff.

In premixed flames, q' is predominantly due to flame surface area perturbations generated by fluctuations in the velocity field along it [5; 6]. Any physics-based model must therefore contain the flame's response to velocity fluctuations. The most accurate flame response can be evaluated with high-fidelity computational fluid dynamics (CFD) simulations but, these are computationally expensive. In this paper we use data to tune the parameters of physics-based reduced-order models for the flame and flow fields in a thermoacoustic system, in order to reduce the computational cost while retaining as much accuracy as possible in flame response.

The physical system under study is a version of the Volvo burner [7; 8], a schematic of which is shown in figure 1 (Sec. 2.1). We model the flame as the zero contour (or level-set) of a continuous function that advects with the flow. This is known as the G-equation model [9] and allows the flame dynamics to be simulated cheaply (Sec. 2.2). The velocity field is given by the discrete vortex method (DVM), which is derived directly from the Navier-Stokes equations and includes the effect of vortical perturbations on the flame front (Sec. 2.3). The flame and flow models are thus physics-based.

In order to render the models quantitatively accurate, their parameters need to be assimilated from experimental data. The ensemble Kalman filter [10] (EnKF) has been used previously to assimilate data into the G-equation model [11; 12]. The EnKF performs Bayesian inference to infer the parameters of the G-equation model by statistically combining model forecasts with experimental measurements. However, the computational requirements of the EnKF render online Bayesian inference unfeasible when measurements are taken at high frequency. Furthermore, the method can fail to infer parameters when these vary quickly in time, and suffers from numerical stability issues when the measurements are noisy. These are both the case for our experimental data, which come from high frequency OH planar laser induced fluorescence (OH PLIF) measurements of a lab-scale version of the Volvo combustor rig [7; 8].

1.2 Bayesian deep learning

Bayesian deep learning refers to the use of deep learning algorithms, such as deep artificial neural networks (NNs), for Bayesian inference. Bayesian NNs [13] replace the point estimates of each of the NN's weights and biases with Gaussian probability distributions, with means and variances learned during training. The distribution of every weight and bias in the NN can be used to infer the outputs from the inputs, for example inferring the parameters of a model from experimental measurements. Unfortunately, Bayesian NNs of practical size are too expensive to train [14]. More recently, ensembles of deep, wide NNs have been used to perform approximate Bayesian inference [15; 16; 17; 18; 19; 20], with the approximation improving with increasing width of the NN's hidden layers. These Bayesian NN ensembles (BayNNEs) learn the mean and variance of the posterior distribution of the outputs given the inputs. When multiple outputs are to be inferred, heteroscedastic BayNNEs learn the means and variances of each output, without assuming a common variance for all outputs. This study uses heteroscedastic BayNNEs to infer the parameters of the velocity field of the G-equation model given experimental observations.

2 The Volvo burner

2.1 Experiment

This data is taken from experiments performed on a version of the Volvo burner shown schematically in figure 1. Premixed air and propane flow into the burner and are burned by a flame stabilized on a triangular bluff body with side length $D = 3.8$ cm. The depth of the duct is $4D$. Images of the flame are recorded at 10 kHz using OH PLIF through a window $3D$ tall and $3.4D$ wide. A sequence of 7998 images each with resolution 69×58 pixels is recorded. Additionally, the horizontal and vertical

velocities within the observation window at a resolution of 69×58 pixels are recorded at each time step. As the air-fuel mixture flows through the burner, vortices are shed periodically from the bluff body. These vortices cause wrinkling and cusping of the flame front. For full details of the Volvo experiment, the reader is referred to Caswell et al. [7] and Fugger et al. [8].

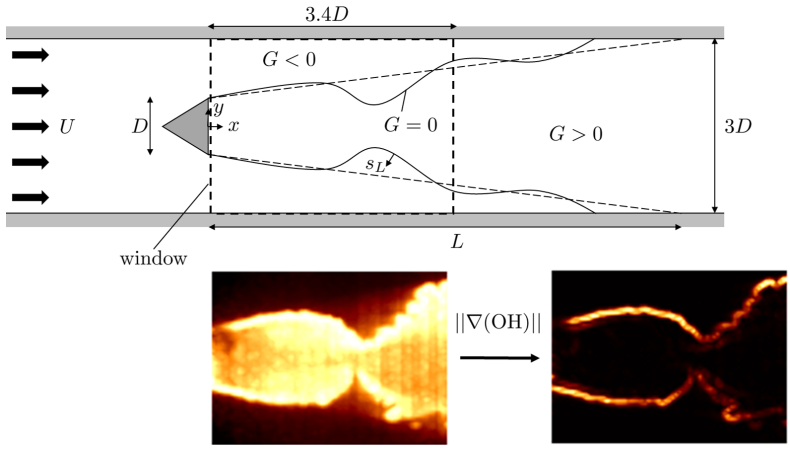

Figure 1: Diagram of the Volvo combustor rig and G-equation model of the flame. As the air-fuel mixture flows through the combustor, vortex shedding causes wrinkling and cusping of the flame front, represented by the $G = 0$ contour of a continuous scalar field $G(x, y, t)$. An example OH intensity image taken through the window is shown, as well as the pre-processing step to find the flame front.

The Volvo flame images are processed and the flame front is extracted as a radial location y, which is a singularly-valued function of the axial co-ordinate x: $y = f(x)$. Although the flame shows greater wrinkling than the Bunsen flame, it does not wrinkle back onto itself and so a $y = f(x)$ discretisation of the flame front is still possible. Starting from an OH PLIF intensity image, the bottom half is discarded because the G-equation model assumes a symmetric flame front. Furthermore, the rightmost 9 columns of data are discarded from each image due to the fact that the flame front is occasionally too faint to be detected here. The flame front is found by thresholding the magnitude of the gradient vector. Next, splines with ten knots are used to smooth $y(x)$. Each flame image is therefore converted into a 60×1 vector of flame front y locations $\mathbf{y}$. The x coordinates are the same for all flames, so are discarded. Observation vectors $\mathbf{z}$ are created by appending ten consecutive $\mathbf{y}$ vectors together. All 7998 images taken from the Volvo burner experiment are processed in this way.

2.2 The flame front model

The flame front is assumed to be a thin boundary between unburnt and burnt gases (see figure 1). The flame travels normal to itself into the unburnt gases with laminar flame speed s_L which depends on the gas composition. The velocity in the burnt gases does not affect the flame kinematics. The unburnt and burnt gases are assumed to travel with velocity $\mathbf{u}(x, y, t)$. Under these assumptions, the flame front is modelled by the $G(x, y, t) = 0$ contour (or level-set) of a continuous scalar field G whose motion is governed by the G-equation:

$$\frac{\partial G}{\partial t} + \mathbf{u} \cdot \nabla G = s_L |\nabla G| \,, \tag{1}$$

where the flame speed s_L is a function of the fuel composition and $\mathbf{u}$ is calculated using a discrete vortex method.

2.3 The velocity field model

The velocity field $\mathbf{u}(\mathbf{x}, t)$ is calculated using the discrete vortex method (DVM) [21; 22]. This method aims to model the large-scale vortical structures that are known to significantly contribute to

flame front oscillations during thermoacoustic instability [23; 24; 25; 26; 27]. It uses the vorticity formulation of the incompressible Euler equations in two dimensions, given as

$$\frac{\partial \boldsymbol{\omega}}{\partial t} + \mathbf{u} \cdot \nabla \boldsymbol{\omega} = \boldsymbol{\omega} \cdot \nabla \mathbf{u} \ , \tag{2}$$

$$\nabla \cdot \mathbf{u} = 0 \ , \tag{3}$$

where $\boldsymbol{\omega}(\mathbf{x}, t)$ is the vorticity field. This vorticity field is discretised into N_V individual fluid elements, each with position $\mathbf{x}_i(t)$, core radius $r_{D,0}$ and circulation Γ_i, where $i = 1, 2 \ldots N_V$:

$$\boldsymbol{\omega} = \sum_{i}^{N_V} \Gamma_i f_\delta(r) \ , \tag{4}$$

where the core function $f_\delta(r)$ characterises the distribution of vorticity within the vortex element. The vortices are advected by the flow, whose velocity field is the solution to the Poisson equation $\nabla^2 \mathbf{u} = -\nabla \times \boldsymbol{\omega}$. A Helmholtz decomposition is applied to $\mathbf{u}$:

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}(\mathbf{x}, t)_{\text{irr}} + \mathbf{u}(\mathbf{x}, t)_{\text{rot}} \ . \tag{5}$$

The irrotational velocity field $\mathbf{u}(\mathbf{x}, t)_{\text{irr}}$ is the solution to a potential flow problem, calculated via a Schwartz-Christoffel conformal mapping [28]. The acoustic forcing on the flow field is modelled with a harmonically oscillating inlet flow velocity $u(t)/U = 1 + \alpha \sin(\text{St}\, t)$ where α is the amplitude of forcing. The rotational velocity field is obtained with a Green's function solution of the Poisson equation above:

$$\mathbf{u}(\mathbf{x}, t)_{\text{rot}} = \nabla \mathcal{G}(\mathbf{x}, \mathbf{x}') * \boldsymbol{\omega}^h \ , \tag{6}$$

where $*$ is the convolution operator and $\mathcal{G}(\mathbf{x}, \mathbf{x}')$ is the Green's function for the Laplacian operator in two dimensions:

$$\nabla^2 \mathcal{G}(\mathbf{x}, \mathbf{x}') = \delta(\mathbf{x} - \mathbf{x}') \ . \tag{7}$$

The vortices advect with the flow according to:

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{u}(\mathbf{x}_i, t) \quad i = 1, \ldots, N_V \tag{8}$$

Vortices enter the domain at each time step Δt with an initial circulation Γ_0. The individual vortices roll-up to form large-scale vortical structures. For further details on the method, see appendix B.

3 Parameter inference with Bayesian neural network ensembles

3.1 The training library of simulated flames

LSGEN2D [29] is a level-set solver that iterates the G field of the G-equation model for a known set of parameters $\alpha, r_{D,0}, z_{0,\text{i}}, \text{St}, \beta$ and f/f_s. The velocity field in LSGEN2D is calculated using the DVM (Eq. 5). In this study, the G field is iterated until convergence to a set of 200 different periodic solutions. This is repeated for 10^4 combinations of parameters sampled from the ranges shown in Table 1. The forced cycle states are processed to find a $y = f(x)$ discretisation of the $G = 0$ contour, for all x in the range of the experiment observation window. This is done by interpolating the G field values for every vertical coordinate, and recording the positions y in vectors $\mathbf{y}$. To create a single observation vector $\mathbf{z}$ representing a sequence of 10 flame front positions, 10 position vectors are appended together. This is repeated for every state in the forced cycle. The result is a library of 2×10^6 observation - parameter pairs.

3.2 Architecture and training of the BayNNE

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

$$p_\theta(\mathbf{p}|\mathbf{z}) = \mathcal{N}\left(\boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\Sigma}(\mathbf{z})\right) , \tag{9}$$

Table 1: Parameters of the discrete vortex method used in the G-equation model that are varied in this work and the range over which they are varied in order to generate the simulated flame front library.

Parameter ranges	Description
$0.05 \leq \alpha \leq 0.5$	Amplitude of base flow forcing
$0.01 \leq r_{D,0} \leq 0.2$	Vortex radius
$0.6 \leq z_{0,\mathrm{i}} \leq 0.7$	Initial velocity coordinate
$4 \leq \beta \leq 8$	Aspect ratio of the unperturbed flame
0.03 m $\leq D \leq 0.04$ m	Bluff-body side length
10 m/s $\leq U \leq 15$ m/s	Mean base flow speed
100 Hz $\leq f \leq 150$ Hz	Frequency of base flow forcing

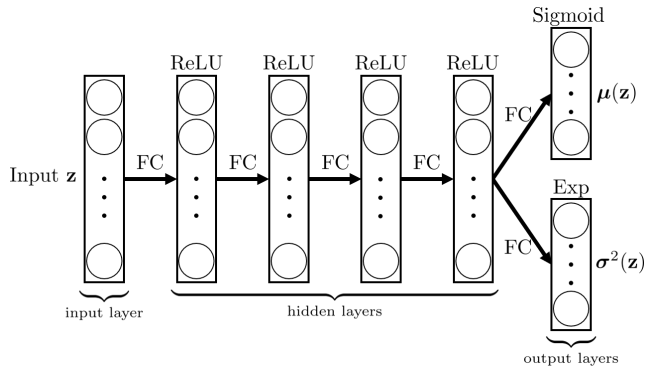

Figure 2: Architecture of each neural network in the ensemble of 20. The input and hidden layers have 900 nodes each, while each output layer has six nodes each. All layers are fully connected (FC). Rectified Linear Unit (ReLU) activation functions are used for the hidden layers and sigmoid and exponential (Exp) activation functions are used for the mean and variance output layers respectively.

where $\boldsymbol{\Sigma}(\mathbf{z})$, the posterior covariance matrix of the parameters given the data, is diagonal with $\boldsymbol{\sigma}^2(\mathbf{z})$ on its diagonal. This enforces our assumption that the parameters are mutually independent, given the observations $\mathbf{z}$. We use an ensemble of $M = 20$ neural networks. The architecture of each neural network is shown in figure 2. Each neural network comprises an input layer, four hidden layers with ReLU activations and two output layers: one for the mean vector $\boldsymbol{\mu}(\mathbf{z})$ and one for the variance vector $\boldsymbol{\sigma}^2(\mathbf{z})$. The output layer for the mean uses a sigmoid activation to restrict outputs to the range $(-1, 1)$. The output layer for the variance uses an exponential activation to ensure positivity. Each neural network in the ensemble is initialised with unique weights $\boldsymbol{\theta}_{j,anc}$ sampled from a Gaussian prior distribution $\mathcal{N}(0, \frac{1}{N_H})$ and biases $\mathbf{b}_{j,anc}$ sampled from a uniform prior distribution in the range $[-\frac{1}{\sqrt{N_H}}, \frac{1}{\sqrt{N_H}}]$ according to [30].

For a single observation $\mathbf{z}$, the j-th neural network in the ensemble produces a mean and variance estimate of the G-equation parameters:

$$\boldsymbol{\mu}_j(\mathbf{z}), \boldsymbol{\sigma}_j^2(\mathbf{z}). \tag{10}$$

This is achieved by minimising the loss function $\mathfrak{L}_j$:

$$\begin{aligned}\mathfrak{L}_j = \left(\boldsymbol{\mu}_j(\mathbf{z}) - \mathbf{p}\right)^T \boldsymbol{\Sigma}_j\left(\mathbf{z}\right)^{-1} \left(\boldsymbol{\mu}_j(\mathbf{z}) - \mathbf{p}\right) + \log\left(\left|\boldsymbol{\Sigma}_j\left(\mathbf{z}\right)\right|\right) \\ + \left(\boldsymbol{\theta}_j - \boldsymbol{\theta}_{anc,j}\right)^T \boldsymbol{\Sigma}_{prior}^{-1} \left(\boldsymbol{\theta}_j - \boldsymbol{\theta}_{anc,j}\right).\end{aligned} \tag{11}$$

The first two terms of the loss function are the negative logarithm of the normalised Gaussian likelihood function up to an additive constant. The third term is a regularising term that penalises deviation from prior anchor values $\boldsymbol{\theta}_{anc,j}$. The NNs produce samples from the posterior distribution. This is called randomised maximum a-posteriori (MAP) sampling [15].

For a single observation vector $\mathbf{z}$, the prediction from the ensemble of neural networks is therefore a distribution of M Gaussians, each centred at their respective means $\boldsymbol{\mu}_j(\mathbf{z})$. Following similar

treatment in [31], this distribution is then approximated by a single multivariate Gaussian posterior distribution $p(\mathbf{p}|\mathbf{z}) \approx \mathcal{N}(\boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\Sigma}(\mathbf{z}))$ with mean and variance:

$$\boldsymbol{\mu}(\mathbf{z}) = \frac{\Sigma_j \boldsymbol{\mu}_j(\mathbf{z})}{M}, \;\; \boldsymbol{\Sigma}(\mathbf{z}) = \text{diag}\left(\boldsymbol{\sigma}^2(\mathbf{z})\right), \tag{12}$$

$$\boldsymbol{\sigma}^2(\mathbf{z}) = \frac{\Sigma_j \boldsymbol{\sigma}_j^2(\mathbf{z})}{M} + \frac{\Sigma_j \boldsymbol{\mu}_j^2(\mathbf{z})}{M} - \left(\frac{\Sigma_j \boldsymbol{\mu}_j(\mathbf{z})}{M}\right)^2. \tag{13}$$

This is repeated for every observation vector $\mathbf{z}$. The posterior distribution $p(\mathbf{p}|\mathbf{z}_i)$ with the smallest total variance $\sigma^2_{i,\text{tot}} = ||\boldsymbol{\sigma}^2(\mathbf{z}_i)||_1$ is chosen as the best guess to the true posterior. The M parameter samples from the chosen posterior are used for re-simulation, which allows us to check the predicted flame shapes and to calculate the normalised area variation over one cycle.

4 Results

Figure 3 shows means and uncertainties of the posterior of parameters (α, $r_{D,0}$, z_0, St, β) given the input experimental data. The BayNNE is able to infer the parameters reliably through abrupt changes in the input data. We see that the uncertainties in $r_{D,0}$, the vortex core radius, are high: the two standard deviation range covers nearly the whole range from which the parameter was sampled (see table 1). This suggests that this parameter is difficult to recover: it only has a weak effect on the flame front over ten consecutive images. Between $t \approx 100$ and $t \approx 340$ we observe approximately periodic behaviour in the plots of the inferred values of α and z_0 . The period of this behaviour is 80 time steps, which is equal to the period of the oscillations observed in the experiment. This implies that there is some model error present: the BayNNE predicts different flame shapes for different moments in the period because no single shape describes the whole period well. However, unlike the core radius parameter $r_{D,0}$, these two parameters are more easily recovered, implying that their effect on the flame front over ten consecutive images is strong.

Using the parameter predictions at the five time steps, the flame can be re-simulated using LSGEN2D and the DVM downstream of the observation window, predicting the flame shape in its entirety as shown in figure 4. This means that the $h(\mathbf{x})$ and $\tau(\mathbf{x})$ fields of the distributed $n - \tau$ heat release model can be calculated, which was not possible from the experimental images alone. For details on how to calculate h and τ from the re-simulated flame shapes, see Supplementary materials E. Figures8 to 12 show the re-simulated flames as well as the h and τ fields. Additionally, the flame shapes are re-simulated at $1/4, 1/2$ and $3/4$ the amplitude inferred by the BayNNE (all other inferred parameters are kept constant). The h and τ fields calculated from these flame shapes are shown in figures 13 and 14. These are used in a Helmholtz solver [32; 33] to determine the eigenfrequencies and growth rates of the thermo-acoustic system.

Table 2: Parameter values and uncertainties inferred by the BayNNEs at five different time steps.

Parameter	t = 32	t = 60	t = 304	t = 376	t = 416
α	0.433 ± 0.105	0.303 ± 0.142	0.435 ± 0.092	0.452 ± 0.089	0.305 ± 0.137
$r_{D,0}$	0.099 ± 0.049	0.096 ± 0.061	0.100 ± 0.069	0.092 ± 0.052	0.078 ± 0.035
$z_{0,i}$	0.615 ± 0.022	0.614 ± 0.017	0.612 ± 0.021	0.613 ± 0.021	0.635 ± 0.028
St	14.30 ± 3.458	16.19 ± 4.413	15.78 ± 5.411	15.29 ± 5.860	13.17 ± 4.156
β	7.974 ± 0.050	7.931 ± 0.156	7.918 ± 0.221	7.913 ± 0.210	7.957 ± 0.109

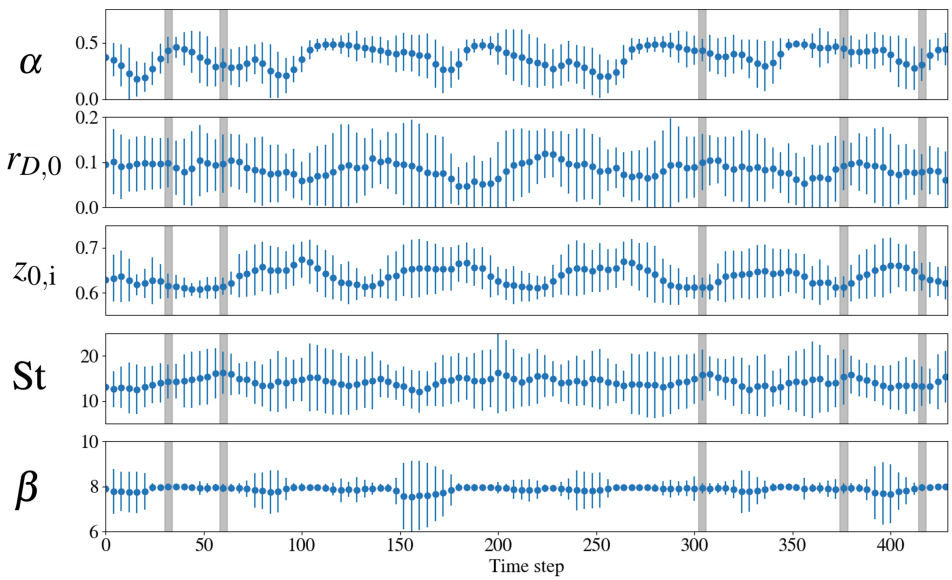

Figure 3: Graphs of the learned posterior of the parameters of the velocity model, represented as means and $\pm$ 2 standard deviations. The inferred parameter values at five different times (grey columns, corresponding to $t = 32$, 60, 304, 376 and 416) are listed in Table 2.

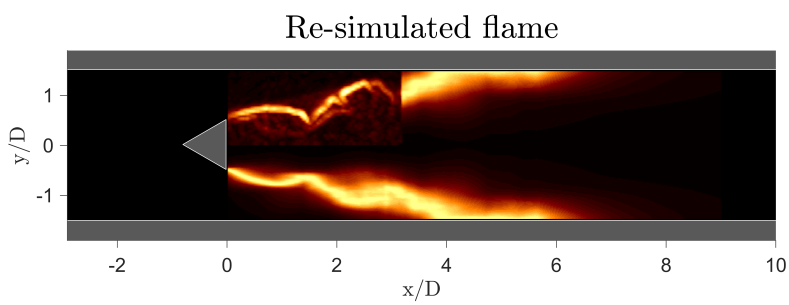

Figure 4: The re-simulated flame shape at $t = 416$. The top left quadrant from 0 to 3.4D is the experimentally-determined flame front. The other three quadrants show the flame shape predicted by re-simulating the flame using LSGEN2D and DVM with their inferred posterior parameter values.

4.1 Helmholtz solver results

In figure 5, the red square plot-markers are the fundamental duct acoustic mode in the absence of a flame, and the grey plot-markers are the thermoacoustic eigenmodes with flame presence. These are shown at the four amplitudes and five time steps. The grey symbols deviate from the red square, showing that the thermoacoustic effect is active. At all time steps, the thermoacoustic eigenmodes correspond to negative growth rates, meaning the system is thermoacoustically stable. Furthermore, the modes for cases $t = 32,\ 60,\ 304$ and 376 are clustered around a single point, meaning they are weakly dependent on velocity forcing amplitude. This shows that the flame transfer function (FTF) is derived with considerable accuracy, since it nearly evaluates to the same mode regardless of the amplitude used to derive it. Thus, with this model, an FTF derived at high amplitude, when it is observable, is also valid at low amplitude, when it is not observable. In the $t = 416$ case, the scatter in the values of the growth rate and frequency for amplitude $\epsilon/4$ and $\epsilon/2$ cases could be due to the larger jumps seen in all regions of the corresponding $\tau(\mathbf{x})$ fields (figure 13).

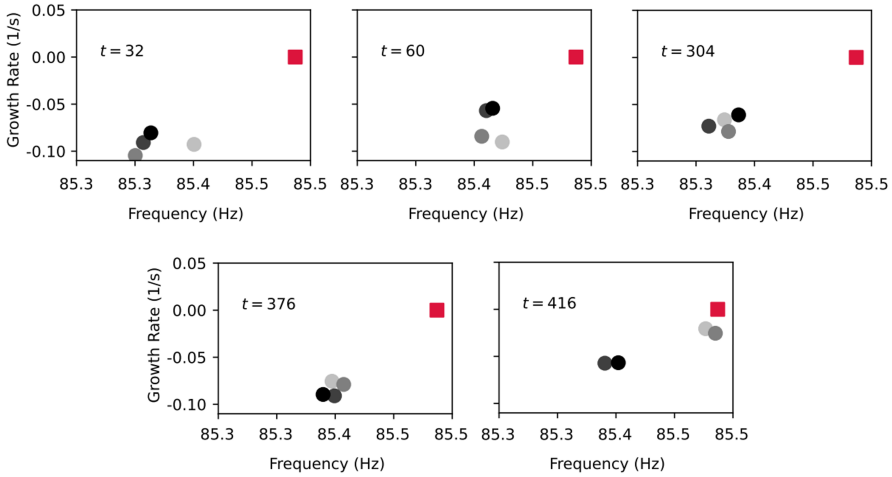

Figure 5: The eigenvalues for the first mode at five time steps and four amplitudes. The red squares in each subplot show the eigenfrequecies of the system when no flame is present. Each frame corresponds to a different time step, while the greyscale within each frame corresponds to a different velocity forcing amplitude, ϵ. The lightest grey corresponds to $\epsilon/4$ and the darkest to ϵ.

5 Conclusions

This work uses a Bayesian machine learning method to infer the parameters of a physics-based model of a bluff-body stabilised flame. With this method, parameters can be inferred with known uncertainty in milliseconds, which is fast enough to be used in real-time. For the velocity field, a discrete vortex method (DVM) is used that models large-scale vortical perturbations on the flame front which play a significant role in thermoacoustic instability. Heteroscedastic Bayesian neural network ensembles (BayNNE), trained on a library of simulated flame fronts, learn to infer five parameters and their uncertainties from ten consecutive flame front snapshots. Both the flame and velocity field models are physics-based and so they can extrapolate in physical space (i.e. beyond the experimental observation window) and also in parameter space (to different perturbation amplitudes) to obtain distributed flame transfer functions (FTFs). The FTFs are then successfully used in a thermoacoustic model to predict frequencies and growth rates. Results show that the BayNNE is able to reliably infer the parameters through changes in the input data. The confidence intervals of each of the inferred parameters of the DVM are different, showing which ones have a more significant impact on the flame response. The predicted thermoacoustic frequencies and growth rates do not show a strong dependency on the amplitude of forcing. This suggests that, with this model, an FTF derived at high amplitude, when it is observable, is also valid at low amplitude, when it is not observable.

Broader impact

This work will lead to cost savings in high-energy density combustor design by enabling engineers to quickly and reliably tune their models to match experiments.

Declaration of funding

This project has received funding from the UK Engineering and Physical Sciences Research Council (EPSRC) award EP/N509620/1 and from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement number 766264.

References

[1] Matthew P. Juniper and R. I. Sujith. Sensitivity and Nonlinearity of Thermoacoustic Oscillations. *Annual Review of Fluid Mechanics*, 50:661–689, 2018. ISSN 00664189. doi: 10.1146/annurev-fluid-122316-045125.

[2] Ying Huang and Vigor Yang. Dynamics and stability of lean-premixed swirl-stabilized combustion. *Progress in Energy and Combustion Science*, 35(4):293–364, 2009. ISSN 0360-1285. doi: 10.1016/j.pecs.2009.01.002. URL `http://dx.doi.org/10.1016/j.pecs.2009.01.002`.

[3] Jakob J. Keller. Thermoacoustic Oscillations in Combustion Chambers of Gas Turbines disturbance of a quantity if placed in front of the symbol for the corresponding quantity A g* = fluctuation in rate of heat addition that would occur in the absence of fluctuations in eff. *AIAA Journal*, 33(12), 1995.

[4] John W. S. B. Rayleigh. The explanation of certain acoustical phenomena. *Nature*, 18(455): 319–321, 1878.

[5] D A Smith and E E Zukoski. Combustion Instability Sustained by Unsteady Vortex Combustion. In *AIAA 21st Joint Propulsion Conference*, pages 1–9, Monterey, California, 1985.

[6] Thierry J. Poinsot, Arnaud C. Trouve, Denis P. Veynante, Sebastien M. Candel, and Emile J. Esposito. Vortex-driven acoustically coupled combustion instabilities. *Journal of Fluid Mechanics*, 177:265–292, 1987. ISSN 14697645. doi: 10.1017/S0022112087000958.

[7] Andrew W. Caswell, Brent A. Rankin, Bethany C. Huelskamp, Amy C. Lynch, Vincent Belovich, and James R. Gord. Spatiotemporal characterization of bluff-body stabilized turbulent premixed flames using simultaneous high-repetition-rate OH PLIF and PIV measurements. *53rd AIAA Aerospace Sciences Meeting*, pages 1–12, 2015. doi: 10.2514/6.2015-0424.

[8] C. A. Fugger, Tongxun Yi, Joshua P. Sykes, Andrew W. Caswell, Brent A. Rankin, Joseph D. Miller, and James R. Gord. The structure and dynamics of a bluff-body stabilized premixed reacting flow. *AIAA Aerospace Sciences Meeting, 2018*, 1(210059), 2018. doi: 10.2514/6.2018-1190.

[9] Forman A. Williams. 3. Turbulent Combustion. *Frontiers in Applied Mathematics*, pages 97–131, jan 1985. doi: 10.1137/1.9781611971064.CH3. URL `https://epubs.siam.org/doi/abs/10.1137/1.9781611971064.ch3`.

[10] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143–10162, may 1994. ISSN 01480227. doi: 10.1029/94jc00572. URL `https://onlinelibrary.wiley.com/doi/full/10.1029/94JC00572https://onlinelibrary.wiley.com/doi/abs/10.1029/94JC00572https://agupubs.onlinelibrary.wiley.com/doi/10.1029/94JC00572`.

[11] Hans Yu, Thomas Jaravel, Matthias Ihme, Matthew P. Juniper, and Luca Magri. Data Assimilation and Optimal Calibration in Nonlinear Models of Flame Dynamics. *Journal of Engineering for Gas Turbines and Power*, 141(12):1–10, 2019. ISSN 15288919. doi: 10.1115/1.4044378.

[12] Hans Yu. *Inverse problems in thermoacoustics*. PhD thesis, University of Cambridge, 2020.

[13] David J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

[14] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning Zoubin Ghahramani. In *Proceedings of the 33rd International Conference on Machine Learning*, New York, NY, 2016. W&CP.

[15] Tim Pearce, Felix Leibfried, Alexandra Brintrup, Mohamed Zaki, and Andy Neely. Uncertainty in Neural Networks: Approximately Bayesian Ensembling. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 108, 2020. URL `http://arxiv.org/abs/1810.05546`.

[16] Ushnish Sengupta, Matt Amos, J. Scott Hosking, Carl Edward Rasmussen, Matthew P. Juniper, and Paul J. Young. Ensembling geophysical models with Bayesian neural networks. *Advances in Neural Information Processing Systems*, 2020-Decem(December), 2020. ISSN 10495258.

[17] Ushnish Sengupta, Maximilian L. Croci, and Matthew P. Juniper. Real-time parameter inference in reduced-order flame models with heteroscedastic Bayesian neural network ensembles. In *NeurIPS*, 2020. URL `http://arxiv.org/abs/2011.02838`.

[18] Maximilian L. Croci, Ushnish Sengupta, and Matthew P. Juniper. Online parameter inference for the simulation of a Bunsen flame using heteroscedastic Bayesian neural network ensembles. In *Deep Learning for Simulation workshop at the 9th International Conference Conference on Learning Representations, ICLR 2021*, Virtual, 2021.

[19] Maximilian L. Croci, Ushnish Sengupta, and Matthew P. Juniper. Data Assimilation Using Heteroscedastic Bayesian Neural Network Ensembles for Reduced-Order Flame Models. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12746 LNCS:408–419, 2021. ISSN 16113349. doi: 10.1007/978-3-030-77977-1_33.

[20] Maximilian L. Croci, Ushnish Sengupta, and Matthew P. Juniper. Bayesian Inference in Physics-Based Nonlinear Flame Models. In *Deep Learning and Inverse Problems workshop at the 35th Conference on Neural Information Processing Systems, NeurIPS 2021*, Virtual, 2021.

[21] A. Leonard. Vortex methods for flow simulation. *Journal of Computational Physics*, 37(3): 289–335, oct 1980. ISSN 0021-9991. doi: 10.1016/0021-9991(80)90040-6.

[22] G.-H Cottet and Petros Koumoutsakos. *Vortex methods - theory and practice.* Cambridge University Press, 03 2000. ISBN 978-0-521-62186-1. doi: 10.1017/CBO9780511526442.

[23] Satyanarayanan R Chakravarthy, Obla J Shreenivasan, Benjamin Boehm, Andreas Dreizler, and Johannes Janicka. Experimental characterization of onset of acoustic instability in a nonpremixed half-dump combustor. *Journal of the Acoustical Society of America*, 2007. doi: 10.1121/1.2741374.

[24] H. Murat Altay, Raymond L. Speth, Duane E. Hudgins, and Ahmed F. Ghoniem. Flame-vortex interaction driven combustion dynamics in a backward-facing step combustor. *Combustion and Flame*, 156(5):1111–1125, 2009. ISSN 00102180. doi: 10.1016/j.combustflame.2009.02.003. URL `http://dx.doi.org/10.1016/j.combustflame.2009.02.003`.

[25] Seunghyuck Hong, Raymond L. Speth, Santosh J. Shanbhogue, and Ahmed F. Ghoniem. Examining flow-flame interaction and the characteristic stretch rate in vortex-driven combustion dynamics using PIV and numerical simulation. *Combustion and Flame*, 160(8):1381–1397, 2013. ISSN 00102180. doi: 10.1016/j.combustflame.2013.02.016. URL `http://dx.doi.org/10.1016/j.combustflame.2013.02.016`.

[26] Santosh Hemchandra, Santosh Shanbhogue, Seunghyuck Hong, and Ahmed F. Ghoniem. Role of hydrodynamic shear layer stability in driving combustion instability in a premixed propane-air backward-facing step combustor. *Physical Review Fluids*, 3(6):1–20, 2018. ISSN 2469990X. doi: 10.1103/PhysRevFluids.3.063201.

[27] Joel V. Vasanth and Satyanarayanan R. Chakravarthy. A reduced-order model for lock-on via vortex-combustion-acoustic closed-loop coupling in a step combustor. *Combustion Science and Technology*, pages 1–23, 2021. doi: 10.1080/00102202.2021.1909578. URL `https://doi.org/10.1080/00102202.2021.1909578`.

[28] Tobin A. Driscoll and Lloyd N. Trefethen. *Schwartz-Christoffel Mapping*. Cambridge University Press, 2002. doi: 10.1017/CBO9780511546808.

[29] Santosh Hemchandra. *Dynamics of Turbulent Premixed Flames in Acoustic Fields*. PhD thesis, University of Cambridge, 2009.

[30] Siyu He, Yin Li, Yu Feng, Shirley Ho, Siamak Ravanbakhsh, Wei Chen, and Barnabás Póczos. Learning to predict the cosmological structure formation. *Proceedings of the National Academy of Sciences of the United States of America*, 116(28):13825–13832, 2019. ISSN 10916490. doi: 10.1073/PNAS.1821458116.

[31] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *Proceedings of the 31st Conference on Neural Information Processing Systems (NeurIPS)*, 2017.

[32] Stefano Falco and Matthew P. Juniper. Shape optimization of thermoacoustic systems using a 2D adjoint Helmholtz solver. In *Proceedings of the ASME 2020 (Turbomachinery Technical Conference & Exposition)*, 2020.

[33] Stefano Falco. *Shape optimization for thermoacoustic instability with an adjoint Helmholtz solver*. PhD thesis, University of Cambridge, aug 2021. URL `https://www.repository.cam.ac.uk/handle/1810/336978`.

[34] Tobin A Driscoll and Lloyd N Trefethen. *Schwarz-Christoffel Mapping*. Cambridge University Press, 2002.

[35] L. M. Milne-Thomson. *Theoretical Hydrodynamics*. Dover Publications, New York, 5 edition, 1968.

[36] Mark Pastoor, Lars Henning, Bernd R. Noack, Rudibert King, and Gilead Tadmor. Feedback shear layer control for bluff body drag reduction. *Journal of Fluid Mechanics*, 608:161–196, 2008. ISSN 00221120. doi: 10.1017/S0022112008002073.

[37] Mirta Perlman. On the accuracy of vortex methods. *Journal of Computational Physics*, 59(2): 200–223, 1985. ISSN 10902716. doi: 10.1016/0021-9991(85)90142-1.

[38] R. R. Clements. An inviscid model of two-dimensional vortex shedding. *Journal of Fluid Mechanics*, 57(2):321–336, 1973. ISSN 14697645. doi: 10.1017/S0022112073001187.

[39] Ahmed F. Ghoniem and Yves Cagnon. Vortex simulation of laminar recirculating flow. *Journal of Computational Physics*, 68:346–377, 1987. ISSN 10902716. doi: 10.1016/0021-9991(87)90062-3.

[40] Miguel Arevallilo Herráez, David R. Burton, Michael J. Lalor, and Munther A. Gdeisat. Fast two-dimensional phase-unwrapping algorithm based on sorting by reliability following a noncontinuous path. *Applied optics*, 41(35):7437, dec 2002. ISSN 1559-128X. doi: 10.1364/AO.41.007437. URL `https://pubmed.ncbi.nlm.nih.gov/12502301/`.

[41] Matthew P. Juniper. Sensitivity analysis of thermoacoustic instability with adjoint Helmholtz solvers. *Physical Review Fluids*, 3:110509, 2018. doi: 10.1103/PhysRevFluids.3.110509.

A Supplementary material: Non-dimensionalisation

The characteristic scales of the problem are the side length D of the bluff body, the length of the unstretched unforced flame, L, the spatial frequency U/L where U is mean base flow speed and the excitation angular frequency $2\pi f$. The Strouhal number is the ratio of the spatial time scale to the excitation time scale (where the time scales are the reciprocal of the frequencies) $\mathrm{St} = 2\pi f L/U$. By defining the aspect ratio of the unstretched, unforced flame $\beta = L/D$, the Strouhal number may be written equivalently as $\mathrm{St} = 2\pi f \beta D/U$.

B Supplementary material: The discrete vortex method

We begin with the incompressible 2D Euler equations in vorticity form:

$$\frac{\partial \boldsymbol{\omega}}{\partial t} + \mathbf{u} \cdot \nabla \boldsymbol{\omega} = \boldsymbol{\omega} \cdot \nabla \mathbf{u}, \quad \nabla \cdot \mathbf{u} = 0, \tag{14}$$

where the term on the right hand side is the stretching of vorticity due to velocity gradients and the terms on the left form the material derivative $\mathrm{D}\boldsymbol{\omega}/\mathrm{D}t$, or the Lagrangian transport term. The vorticity in the domain is discretised into a number of vortex elements at locations $\mathbf{x}_i$, each carrying an elementary amount of vorticity within a core of radius δ. The distribution of vorticity within the vortex element is given by a core function f_δ, which ensures a finite velocity at the core centre. The discrete vorticity field $\boldsymbol{\omega}^h$ is then

$$\boldsymbol{\omega}^h = \sum_{i=1}^{N_V} \Gamma_i f_\delta(r), \tag{15}$$

where Γ_i are the vortex circulations and the summation is over the total number of vortices in the domain. The advection of the vortices are given by the solution of the N_V ODEs

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{u}(\mathbf{x}_i, t). \tag{16}$$

The local flow velocities $\mathbf{u}(\mathbf{x}_i, t)$ are obtained from the solution of the Poisson equation $\nabla^2 \mathbf{u} = -\nabla \times \boldsymbol{\omega}$:

$$\mathbf{u} = \nabla G(\mathbf{x}, \mathbf{x}') \star \boldsymbol{\omega}^h + \mathbf{u}_p, \tag{17}$$

where $G(\mathbf{x}, \mathbf{x}')$ is the Green's function for the Laplacian operator in 2D, i.e., $\nabla^2 G(\mathbf{x}, \mathbf{x}') = \delta(\mathbf{x} - \mathbf{x}')$; and $\star$ is the convolution operator. In the absence of any boundaries, $G(\mathbf{x} - \mathbf{x}') = \ln(|\mathbf{x} - \mathbf{x}'|)/2\pi$. The first term is the rotational component of the flow and the velocity $\mathbf{u}_p$ is the irrotational part of the flow field, which is uniquely chosen such that the wall-normal boundary condition $\mathbf{u} \cdot \hat{\mathbf{n}} = 0$ is satisfied.

The solution of the potential flow problem can be found via a Schwartz-Christoffel transformation (SCT). Here, the physical domain is modelled in a complex plane $z = x + yj$ (figure 6(a)). Symmetry about the x-axis is assumed, since the experimental flame shapes are similarly symmetric. The z-plane is mapped onto the upper-half of a transformed s-plane (figure 6(c)). If $s_{V,\,i} \quad i = 1, \ldots, N_V$ represents the vortex locations, an image system of vortices at their conjugates $s^*_{V,\,i}$ is then defined in the lower-half of the s-plane, such that the resulting wall-normal velocity cancels to 0. The mapping is performed in two steps via an intermediate mapping to an infinite-strip in the ξ-plane (figure 6(b)), which simplifies the mathematical expression for the SCT [34]. The mapping from z to ξ is obtained by integrating the expression

$$\frac{dz}{d\xi} = 1.5 \left[\sinh \frac{\pi}{2}\xi\right]^{-1/4} \left[\sinh \frac{\pi}{2}(\xi - \xi_E)\right]^{3/4} \left[\sinh \frac{\pi}{4}(\xi - \xi_F)\right]^{-1/2}, \tag{18}$$

and the mapping from ξ to s is simply $s = e^{\pi\xi}$. The values ξ_E, ξ_F and ξ_B in figure 6(b) are obtained numerically using the SCT toolbox developed by [34]. These are 0.95, 0.63 and $0.58 + j$ respectively.

The potential flow velocity field is then given by a spatially-uniform inlet flow $U(t)$ in the ξ-plane and induced velocities from vortex elements and their images in s-plane. Both are transformed to the

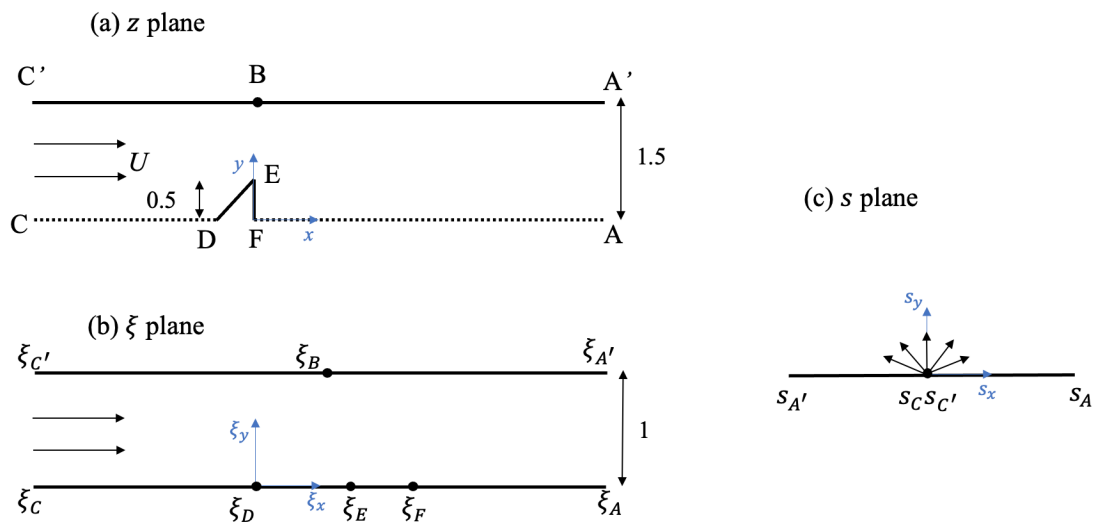

Figure 6: Illustration of the Schwartz-Christoffel mappings from (a) the physical z plane to (b) the transformed ξ plane, which is an infinite strip, and then to (c) the upper half of the s plane. The dotted line in (a) is a line of symmetry. The subscripts in the labels in the ξ and s planes correspond to the respective vertices in the z-plane. Marked dimensions in the z plane in (a) are scaled by the size D of the bluff body.

physical z-plane using the chain rule.The chain rule is applied to transform these velocities into the physical domain. The inlet flow is given as

$$\sigma_{inlet} = U(t)\frac{d\xi}{dz}, \tag{19}$$

where σ represents complex velocities. The flow due to the vortices and their images as [35]

$$\sigma_{vort} = \sum_{i=1}^{N_V} j\frac{\Gamma_i}{2\pi}\left[-\frac{1}{s - s_{V,\,i}} + \frac{1}{s - s^*_{V,\,i}}\right]\frac{ds}{d\xi}\frac{d\xi}{dz}. \tag{20}$$

In Eq. 20, the sum is over all the vortex elements each with circulation Γ_i and location $s_{V,\,i}$. The x and y components of the physical velocity field $\mathbf{u}$ are then $\mathrm{Re}(\sigma_{inlet}+\sigma_{vort})$ and $-\mathrm{Im}(\sigma_{inlet}+\sigma_{vort})$ respectively. The vortices advect with the flow as

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{u}(\mathbf{x}_i, t) \quad i = 1, \ldots, N_V \tag{21}$$

A generalised Kutta condition is used to model the vortex generation process. According to this, the vorticity is discretised by introducing a single vortex element into the domain at each time step Δt with a circulation $\Gamma_0 = -u_0(t)^2\Delta t/2$. The velocity u_0 is computed at a location $z_0 = z_{0,\mathrm{r}} + z_{0,\mathrm{i}}j$ directly above the bluff body edge E as in figure 6(a) ($z_{0,\mathrm{r}} = 0$). The vortex elements are then advected as per Eq. 21.

The above gives a point vortex description of the potential flow field. The singularity at the center of the vortex elements is removed using a refined Lamb-Oseen model. According to this model, a vortex 'core' of size r_D is defined, along with an azimuthal velocity u_θ that varies smoothly with distance r from the center as

$$u_\theta = \frac{\Gamma_0}{2\pi r}(1 - e^{-r^2/r_D^2}) \tag{22}$$

In the model in Eq. 22 the velocity field tends to that of a rotational vortex as $r \to 0$, and it tends to that of the potential flow solution as $r \to \infty$. We implement a linear temporal growth of the core radius as $r_D(t) = r_{D,0} + 2\sqrt{\nu t}$, where ν is the kinematic viscosity and $r_{D,0}$ is the initial core radius

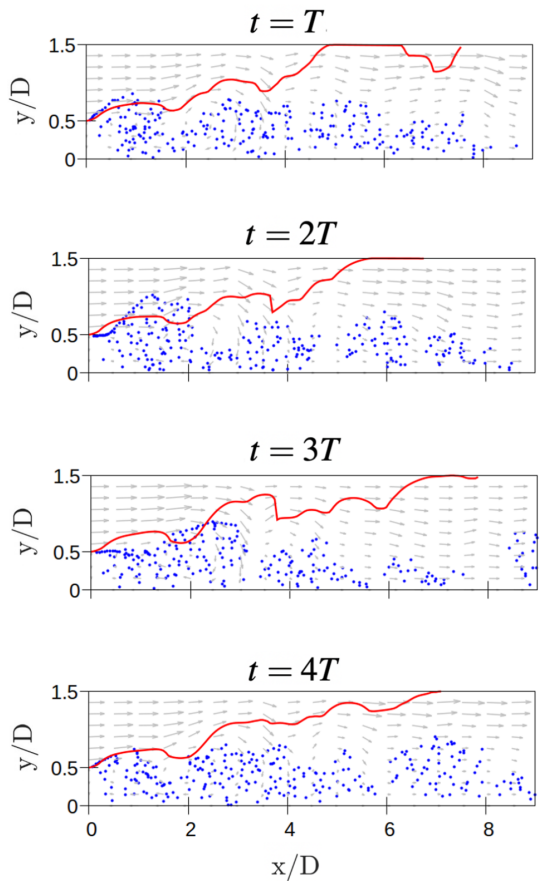

Figure 7: Flame fronts (red) one period apart simulated using the G-equation and the discrete vortex method. The blue dots are the centres of the vortices, advected by the velocity field shown in grey.

of the vortex elements at the instant of being introduced into the domain. The linear temporal growth model thus accounts for the viscous diffusion of vorticity [36].

To incorporate the effect of the boundary layer on the vorticity, we model the decay of vorticity at the walls as $d\Gamma/dt = -\Gamma$, when the distance between the vortex center of any vortex element and the wall is less than r_D. All lengths in the DVM are non-dimensionalised with the bluff body side length D, velocities with the mean inlet $\bar{U}$ and time with $\bar{U}/D$.

The acoustic forcing on the flow field is modelled with a harmonically oscillating inlet flow velocity $U(t)/\bar{U} = 1 + \alpha \sin(\mathrm{St}\ t)$ in Eq. 19, where $\mathrm{St} = 2\pi f \beta D/\bar{U}$ is the forcing Strouhal number expressed in terms of the forcing frequency f and the nominal aspect ratio β. This adds a time-dependent potential to the existing potential flow, while still fulfilling the boundary conditions. The oscillatory flow $u_0(t)$ at z_0 produces vortices of varying Γ_0 and so the strengths of the vortex elements are also affected by the forcing. The effect of the forcing is thereby incorporated into the vortex shedding process. In order to ensure that N_V does not increase without bound, vortices are deleted when they cross a streamwise location $x_{\max} = 25$.

The time integration in Eq. 21 is solved using a 4th-order Runge-Kutta method, and the numerical integration of the SCT in Eq. 18 is performed using Simpson's rule.

The accuracy of the DVM in predicting the velocity field depends especially on its discretisation parameters - the core size, the core function and the number of vortex elements shed per time step [37]. In many studies [38; 39; 36], these parameters are hand-tuned and pre-set before the model is used for predictions. In order to render the flame-vortex model more quantitatively accurate in its predictions, the parameters $r_{D,0}$, $z_{0,\mathrm{i}}$, St, α and β are inferred from a machine learning model in this work.

The evolution of the flame front along with the velocity field from the DVM is shown in figure 7. The generation of wrinkles on the flame front due to the vortical perturbations is evident, and is one of the major mechanisms of flame area perturbations, to cause heat release rate fluctuations q', that drive thermoacoustic instability.

C Hyperparameters

Table 3: Hyperparameter settings used for neural network training.

Hyperparameter	Value
Training	
Train-test split	80:20
Batch size	256
Epochs	100
Optimiser	Adam
Learning rate	10^{-4}
Architecture	
Input units	600
Hidden layers	4
Units per hidden layer	600
Output layers	2
Units per output layer	5
Ensemble size	20

D Supplementary material: Results

D.1 Re-simulated flame results

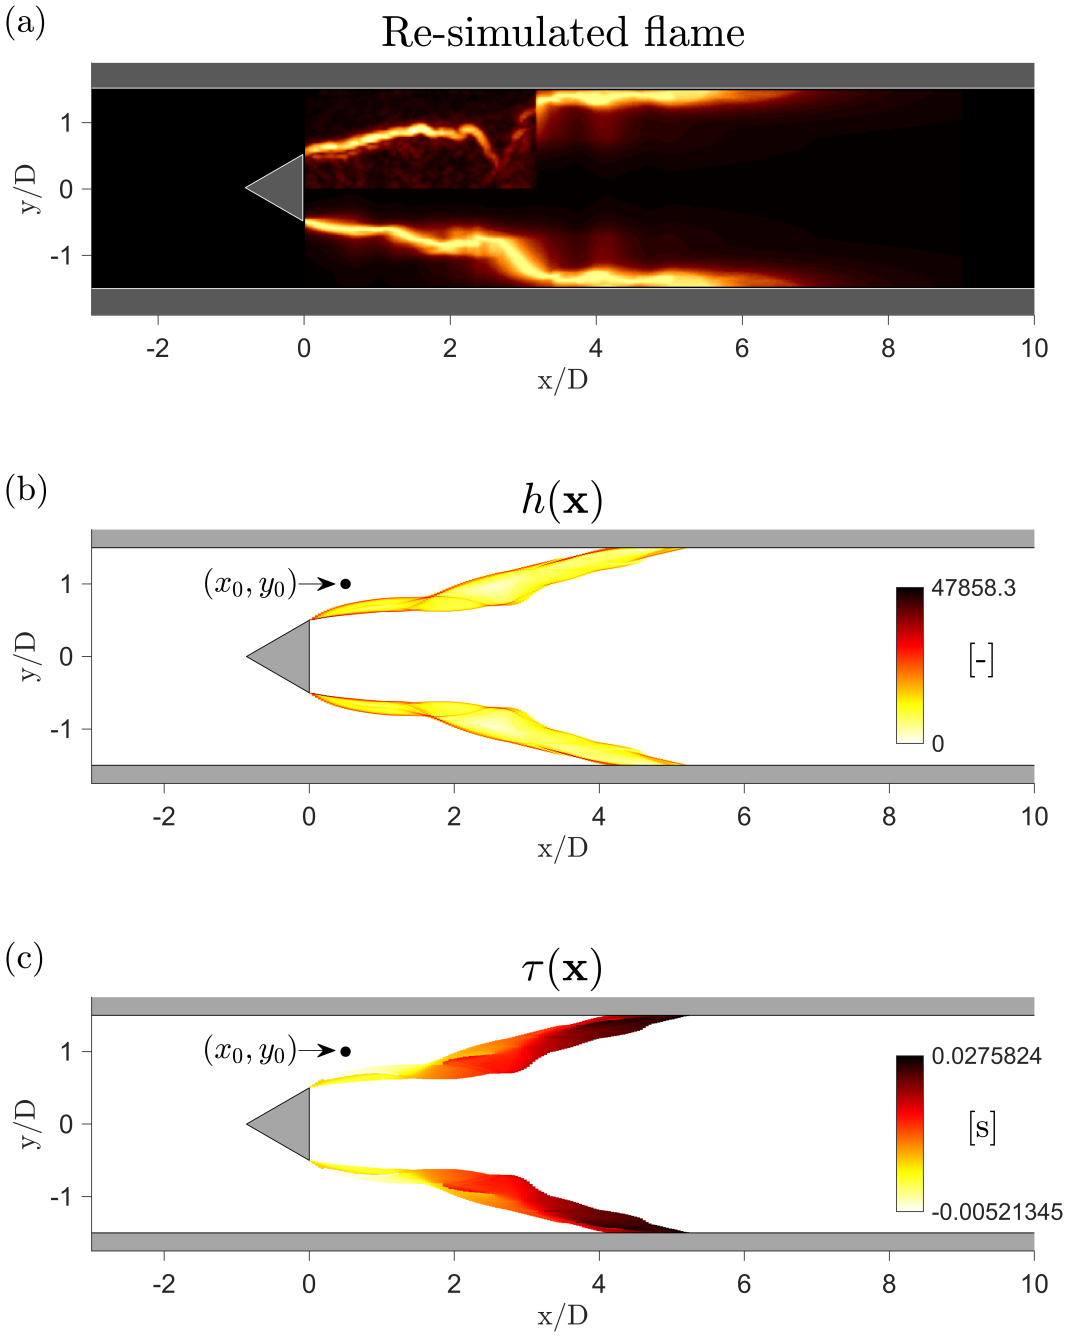

Figure 8: (a): Re-simulated flame shape at $t = 32$. (b) and (c): n and τ fields calculated from the re-simulated flame shape at $t = 32$. The centre (x_0, y_0) of the measurement region $w(x, y)$ is also shown. Part of the flame is upstream of the measurement point, which leads to a negative time delay.

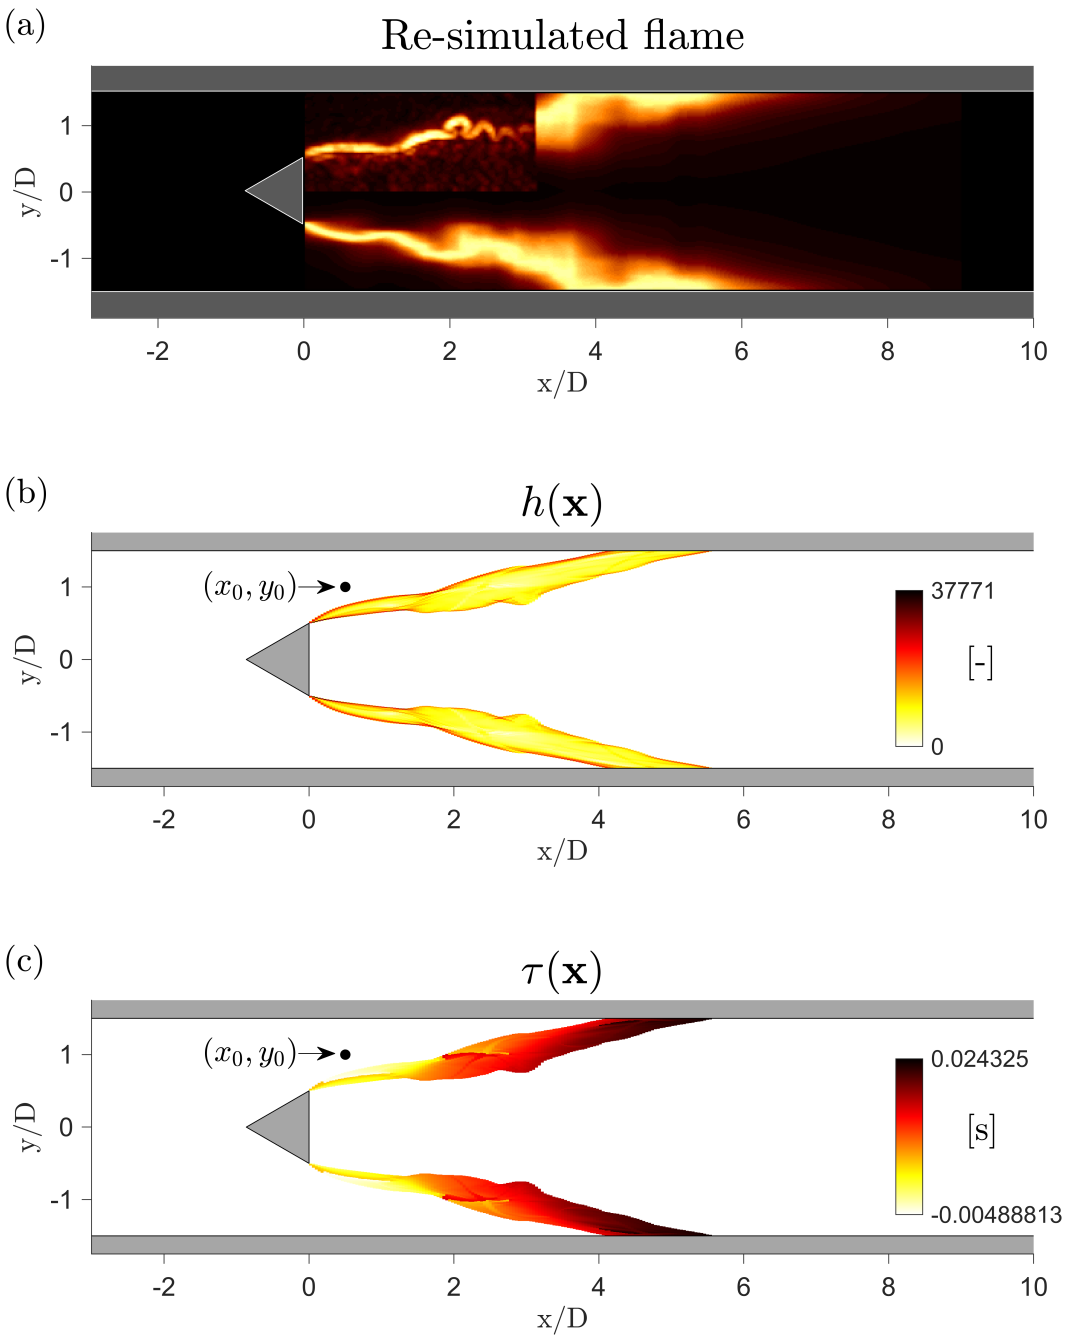

Figure 9: (a): Re-simulated flame shape at $t = 60$. (b) and (c): n and τ fields calculated from the re-simulated flame shape at $t = 60$. The centre (x_0, y_0) of the measurement region $w(x, y)$ is also shown. Part of the flame is upstream of the measurement point, which leads to a negative time delay.

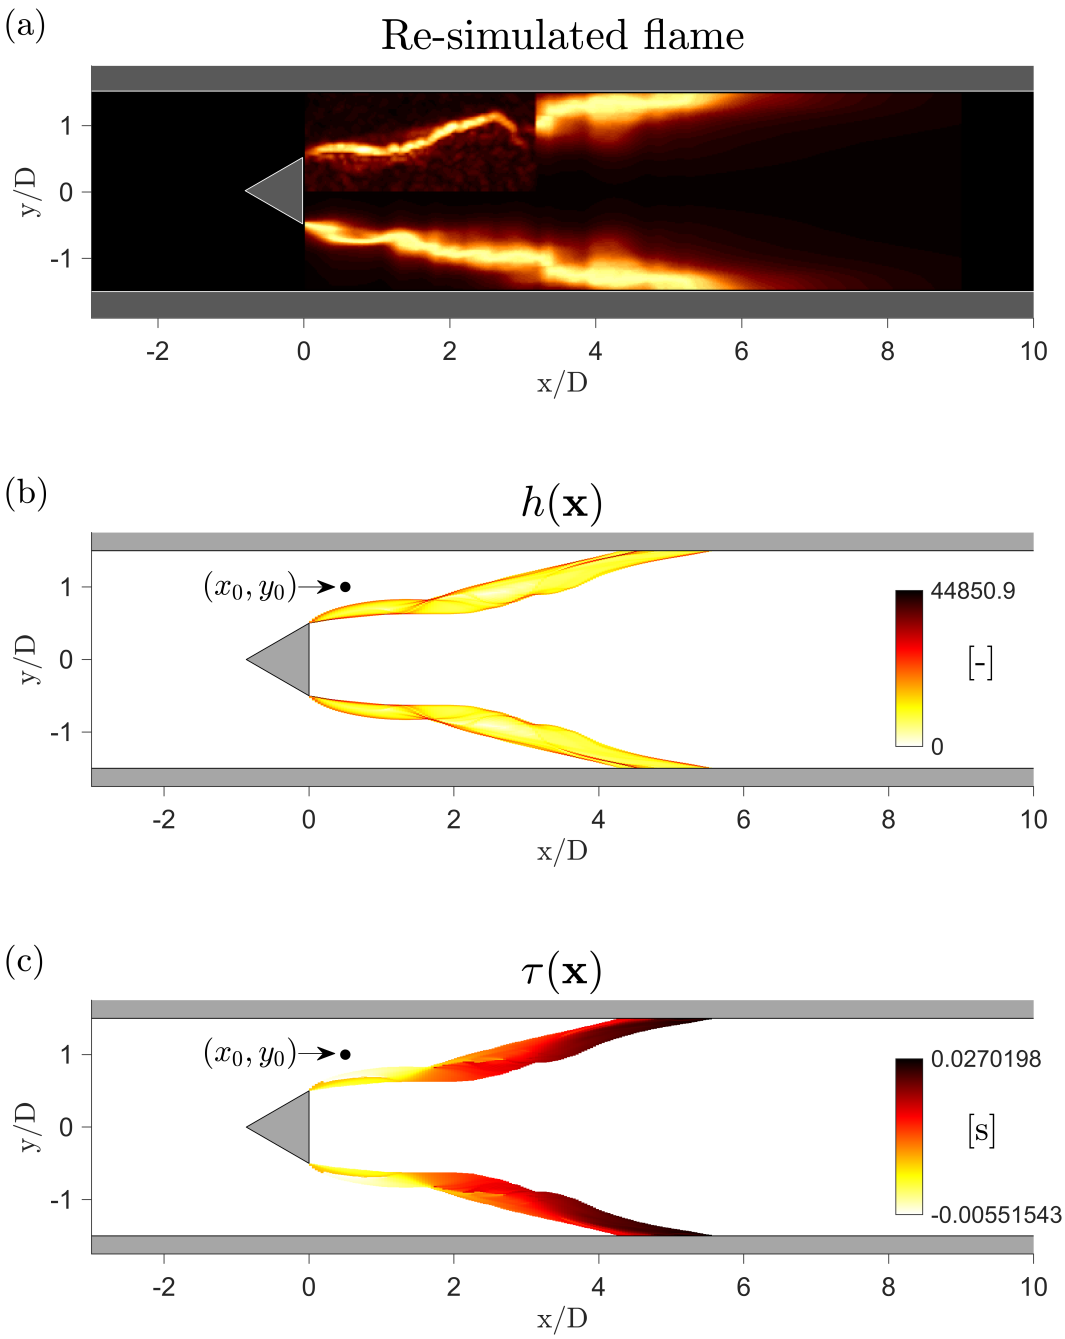

Figure 10: (a): Re-simulated flame shape at $t = 304$. (b) and (c): n and τ fields calculated from the re-simulated flame shape at $t = 304$. The centre (x_0, y_0) of the measurement region $w(x, y)$ is also shown. Part of the flame is upstream of the measurement point, which leads to a negative time delay.

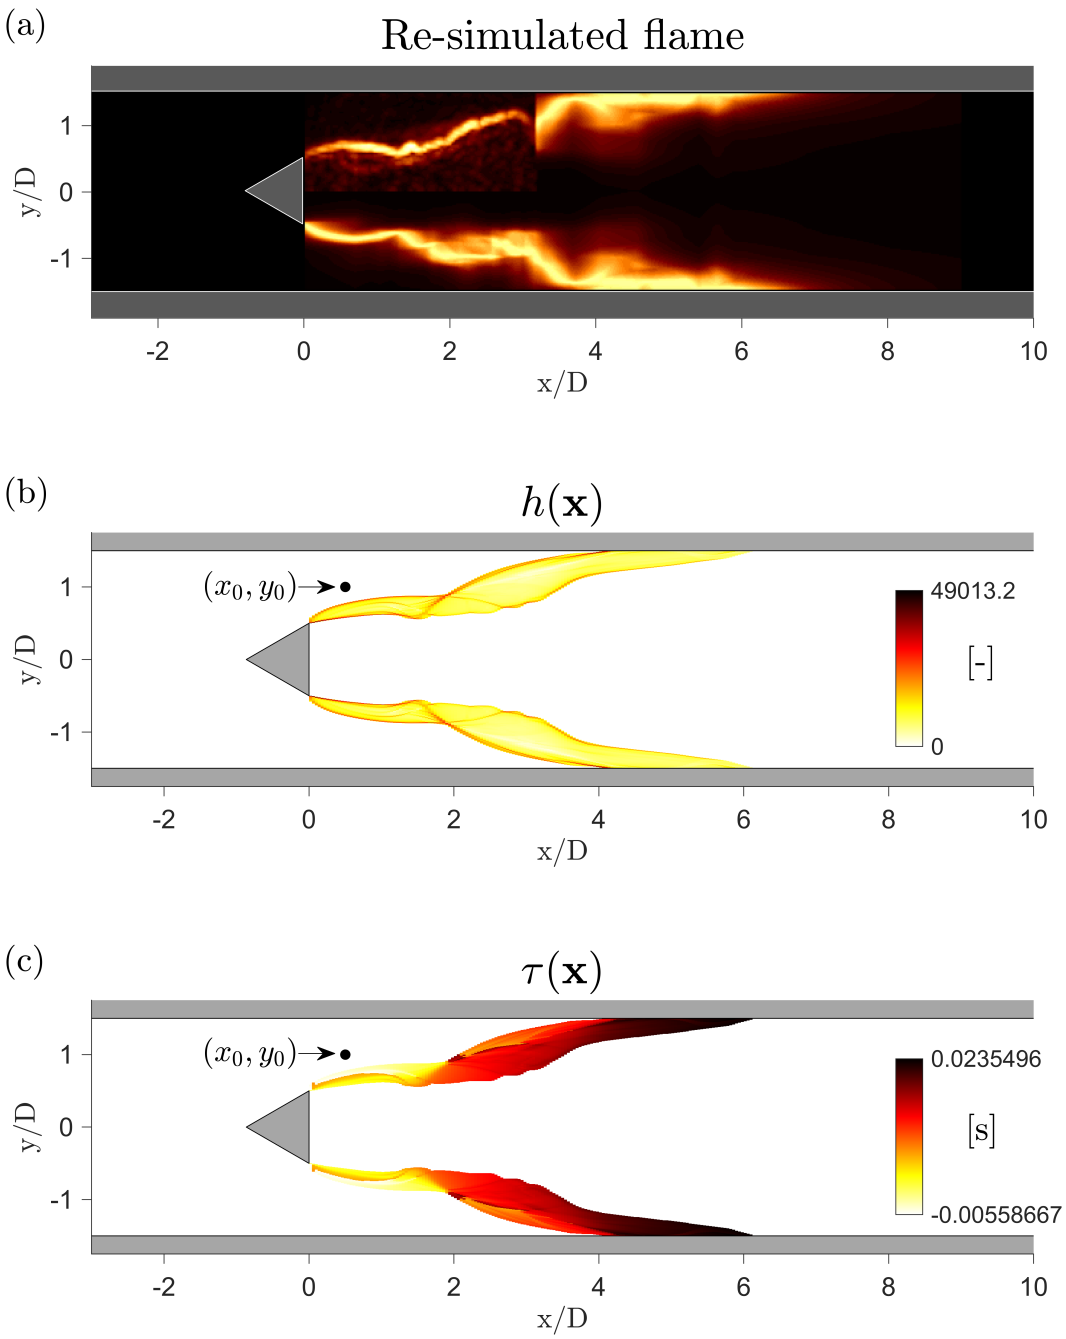

Figure 11: (a): Re-simulated flame shape at $t = 376$. (b) and (c): n and τ fields calculated from the re-simulated flame shape at $t = 376$. The centre (x_0, y_0) of the measurement region $w(x, y)$ is also shown. Part of the flame is upstream of the measurement point, which leads to a negative time delay.

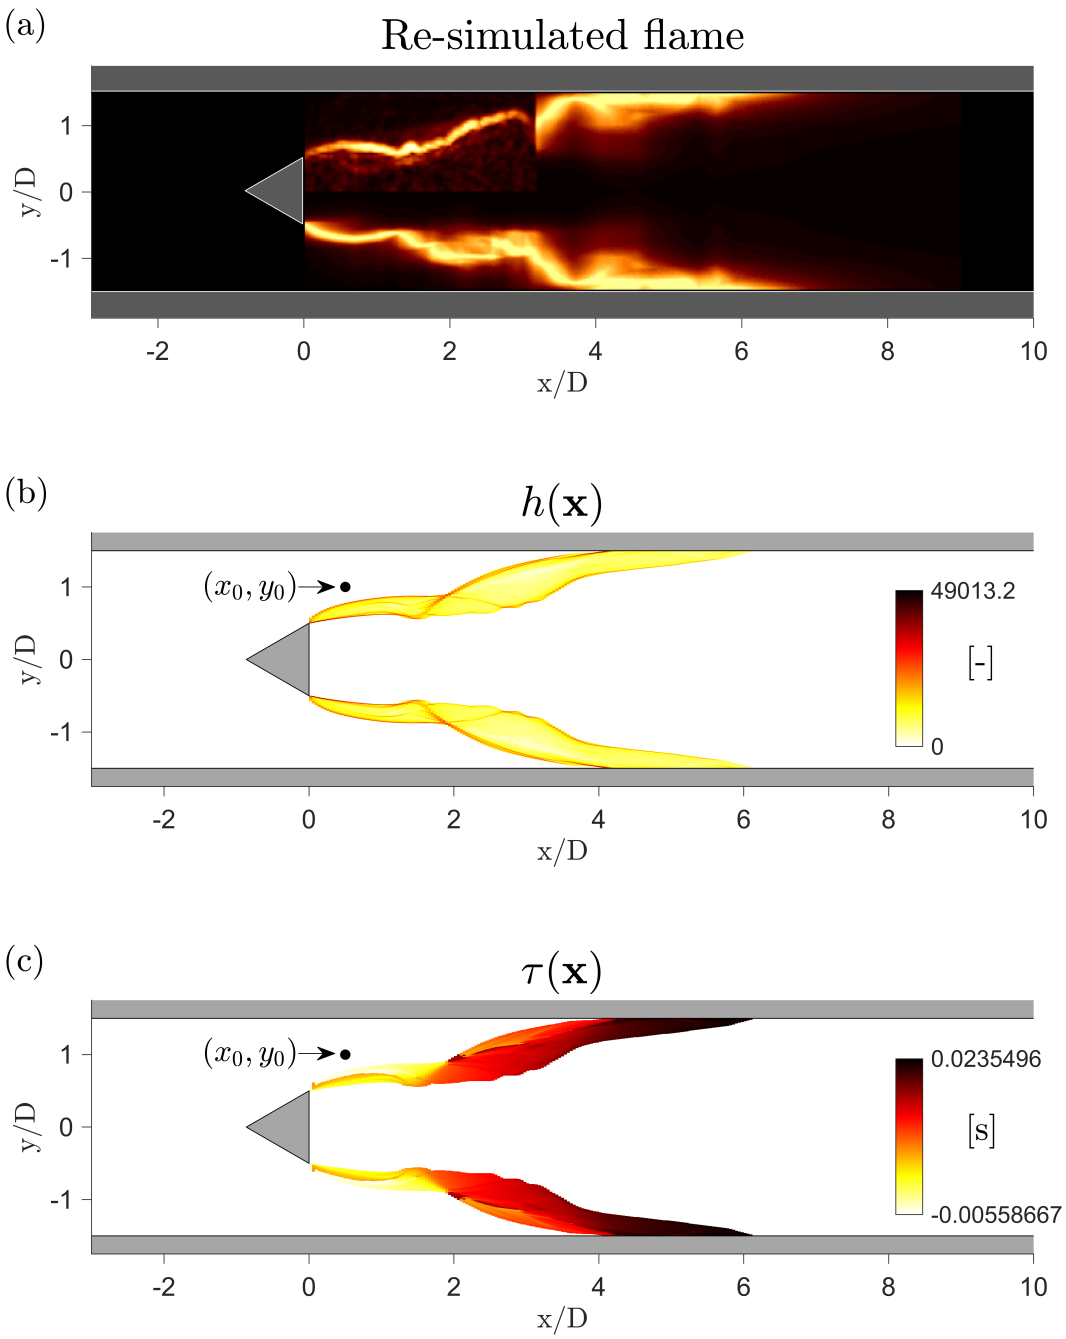

Figure 12: (a): Re-simulated flame shape at $t = 376$. (b) and (c): n and τ fields calculated from the re-simulated flame shape at $t = 376$. The centre (x_0, y_0) of the measurement region $w(x, y)$ is also shown. Part of the flame is upstream of the measurement point, which leads to a negative time delay.

D.2 Parameter-space extrapolation

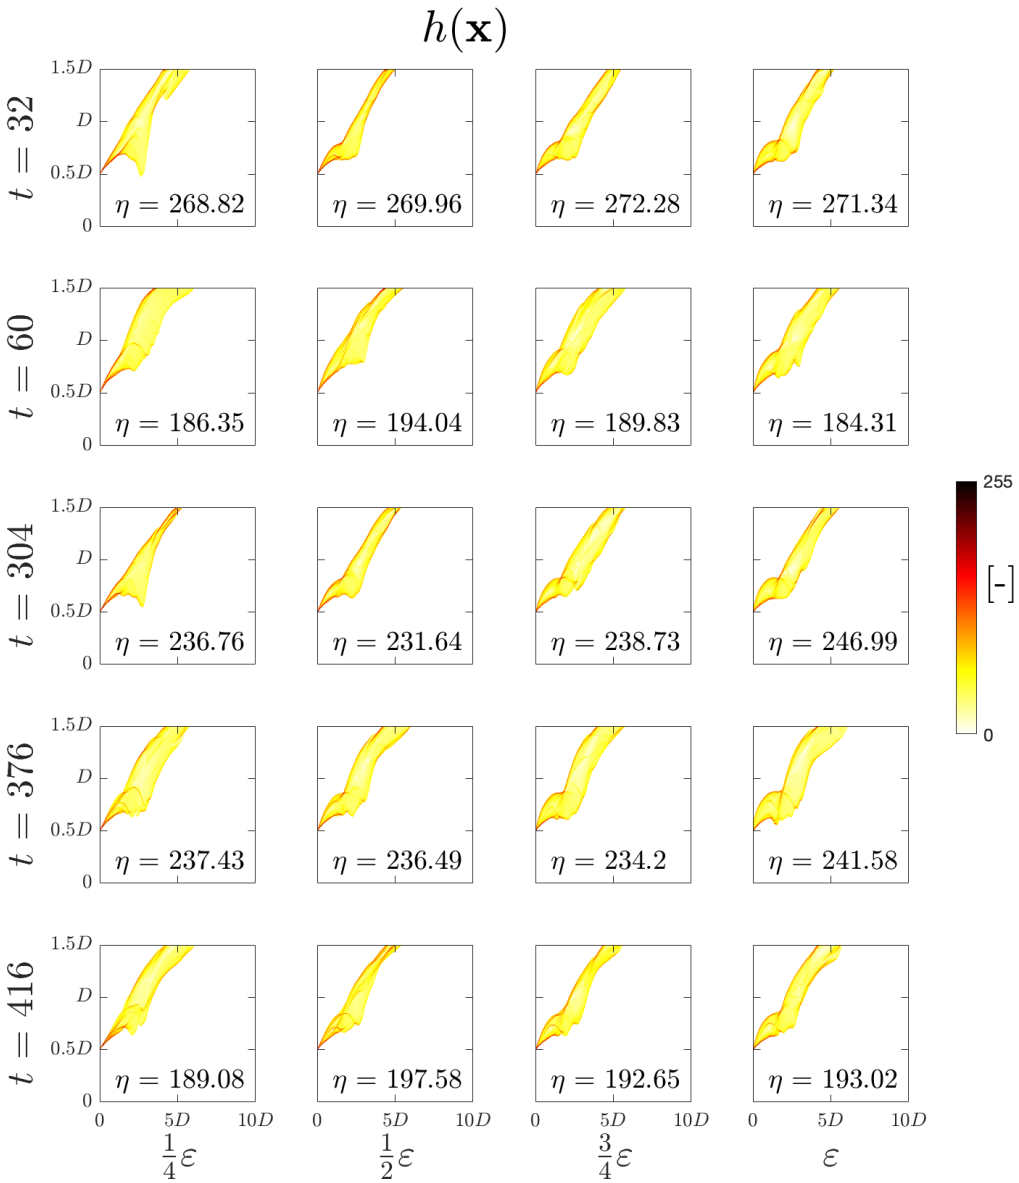

Figure 13: Plots of $h(\mathbf{x})$ at five different time steps t and at four amplitudes: $\varepsilon/4, \varepsilon/2, 3\varepsilon/4$ and ε, the amplitude inferred by the BayNNE. The values of η are also displayed for each time step and amplitude.

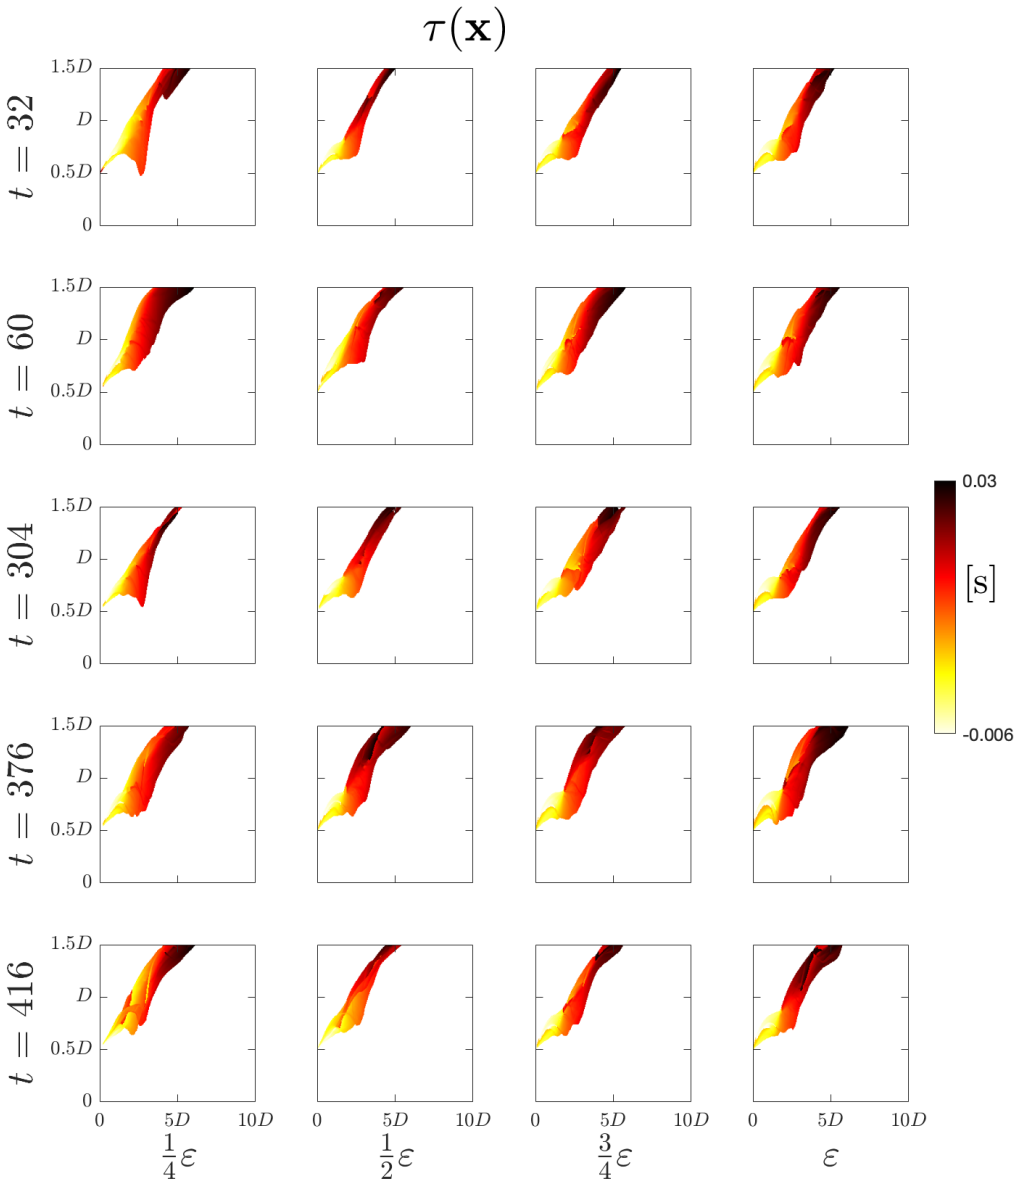

Figure 14: Plots of $\tau(\mathbf{x})$ at five different time steps t and at four amplitudes: $\varepsilon/4, \varepsilon/2, 3\varepsilon/4$ and ε, the amplitude inferred by the BayNNE.

E Supplementary material: The flame behaviour expressed as a distributed $n - \tau$ model

The above analysis gives the local heat release rate field, $q(\mathbf{x}, t)$, and the local velocity field, $u(\mathbf{x}, t)$, over one cycle at a known frequency, ω, where $\mathbf{x} = (x, y)$. At each pixel, we take the Fourier transform of q and u at ω, for example $\hat{q}(\mathbf{x}, \omega) = \int q(\mathbf{x}, t) \exp(i\omega t) dt$, and convert this to the heat release rate per unit volume by dividing by the volume corresponding to each pixel. The n and τ fields are calculated as follows:

$$n(\mathbf{x}) = \left| \frac{\hat{q}(\mathbf{x}, \omega)}{I} \right| \ , \ \ \phi(\mathbf{x}) = \angle \left(\frac{\hat{q}(\mathbf{x}, \omega)}{I} \right) \ , \tag{23}$$

where $I = \int_\Omega w(\mathbf{x}) |\hat{u}(\mathbf{x}, \omega)| d\mathbf{x}$ and $w(\mathbf{x})$ is the measurement zone:

$$w(x, y) = \exp\left(-a_r (x - x_0)^2 - a_r (|y| - y_0)^2 \right) \ , \tag{24}$$

where $a_r = 6200$, and $(x_0, y_0) = (D/2, D)$ where D is the bluff-body side length. This is somewhat arbitrary so long as the measurement region is the same in the thermoacoustic model as it is in the heat release model. The phase, $\phi(\mathbf{x})$, of q is wrapped between $\pm\pi$ and is unwrapped with the algorithm in Ref. [40]. The phase is then divided by ω to give the spatially-distributed time delay $\tau(\mathbf{x})$. Finally the n field is scaled to a field that integrates to 1 multiplied by a scalar, η in units of J/m, in order to be consistent with equation [4] of [41].