

Hard-constrained neural networks for modeling nonlinear acoustics

Original

Hard-constrained neural networks for modeling nonlinear acoustics / Ozan, D. E.; Magri, L.. - In: PHYSICAL REVIEW FLUIDS. - ISSN 2469-990X. - 8:10(2023), pp. 1-33. [10.1103/PhysRevFluids.8.103201]

Availability:

This version is available at: 11583/2995074 since: 2024-12-07T19:17:53Z

Publisher:

American Physical Society - APS

Published

DOI:10.1103/PhysRevFluids.8.103201



Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Hard-constrained neural networks for modeling nonlinear acoustics

Defne E. Ozan ¹ and Luca Magri ^{1,2,*}¹*Department of Aeronautics, Imperial College London, London SW7 2AZ, United Kingdom*²*The Alan Turing Institute, London NW1 2DB, United Kingdom*

(Received 24 May 2023; accepted 3 October 2023; published 25 October 2023)

In this computational paper, we model acoustic dynamics in space and time from synthetic sensor data. The tasks are (i) to predict and extrapolate the spatiotemporal dynamics and (ii) to reconstruct the acoustic state from partial observations. To achieve this, we develop acoustic neural networks. These are networks that learn from sensor data, while being constrained by prior knowledge on acoustic and wave physics. The prior knowledge is constrained as a soft constraint, which informs the training, and as a hard constraint (Galerkin neural networks), which constrains parts of the network’s architecture as an inductive bias. First, we show that standard feedforward neural networks are unable to extrapolate in time, even in the simplest case of periodic oscillations. This motivates the constraints on the prior knowledge. Second, we constrain the prior knowledge on acoustics in increasingly effective ways by (i) employing periodic activations (periodically activated neural networks), (ii) informing the training of the networks with a penalty term that favors solutions that fulfill the governing equations (soft constrained), (iii) constraining the architecture in a physically motivated solution space (hard constrained), and (iv) a combination of these. Third, we apply the networks on two test cases for two tasks in nonlinear regimes, from periodic to chaotic oscillations. The first test case is a twin experiment, in which the data are produced by a prototypical time-delayed model. In the second test case, the data are generated by a higher-fidelity model with mean-flow effects and a kinematic model for the flame source. We find that (i) constraining the physics in the architecture improves interpolation while requiring smaller network sizes, (ii) extrapolation in time is achieved by periodic activations, and (iii) velocity can be reconstructed accurately from only pressure measurements with a combination of physics-based hard and soft constraints. In acoustics and thermoacoustics, this work opens possibilities for physics-constrained data-driven modeling. Beyond acoustics, this work opens strategies for constraining the physics in the architecture, rather than the training.

DOI: [10.1103/PhysRevFluids.8.103201](https://doi.org/10.1103/PhysRevFluids.8.103201)

I. INTRODUCTION

When modeling, reconstructing, and forecasting dynamics from data, constraining prior knowledge into machine learning methods can significantly improve prediction, robustness, and generalizability [e.g., 1–4]. On the one hand, constraints that are imposed in the loss function as penalty terms, which act during training, are referred to as “soft constraints.” Soft constraints

*l.magri@imperial.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

can include the governing partial differential equations (PDEs) derived from conservation laws and initial and boundary conditions [5]. Deep feedforward neural networks with soft-constrained physics information, coined as physics-informed neural networks (PINNs) [6], have been employed to infer flow fields from synthetic data from prototypical flows [7,8], from puffing pool fires [9], experimental data from a flow over an espresso cup [10], and clinical magnetic resonance imaging data [11], to name only a few. Beyond PINNs, physics information has enabled superresolution tasks without high-resolution labels in deep feedforward neural networks [12–14] and in convolutional neural networks [15,16]. On the other hand, constraints that are imposed in the architecture (as opposed to the training) are referred to as “hard constraints.” Hard constraints span the areas of the function space in which physical solutions live; i.e., they create an *inductive bias* [17]. Known PDEs [18–20], invariances [21], Dirichlet boundary conditions [5], and periodic boundary conditions [16,22,23] have been incorporated in the architecture of neural networks as hard constraints. In this paper, we design Galerkin neural networks, which hard-encode the solution structure into the architecture. These networks are inspired by Galerkin projection, which is a common technique for solving PDEs by projecting the equations onto a finite set of modes, which transforms the problem into a set of ordinary differential equations [24]. The library of modes can be a generic basis such as Fourier, or data driven such as proper orthogonal decomposition (POD) modes. In this paper, we take advantage of a physical basis.

We focus on solutions from acoustics and thermoacoustics, which originate from nonlinear wave equations. Thermoacoustic systems contain nonlinearities in the heat-release model, which, when coupled with the acoustics in a positive feedback loop, can generate self-excited oscillations and rich nonlinear behaviors via bifurcations (e.g., [25–29]). Because these oscillations can have detrimental effects on the system’s structure and performance, their prediction and control are active areas of research, for example, in gas turbines [30,31] and rocket engines [32]. Traditionally, the prediction of thermoacoustics has been achieved with first principles. In the time domain, a direct approach is the brute-force time integration of the governing equations following a discretization scheme. High-fidelity models such as large-eddy simulations that model the acoustics and the flame simultaneously on fine grids provide highly accurate solutions, but they are computationally expensive [31]. Low-fidelity models reduce the computational effort at the expense of accuracy while obtaining models that can be used for stability, bifurcation analysis, and parametric studies. Generally, these approaches combine a linear acoustic solver with a heat-release model. In this direction, nonlinear behavior of longitudinal and annular combustors has been investigated by employing network models based on the traveling wave (e.g., [33–37]), Galerkin decomposition of pressure and velocity using only a finite number of acoustic eigenmodes and projection of the PDEs onto those modes (e.g., [38,39]), and numerical discretizations of the PDEs [40,41]. Predictions with first principles only are either computationally expensive (e.g., large-eddy simulation), or as good as the model assumptions. Because thermoacoustics is a multiphysics phenomenon, the model assumptions made in low-order models are inevitably substantial. This motivates adding data into the first-principles modeling of thermoacoustics, for which machine learning methods excel. The data typically come from laboratory experiments that are conducted with setups consisting of a duct with a heat source, which can be a flame or an electrically heated wire mesh (e.g., [42,43]). In the experiments, the collected data are usually the acoustic pressure measured by microphones at high sampling rates. In this direction, data assimilation techniques have been applied to improve physics-based qualitatively accurate models of a prototypical thermoacoustic system, while estimating model parameters. Nova and Magri [44] combined a thermoacoustic low-order model with an ensemble Kalman filter, and an echo-state network for real-time bias-aware data assimilation of the acoustic state [45]. In other applications, neural networks have been developed for heat-release models or flame response inference in thermoacoustic applications [46–48]. In [49], a physics-informed feedforward neural network approach was developed for learning thermoacoustic limit cycles by using periodic activation functions. Physics-informed neural networks were also shown to be successful for the reconstruction of acoustic fields with nonideal boundary conditions [50].




FIG. 1. Example of a standard feedforward neural network (FNN) (e.g., [52]) with two hidden layers. This is the data-driven-only model, which has no prior knowledge embedded in the architecture. The trainable parameters are the weights, $\mathbf{W}^{(l)}$, and the biases, $\mathbf{b}^{(l)}$.

The overarching goal of this paper is to generalize and develop acoustic neural networks to embed the structure of the nonlinear wave solution into the architecture. The specific objective of this paper is threefold: (i) to predict and extrapolate in time thermoacoustic oscillations, (ii) to reconstruct pressure and velocity over the entire domain from pressure sensors only, and (iii) to obtain a model that is robust to noise and generalizable to unseen scenarios. This paper is organized as follows. In Sec. II, we provide the mathematical background for feedforward neural networks and Galerkin decompositions. In Sec. III, we introduce Galerkin neural networks. In Sec. IV, we discuss the application of Galerkin neural networks to acoustics and thermoacoustics. In Secs. V and VI, we show results for twin experiments on synthetic data from a Rijke tube and on synthetic data from a higher-fidelity model. The paper ends with a conclusion section.

II. BACKGROUND

A. Standard feedforward neural networks

Let $\mathbf{y} \in \mathbb{R}^{N_y}$ represent some vector of physical quantities pertaining to a system that depends on space, $\mathbf{x} \in \mathbb{R}^{N_d}$, where N_d is the dimension, and time, $t \in \mathbb{R}$. Then, given full or partial observations of \mathbf{y} , our goal is to learn a model \mathbf{f} that predicts an output vector $\hat{\mathbf{y}} \in \mathbb{R}^{N_y}$ from an input vector (\mathbf{x}, t) while minimizing an error metric, \mathcal{L} . A feedforward neural network (FNN) is defined by a composition of functions, \mathbf{f} , which, when appropriately designed, can approximate any continuous function in a specified range [51]:

$$\mathbf{f}(\mathbf{x}, t) := \mathbf{f}^{(L)}(\mathbf{f}^{(L-1)}(\dots \mathbf{f}^{(1)}(\mathbf{x}, t))), \quad (1a)$$

$$\mathbf{f}^{(l)}(\mathbf{z}) := \boldsymbol{\phi}^{(l)}(\mathbf{W}^{(l)}\mathbf{z} + \mathbf{b}^{(l)}), \quad (1b)$$

where $\mathbf{f}^{(l)}$ is the function that maps the layer $l-1$ to the layer l , where $l = 1, \dots, L$ and L is the number of layers; \mathbf{W} are the weight matrices and \mathbf{b} are the biases, which are the trainable parameters; and $\boldsymbol{\phi}^{(l)}$ are the activation functions from the layer $l-1$ to the layer l , which are applied to each component of the argument. For regression, in the last layer, the activation $\boldsymbol{\phi}^{(L)}$ is linear. The neural network offers an ansatz for a continuous function through linear operations and simple nonlinear activations (please refer to [52] for a pedagogical and geometric explanation).

The FNN is a standard architecture that is fully data driven; i.e., no prior knowledge is embedded in the network (Fig. 1). The network's weights and biases, collectively grouped in a variable $\boldsymbol{\chi}$, are optimized via gradient descent to minimize an error, which is provided by a loss function \mathcal{L}

$$\boldsymbol{\chi}^* = \arg \min_{\boldsymbol{\chi}} \mathcal{L}(\boldsymbol{\chi}). \quad (2)$$

When no physics constraint is imposed as in standard neural networks, the data-driven loss is quantified by the mean-squared error (MSE) between the measured data and predictions of the network. In the formulation of the data-driven loss, we also account for the cases of partial state observations with a measurement matrix $\mathbf{M} \in \mathbb{R}^{N_y \times N_y}$ that indicates which states are measured:

$$M_{ij} = \begin{cases} 1 & \text{when } i = j \text{ and } y_i \text{ is measured} \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

When full state measurements are available, then \mathbf{M} is the identity matrix, $\mathbf{I}_{N_y \times N_y}$. The loss is given as

$$\mathcal{L} \equiv \mathcal{L}_{DD} = \frac{1}{NN_m} \sum_{k=1}^N \|\mathbf{M}\mathbf{y}_k - \mathbf{M}\hat{\mathbf{y}}_k\|_2^2, \quad (4)$$

where the subscript DD stands for data driven, $N_m \leq N_y$ is the number of measured states, the subscript 2 denotes the ℓ_2 norm, and the subscript k denotes the k th element in the data set of N pairs of input and output vectors. In Sec. IV B, we tailor the neural networks to acoustic problems by specifying the data, input vectors, task, activation functions, nonlinear maps, and loss functions. The prior knowledge will be embedded in different ways in the neural networks.

B. Separation of variables and Galerkin methods

Separation of variables seeks to find solutions to partial differential equations in a special form as the product of functions of the individual independent variables, e.g., time and single space coordinates, which results in an eigenvalue problem [53]. Linear wave equations and wave solutions can thus be represented in an acoustic eigenbasis, namely, Fourier modes, which is complete. On the other hand, the thermoacoustic problem is governed by a wave equation with a nonlinearity arising from the heat release (please refer to Sec. IV A for details). However, the nonlinearity will only slightly change the frequency and mode coupling, which motivates the use of the acoustic eigenspace as an expressive basis to represent the nonlinear solutions as well [38]. A weak solution to the nonlinear problem is provided by the Galerkin method, which approximates the solution to the PDE by projecting the equations onto a finite-dimensional subspace [24]. The solution is written as a linear combination of the basis functions that span this subspace [24],

$$y_i(x, t) = \sum_{j=1}^{N_g} \alpha_j^i(t) \Psi_j^i(x), \quad (5)$$

where $\mathbf{y} \in \mathbb{R}^{N_y}$ is a vector of physical quantities (Sec. II A) and $i = 1, 2, \dots, N_y$ denotes the i th element of \mathbf{y} , $\Psi_j^i(x)$ are the basis functions, or Galerkin modes, N_g is the number of Galerkin modes, and $\alpha_j^i(t)$ are the Galerkin amplitudes. The orthogonality of the Galerkin modes guarantees that the approximation error is orthogonal to the space spanned by the finite number of modes retained. This property makes the Galerkin method a suitable choice for the modeling of acoustic problems because the low-frequency modes contain most of the energy and, thus, the truncation error can be minimal.

III. GALERKIN NEURAL NETWORKS

FNNs are flexible and useful tools for function approximation; however, they can suffer from under- or overfitting, especially in the case of scarce or noisy data. This requires a careful hyperparameter tuning, which becomes computationally expensive in proportion with the size of network. Even then, the network may not be able to capture some features of the solution because of missing data. In order to counteract these shortcomings, we exploit the physical knowledge about the spatiotemporal basis of the system in question, and propose a network structure inspired




FIG. 2. Galerkin neural network (GalNN). This network is composed of two branches: (a) a spatial branch (hard constraint) and (b) a temporal branch, which learns the temporal behavior. The trainable parameters are the weights, $\mathbf{W}^{(\cdot)}$, and the biases, $\mathbf{b}^{(\cdot)}$

by the Galerkin decomposition of the system as motivated in Sec. II B. The chosen Galerkin modes are a known nonlinear transformation of the spatial coordinates, which is introduced in the network. Therefore, by design, this network can be configured to automatically satisfy the boundary conditions. We will refer to this network architecture as the *Galerkin neural network* (GalNN). The Galerkin network is composed of a spatial branch that transforms \mathbf{x} into the *a priori* Galerkin modes, $\Psi^i(\mathbf{x})$ [Eq. (5)], and a temporal branch that is an FNN [Eq. (1a)] that takes only t as an input and predicts the unknown Galerkin amplitudes, $\alpha^i(t)$ [Eq. (5)], as outputs. The final outputs, $\hat{\mathbf{y}}$, are computed by Eq. (5). Formally, the GalNN is defined as

$$f_i(\mathbf{x}, t) = \sum_{(i-1)N_g+1}^{iN_g} g_j(\mathbf{x})h_j(t), \quad (6)$$

where \mathbf{f} is the map from (\mathbf{x}, t) to $\hat{\mathbf{y}}$, $i = 1, 2, \dots, N_y$ denotes the i th element of the output vector, $\mathbf{g}(\mathbf{x})$ is the spatial branch given by $\mathbf{g}(\mathbf{x}) = (\Psi^1(\mathbf{x}), \Psi^2(\mathbf{x}), \dots, \Psi^{N_y}(\mathbf{x}))$, and $\mathbf{h}(t)$ is the temporal branch modeled as an FNN, the outputs of which are denoted as $(\alpha^1(t), \alpha^2(t), \dots, \alpha^{N_y}(t))$. The data-driven loss is defined the same as the standard FNN case [Eq. (4)] over the prediction error on \mathbf{y} . The time evolution of the Galerkin amplitudes is thus obtained as an intermediary step. This architecture is shown in Fig. 2.




FIG. 3. Schematic of the thermoacoustic system. A straight duct with open ends and a heat source, q , located at x_f . The acoustic reflection coefficients at the inlet and outlet are R_{in} and R_{out} . The Rijke tube is employed in Sec. V for twin experiments with the following assumptions: (i) zero Mach number, i.e., $\bar{u}_1 = \bar{u}_2 = 0$; (ii) no change in the mean flow variables before and after the flame, e.g., $\bar{\rho}_1 = \bar{\rho}_2 = 1$; and (iii) ideal boundary conditions, i.e., $R_{in} = R_{out} = -1$. A higher-fidelity model with a mean-flow and a kinematic model for the flame is employed in Sec. VI for analyzing model generalization.

IV. APPLICATION TO ACOUSTICS AND THERMOACOUSTICS

In this paper, we develop a physics-constrained data-driven model of the acoustic variables as a function of time and space. Given enough capacity by means of number of neurons and layers, FNNs can fit the data within the training range. This can be considered as an interpolation problem, by which we mean that the network predicts on data points that are within the bounds of the training input range. When the prediction is performed for scenarios outside the training input range, we call this an extrapolation problem. Thermoacoustic systems exhibit oscillations, which can be periodic, quasiperiodic, or chaotic in time. First, given time-series data from a time window, our task is to obtain a model that can extrapolate such oscillatory behavior. In experiments, full-state observations may be unavailable or prohibitively expensive; e.g., in acoustics, only pressure data may be available through measurements with microphones. Second, our task is to reconstruct the flow variables over the entire spatial domain from full-state or partial noisy state observations, which poses a challenge for purely data-driven models. Third, we seek a robust, generalizable, low-order model for acoustic and thermoacoustic solutions using as few data as possible.

A. Background on thermoacoustics

We consider a thermoacoustic system composed of a straight duct of length \tilde{L} with a compact heat source located at $\tilde{x} = \tilde{x}_f$, where $\tilde{(\cdot)}$ denotes a dimensional variable (Fig. 3). We make the following assumptions about the system: (i) the acoustics are one dimensional, i.e., the tube is sufficiently longer than its diameter for the cut-on frequency to be large enough for longitudinal acoustics only to propagate; (ii) the heat release from the flame acts as a pointwise monopole source of sound (compact assumption); (iii) the mean flow is low Mach number and there is no entropy wave convection; and (iv) effects of viscosity and heat conduction are negligible (e.g., [28]). The reflection of the acoustic waves at the boundaries is given by the reflection coefficients R_{in} and R_{out} for inlet and outlet, respectively, which determine the boundary conditions.

The dynamics are governed by the dimensional equations derived from mass, momentum, and energy conservation. Modeling the heat source as a compact source results in two duct segments related by jump conditions that are enforced at the heat source location. For brevity, the subscripts 1 and 2 denote conditions before and after the flame, i.e., $\tilde{x} = \tilde{x}_{f,1}$ and $\tilde{x} = \tilde{x}_{f,2}$, respectively. The jump conditions are found from mass, momentum, and energy fluxes across the flame with the ideal gas law (e.g., [54]). The governing equations and the jump conditions are linearized by assuming that

TABLE I. Summary of the proposed neural networks. All the networks can be equipped with a soft constraint, which penalizes nonphysical solutions in the loss function (see Sec. IV B 2). When the soft constraint is present, the abbreviation will be prefixed with PI (physics informed). For example, a Galerkin network with the soft constraint will be abbreviated as PI-GalNN.

Network	Abbreviation	Constraint	Separation of time and space	Can extrapolate in time	Section
Feedforward neural network	FNN	No	No	No	II A
Periodically activated feedforward neural network	P-FNN	No	No	Yes, via activation function	IV B 1
Galerkin neural network	GalNN	Hard	Yes	Yes	IV B 3
Periodic Galerkin neural network	P-GalNN	Hard	Yes	Yes, for limit cycles it guarantees fundamental frequency	IV B 4

the flow variables can be expressed as infinitesimal perturbations on top of a mean flow, i.e., $(\tilde{\cdot}) = (\bar{\cdot}) + (\tilde{\cdot})'$, where $(\bar{\cdot})$ denotes the steady mean-flow variable and $(\tilde{\cdot})'$ denotes the unsteady infinitesimal perturbations. Under the low-Mach-number assumption, the acoustics are governed by momentum and energy equations (e.g., [28])

$$\tilde{\rho} \frac{\partial \tilde{u}'}{\partial \tilde{t}} + \tilde{\rho} \tilde{u} \frac{\partial \tilde{u}'}{\partial \tilde{x}} + \frac{\partial \tilde{p}'}{\partial \tilde{x}} = 0, \quad (7a)$$

$$\frac{\partial \tilde{p}'}{\partial \tilde{t}} + \tilde{u} \frac{\partial \tilde{p}'}{\partial \tilde{x}} + \gamma \tilde{p} \frac{\partial \tilde{u}'}{\partial \tilde{x}} - (\gamma - 1) \frac{\tilde{q}'}{\tilde{A}} \delta(\tilde{x} - \tilde{x}_f) = 0, \quad (7b)$$

where $\tilde{\rho}$ is the density, \tilde{u} is the velocity, \tilde{p} is the pressure, \tilde{q} is the heat-release rate, \tilde{A} is the cross-section area of the duct, γ is the heat capacity ratio, and δ is the Dirac delta distribution. In Secs. V and VI, the governing equations are employed with different simplifications to generate the synthetic data for our study. We will impose the governing equations in their nondimensional PDE form, denoted by omitting the symbol $(\tilde{\cdot})$, as soft constraints during the training of our neural networks. This approach will be discussed in detail in Sec. IV B 2.

B. Acoustic neural networks

The dynamics of thermoacoustic oscillations are dominated by unstable eigenfunctions (e.g., [28]), which are periodic oscillations in time. We propose neural networks that can naturally infer acoustic dynamics, which are the basis functions of nonlinear thermoacoustic behaviors. The proposed networks embed the prior knowledge through the activation functions (Sec. IV B 1), through a penalization loss function in the training (soft constraint, Sec. IV B 2), and through the architecture (hard constraint, Sec. IV B 3). For the thermoacoustic system described in Sec. IV A, the input vector consists of the one-dimensional spatial coordinate, x , and time, t , and the output vector consists of the acoustic pressure and velocity fluctuations, $\mathbf{y} = (p', u')$. In Sec. IV B 1, we motivate employing periodic activation functions in the standard FNN for acoustic problems, the eigenfunctions of which are periodic. Furthermore, in Sec. IV B 2, we include a physics-based regularization term that penalizes solutions that violate the conservation laws governing the acoustic dynamics. Finally, in Sec. IV B 3, we promote a hard-constrained architecture in the form of a GalNN (Sec. III). Inspired by the physical remarks of Sec. II B, we design a neural network architecture that spans the Hilbert space with the acoustic eigenfunctions as the Galerkin modes, while having trainable parameters for inference and closure of the unknown physical terms. The acoustic neural networks are summarized at the end of this section in Table I.

1. Periodically activated feedforward neural networks (P-FNNs)

In order to augment the extrapolation capability, observe that nonlinear thermoacoustic dynamics originate from the nonlinear coupling of acoustic eigenfunctions (e.g., [28]), which are Fourier modes. From a data-driven perspective, this means that the weights and biases should be periodically activated. A straightforward strategy is to employ periodic activations in the acoustic neural networks, i.e., $\phi^{(l)} = \sin(z)$ in Eq. (1b) [49]. The physics of the system, namely, the periodic nature of the solutions, is embedded in the network itself via the choice of activation function, which provides an inductive bias on the function space of the network and improves extrapolation [55]. The weights of a layer l with the sine activation are initialized from a uniform distribution in the range $[-\sqrt{3/\text{fan}_{\text{in}}}, \sqrt{3/\text{fan}_{\text{in}}}]$, where fan_{in} is the number of neurons in the layer $l - 1$, n_{l-1} [55].

2. Soft constraints with physics-informed losses

Prior knowledge can be embedded as a penalization term in the loss function in Eq. (2) to minimize during training (e.g., [5,6]). This approach improves the generalizability of the network, while promoting physical outputs [3]. In thermoacoustics, neural network predictions should fulfill the acoustic conservation laws; therefore, we penalize solutions that violate momentum and energy equations in the loss function

$$\mathcal{L} = \lambda_{DD}\mathcal{L}_{DD} + \lambda_M\mathcal{L}_M + \lambda_E\mathcal{L}_E, \quad (8)$$

where \mathcal{L}_{DD} is the data-driven loss [Eq. (4)] (Sec. II A), and \mathcal{L}_M and \mathcal{L}_E are the residual losses from the conservation of momentum and energy equations, respectively. The non-negative scalars λ_{DD} , λ_M , and λ_E are regularization hyperparameters. Obtaining momentum and energy losses, \mathcal{L}_M and \mathcal{L}_E , requires the evaluation of the physical residuals of the network predictions. We denote the partial differential operators that define the residual from momentum and energy equations as $\mathcal{F}_M(p', u')$ and $\mathcal{F}_E(p', u')$, respectively. A given (p', u') is a solution of the system if the residuals are zero, i.e., $\mathcal{F}_M(p', u') = 0$ and $\mathcal{F}_E(p', u') = 0$. We will exactly define these residuals separately for the Rijke tube in Sec. V A and for the higher-fidelity model in Sec. VI A along with the modeling assumptions and nondimensionalization of the original Eqs. (7). The physics-informed losses are then computed as

$$\mathcal{L}_M = \frac{1}{N + N_s} \sum_{k=1}^{N+N_s} \mathcal{F}_M(\hat{p}'_k, \hat{u}'_k)^2, \quad (9)$$

$$\mathcal{L}_E = \frac{1}{N + N_s} \sum_{k=1}^{N+N_s} \mathcal{F}_E(\hat{p}'_k, \hat{u}'_k)^2, \quad (10)$$

where N is the number of training data points, i.e., $\{x_k, t_k\}_{k=1}^N$, and N_s is the number of uniformly sampled points over the whole training domain, as we can evaluate the physical loss at any location in time and space. For this purpose, we employ automatic differentiation using `tf.GradientTape` functionality from TensorFlow [56]. The automatic differentiation of the output variables with respect to the input variables yields the Jacobian that contains the partial derivatives, $\frac{\partial \hat{p}'}{\partial x}$, $\frac{\partial \hat{p}'}{\partial t}$, $\frac{\partial \hat{u}'}{\partial x}$, $\frac{\partial \hat{u}'}{\partial t}$. These partial derivatives are plugged into the operators, \mathcal{F}_M and \mathcal{F}_E , to calculate the residual of the network.

3. Hard constraints with Galerkin neural networks (GalNNs)

A suitable eigenbasis for the thermoacoustic problem is provided by the natural acoustic eigenfunctions, because most of the energy of the solution is dominated by the first acoustic modes, and effects of mean flow or nonlinearity only slightly change the frequency and the eigen-shapes (Sec. II B). The core eigenfunctions of the thermoacoustic system with zero-Mach-number case are similar to a low-Mach-number case, $M \lesssim 0.1$ [57]. Hence, to derive them, we assume

a zero Mach number for simplification, whose dynamics are governed by the nondimensional equations (e.g., [57])

$$\bar{\rho} \frac{\partial u'}{\partial t} + \frac{\partial p'}{\partial x} = 0, \quad (11a)$$

$$\frac{\partial p'}{\partial t} + \frac{\partial u'}{\partial x} - \dot{q}' \delta(x - x_f) = 0, \quad (11b)$$

where the density is modeled as $\bar{\rho} = \bar{\rho}_1$ when $0 \leq x < x_f$, and $\bar{\rho} = \bar{\rho}_2$ when $x_f < x \leq 1$. The dimensional variables have been scaled as $x = \tilde{x}/\tilde{L}$; $t = \tilde{t}\tilde{c}_{\text{ref}}/\tilde{L}$, where \tilde{c}_{ref} is the mean speed of sound; $u = \tilde{u}/\tilde{c}_{\text{ref}}$; $\rho = \tilde{\rho}/\tilde{\rho}_{\text{ref}}$; $p = \tilde{p}/(\tilde{\rho}_{\text{ref}}\tilde{c}_{\text{ref}}^2)$; and $\dot{q} = \tilde{q}(\gamma - 1)/(\tilde{\rho}_{\text{ref}}\tilde{c}_{\text{ref}}^3)$. A physically justified method to solve the set of PDEs (11) is to decompose pressure and velocity as

$$p'(x, t) = \sum_{j=1}^{N_g} \begin{cases} \mu_j(t)\Pi_j^{(1)}(x), & 0 \leq x < x_f \\ \mu_j(t)\Pi_j^{(2)}(x), & x_f < x \leq 1, \end{cases} \quad (12a)$$

$$u'(x, t) = \sum_{j=1}^{N_g} \begin{cases} \eta_j(t)\Upsilon_j^{(1)}(x), & 0 \leq x < x_f \\ \eta_j(t)\Upsilon_j^{(2)}(x), & x_f < x \leq 1, \end{cases} \quad (12b)$$

and project the equations onto the acoustic eigenfunctions found as below in the case of ideal boundary conditions, i.e., $p'(x=0) = p'(x=1) = 0$ [57],

$$\Pi_j^{(1)}(x) = -\sin(\omega_j\sqrt{\bar{\rho}_1}x), \quad \Pi_j^{(2)}(x) = -\left(\frac{\sin \gamma_j}{\sin \beta_j}\right) \sin(\omega_j\sqrt{\bar{\rho}_2}(1-x)), \quad (13a)$$

$$\Upsilon_j^{(1)}(x) = \frac{1}{\sqrt{\bar{\rho}_1}} \cos(\omega_j\sqrt{\bar{\rho}_1}x), \quad \Upsilon_j^{(2)}(x) = -\frac{1}{\sqrt{\bar{\rho}_2}} \left(\frac{\sin \gamma_j}{\sin \beta_j}\right) \cos(\omega_j\sqrt{\bar{\rho}_2}(1-x)), \quad (13b)$$

where N_g is the number of Galerkin modes retained in the approximation, and $\gamma_j = \omega_j\sqrt{\bar{\rho}_1}x_f$, $\beta_j = \omega_j\sqrt{\bar{\rho}_2}(1-x_f)$. The acoustic angular frequencies ω_j are the solutions of the dispersion relationship, $\sin \beta_j \cos \gamma_j + \cos \beta_j \sin \gamma_j \sqrt{\frac{\bar{\rho}_1}{\bar{\rho}_2}} = 0$. In the limit of no jump in mean-flow density, i.e., $\bar{\rho}_1 = \bar{\rho}_2$, $\omega_j = j\pi$, which are the natural acoustic angular frequencies [57]. Motivated by the Galerkin decomposition (12), we express the acoustic and thermoacoustic solutions as a GalNN [Eq. (6)], where $\Pi_j^{(1)}$, $\Pi_j^{(2)}$ [Eq. (13a)], and $\Upsilon_j^{(1)}$, $\Upsilon_j^{(2)}$ [Eq. (13b)] are the Galerkin modes, and μ_j and η_j are the Galerkin amplitudes. The pressure and velocity are determined from the predicted Galerkin amplitudes using Eqs. (12). This architecture is shown in Figs. 4(a) and 4(b).

Whereas the acoustic pressure is continuous across the flame, the acoustic velocity undergoes a discontinuity due to the dilation from the heat-release rate [28]. When such a discontinuity is approximated with a finite Fourier series, such as the Galerkin modes we are using, the Gibbs phenomenon manifests itself as high-frequency oscillations around the discontinuity [40,58], which causes an unphysical behavior. To physically capture the discontinuity and eliminate unphysical oscillations in the predictions, we add two step functions in the velocity modes,

$$\Upsilon_{N_g+1}^{(1)}(x) = 1, \quad \Upsilon_{N_g+1}^{(2)}(x) = 0, \quad (14a)$$

$$\Upsilon_{N_g+2}^{(1)}(x) = 0, \quad \Upsilon_{N_g+2}^{(2)}(x) = 1, \quad (14b)$$

such that the summation (12b) runs from 1 to $N_g + 2$. The modes $\Upsilon_{N_g+1}^{(\cdot)}$ and $\Upsilon_{N_g+2}^{(\cdot)}$ are weighted by independent coefficients, η_{N_g+1} and η_{N_g+2} , which allows them to capture a jump discontinuity at the flame location.

The temporal branch of the GalNN should be capable of extrapolation in time of thermoacoustic oscillations. These solutions bifurcate from limit-cycle oscillations, which are periodic, to quasiperiodic and chaotic oscillations. As they originate from periodic solutions, even the quasiperiodic and chaotic solutions are in fact dominated by periodicity, which can also be observed in the frequency




FIG. 4. The Galerkin neural network (GalNN) tailored for thermoacoustics. (a) Spatial branch, in which the solution space is spanned by the acoustic eigenfunctions (13) (hard constraint). The spatial modes $\Pi_j^{(1)}$, $\Pi_j^{(2)}$, $\Upsilon_j^{(1)}$, and $\Upsilon_j^{(2)}$, where $j = 1, 2, \dots, N_g$, are collected in the vectors $\Pi^{(1)}$, $\Pi^{(2)}$, $\Upsilon^{(1)}$, and $\Upsilon^{(2)}$, respectively. Step functions that model the discontinuity through the compact heat source (14) are shown at the bottom of the spatial branch. (b) Temporal branch, which is a periodically activated feedforward neural network. The outputs μ_j and η_j are stacked as vectors, which are projected onto the modes $\Pi^{(1)}$, $\Pi^{(2)}$, $\Upsilon^{(1)}$, and $\Upsilon^{(2)}$. The outputs η_{N_g+1} and η_{N_g+2} are associated with the discontinuity modes (14). The trainable parameters are the weights, $\mathbf{W}^{(c)}$, and the biases, $\mathbf{b}^{(c)}$. (c) Temporal branch of the periodic Galerkin neural network further constrained for limit-cycle solutions. It is reconstructed such that the weights are integer multiples of only one trainable fundamental angular frequency, θ . The trainable parameters are θ , the weights, $\mathbf{W}^{(2)}$, and the biases, $\mathbf{b}^{(2)}$.

spectrum of the oscillations (later in Sec. VB 2). Therefore, we deploy a periodically activated FNN (Sec. IV B 1) as the temporal branch of the GalNN.

4. Periodic Galerkin networks (P-GalNNs)

We enforce periodic activations (Sec. IV B 1) in the Galerkin neural network in order to achieve extrapolation for limit-cycle, quasiperiodic, and chaotic solutions. For simplicity, if we consider a single-layer periodically activated GalNN, then the outputs of the temporal branch of the network are given as a sum of sinusoids, where the hidden layer weights represent the angular frequencies and the biases represent the phases, i.e., $\sin(\mathbf{W}^{(1)}t + \mathbf{b}^{(1)})$. For a signal given as a sum of sinusoids,

its angular frequency is the greatest common divisor of the angular frequencies of its components. For a limit cycle, the greatest common divisor is its fundamental angular frequency, and the other frequencies are its harmonics. This applies to a periodically activated neural network as well. However, the weights of the neural network are initialized randomly and then trained on a finite amount of data, so it is numerically unlikely for the learned weights to be exact integer multiples of a fundamental frequency. To overcome this numerical challenge, we add a further constraint on the GalNN to guarantee periodic behavior with the fundamental frequency of the limit cycle. We create a hidden layer that takes time as input and outputs $(\sin(\mathbf{h}\theta), \cos(\mathbf{h}\theta))$, $\mathbf{h} = (1, 2, \dots, N_h)$, such that the weights in the periodic activations are integer multiples of one trainable variable, θ . The variable θ is initialized to the nondimensional angular frequency of the acoustic system, π , which is an educated guess for the actual thermoacoustic frequency as from our physical knowledge we know that the nonlinearity will only slightly change this. Hence, upon training, θ corresponds to the angular frequency of the limit cycle, and the number of harmonics, N_h , can be regarded as a hyperparameter. Because we have both sine and cosine activations for the same angular frequency, the phase information is automatically captured; hence we do not require a bias term. The temporal branch of the GalNN tailored for limit cycles is shown in Fig. 4(c).

V. EXTRAPOLATION AND STATE RECONSTRUCTION: TWIN EXPERIMENTS

A. Data set from the Rijke tube

We perform twin experiments on synthetic data generated from a prototypical thermoacoustic system, the Rijke tube, to demonstrate how the acoustic neural networks developed in Sec. IV B tackle extrapolation and state reconstruction. The Rijke tube (Fig. 3) captures the qualitative nonlinear dynamics and bifurcations observed in real-life applications [39]. In this model, we further assume that (i) the mean flow has a zero Mach number with uniform density, and (ii) the boundary conditions are ideal, i.e., $p'(x=0, t) = p'(x=1, t) = 0$, hence $R_{\text{in}} = R_{\text{out}} = -1$. We solve the associated PDEs (11) using the Galerkin decomposition [Eqs. (12) and (13)], with $\bar{\rho}_1 = \bar{\rho}_2 = 1$, which implies that $\omega_j = j\pi$ [57]. By substituting the pressure and velocity variables in Eqs. (11) with their Galerkin decompositions in Eqs. (12) and projecting the dynamics onto the Galerkin modes, the dynamics of the Galerkin variables η_j and μ_j are described by a $2N_g$ -dimensional system of ordinary differential equations [41]:

$$\dot{\eta}_j - \mu_j j\pi = 0, \quad (15a)$$

$$\dot{\mu}_j + \eta_j j\pi + \zeta_j \mu_j + 2\dot{q}' \sin(j\pi x_f) = 0. \quad (15b)$$

In the models of Rijke tubes in literature (e.g., [59]), a modal damping term acting on the pressure, $\zeta p'(x, t)$, is added to the energy equation (11b) to account for the unmodeled effects of dissipation at the boundaries, and at viscous and thermal boundary layers at the tube walls. This term lets the higher modes be damped out. In the projected dynamics, the damping term is $\zeta_j \mu_j$, where $\zeta_j = c_1 j^2 + c_2 j^{1/2}$ [60]. The heat-release rate is described by a modified King's law,

$$\dot{q}' = \beta(\sqrt{|1 + u'(x_f, t - \tau)|} - 1), \quad (16)$$

where β and τ are the heat-release strength and the flame time delay, respectively [61]. The time-delayed problem is transformed into an initial value problem via an advection function with the dummy variable v so that it can be solved with a standard time-marching scheme [41]:

$$\frac{\partial v}{\partial t} + \frac{1}{\tau} \frac{\partial v}{\partial X} = 0, \quad 0 \leq X \leq 1, \quad v(X=0, t) = u'(x_f, t). \quad (17)$$

The PDE (17) is discretized using a Chebyshev spectral method with N_c points [62].

The physics-informed loss of the Rijke tube system is calculated by evaluating the left-hand side of PDEs (11); i.e., the momentum residual, \mathcal{F}_M , is defined as the left-hand side of Eq. (11a), and the energy residual, \mathcal{F}_E , is defined as the left-hand side of Eq. (11b). When calculating the heat-release




FIG. 5. Extrapolation in time of the acoustic pressure with standard neural networks. Effect of activation functions: [(a), (b)] ReLU, [(c), (d)] tanh, [(e), (f)] sine in all hidden layers, and [(g), (h)] sine-ReLU (sine in the first layer, ReLU for the rest of the hidden layers). Top row: Predictions in both time and space. Bottom row: Predictions in time at $x = 0.25$, training data, and ground truth. Typical activation functions, such as ReLU [(a), (b)] and tanh [(c), (d)], fail to extrapolate in time.

term (16) in the physics-informed loss, the time-delayed velocity at the flame can be predicted from the network as $\hat{u}'(x_f, t_k - \tau)$ [for the GalNN $\hat{u}'(x_f, t_k - \tau) = \sum_{j=1}^{N_g} \hat{\eta}_j(t_k - \tau) \cos(j\pi x_f)$], i.e., by evaluating the network at $(x = x_f, t = t_k - \tau)$. Therefore, since the model takes time as an input for both FNN and GalNN architectures, there is no need to account for the time delay with a separate dummy variable. The evaluation of the heat-release law requires the system parameters x_f , β , and τ to be known.

For the FNN, the implementation details concerning the calculation of heat release and modal damping are provided in Appendix A. For the GalNN, as Galerkin amplitudes are available, the physical residuals of the Rijke tube can be directly computed from the ODEs (15) in these twin experiments.

B. Extrapolation in time

We simulate the Rijke tube system (11) with the parameters $N_g = 10$, $N_c = 10$, $x_f = 0.2$, $\beta = 5.7$, $\tau = 0.2$, $c_1 = 0.1$, and $c_2 = 0.06$. After a transient period, this system settles onto a limit cycle with a period of ≈ 1.927 time units. The data set is composed of pressure and velocity measurements from 13 sensor locations, which are uniformly distributed along the tube including the boundaries. Before training, the inputs are standardized to have a zero mean and unit variance.

1. Standard feedforward neural networks

Figures 5(a) and 5(b) and Figs. 5(c) and 5(d) show the results of FNNs with rectified linear unit (ReLU) and tanh activation functions when trained on this data set. The network hyperparameters (number of layers, number of neurons, learning rate) have been tuned via a grid search, in which we choose the architectures that resulted in the smallest validation losses. The network properties are provided in Table II. The ReLU network is initialized with the method of He *et al.* [63], and the tanh network with the method of Glorot and Bengio [64]. As illustrated in Figs. 5(a) and 5(b) and Figs. 5(c) and 5(d), FNNs with conventional activation functions fail at extrapolating periodic functions in time.

TABLE II. Summary of neural network model properties trained on the Rijke tube data.

	ReLU FNN	tanh FNN	sin FNN	sin-ReLU FNN	GalNN	P-GalNN
Hidden layers	5	5	3	2	2	1
Activations	ReLU	tanh	sine	sine-ReLU	sine	harmonics
Neurons	96	96	32	64	16	40
Learning rate	0.001	0.001	0.001	0.001	0.001	0.001
Optimizer				Adam		
Batch size				64		

Figure 5(e) and 5(f) show the extrapolation capability of an FNN equipped with sine activation, i.e., a periodically activated FNN (P-FNN) as introduced in Sec. IV B 1. (Note that when using the sine activation, the inputs are not normalized before training.) The activation function is modified to include a hyperparameter a such that it becomes

$$\phi(z) = \frac{1}{a} \sin(az). \quad (18)$$

This hyperparameter is used to fine-tune the frequency content of the learned functions. For this data set, it is determined as $a = 10$ with a grid search; the effect of varying a is shown in more detail in Appendix B. Physically, since the weights are initialized as $\sim \text{Uniform}(-1.22, 1.22)$ (Sec. IV B 1), multiplying this by $a = 10$ provides a good initial guess on the order of magnitude of the angular frequency for the training. Figures 5(g) and 5(h) show the case for the sine-ReLU configuration, which can also extrapolate in time. Both sine and sine-ReLU networks can extrapolate also at locations where no sensor measurements are available (result not shown). Deeper architectures are required when training with purely ReLU or tanh networks compared to sine or sine-ReLU networks. For example, in this case we use two layers of sine and sine-ReLU vs five layers of ReLU and tanh, and even then the ultimate training and validation losses of sine and sine-ReLU networks were one order of magnitude smaller. To conclude, using periodic activations produces more physical and expressive neural network models for acoustics than the conventional activations, which can also extrapolate in time, while speeding up training and hyperparameter search as a consequence of the smaller network sizes required.

2. Physics-constrained networks and long-term extrapolation

We showed that networks equipped with periodic activations make accurate predictions on data points chosen in a time range right after the training. Here, we analyze how the learned models extrapolate in the long term. The ideal model of a periodic solution should capture the periodicity that the system exhibits for all times. Network predictions over a long time period are plotted in Figs. 6(a) and 6(b), for P-FNN and GalNN, respectively, when trained on data from approximately two periods. Over a long time range, we observe a beatinglike phenomenon. In the frequency spectra of the acoustic variables, the peaks are at the harmonics of the angular frequency of the limit cycle, which we determined as $\theta^* = 3.2605$ by inspection of the frequency spectrum and autocorrelation of the time series. After the training, this is captured by the weights in the first hidden layer as discussed in Sec. IV B 4. These weights are multiplied with the time input before periodic activation, i.e., $\sin(a\mathbf{W}_1^{(1)}x + a\mathbf{W}_2^{(1)}t + a\mathbf{b}^{(1)})$ for P-FNN, and $\sin(a\mathbf{W}^{(1)}t + a\mathbf{b}^{(1)})$ for GalNN, where a is a hyperparameter of the sine activation [Eq. (18)]. The optimized weights are placed close to the harmonics of the angular frequency of the limit cycle. Figure 7 shows the (sorted) weights of the P-FNN and a GalNN. The harmonics of the angular frequency of the limit cycle are shown in dashed horizontal lines. Although these weights are close to the harmonics of the frequency of the limit cycle, they are not exact integer multiples of each other. Even the weights clustered around one frequency slightly differ from each other. This can be observed in the power spectral density (PSD)




FIG. 6. Long-term extrapolation in time of periodic acoustic pressure time series with physics-constrained neural networks. The training data are small and span the first four time units inside the red vertical line. Predictions on the pressure time series with (a) periodically activated feedforward neural network (P-FNN), (b) Galerkin neural network (GalNN), and (c) periodic Galerkin neural network (P-GalNN). Power spectral density of true and predicted time series with (d) P-FNN, (e) GalNN, and (f) P-GalNN. Constraining prior knowledge on periodicity is key to an accurate long-term prediction.




FIG. 7. In reference to Fig. 6, trained weights of the first hidden layers of (a) a periodically activated feedforward neural network (P-FNN), (b) a Galerkin neural network (GalNN), and (c) a periodic Galerkin neural network (P-GalNN). These weights are multiplied with the time input before periodic activation, i.e., $\sin(aW_1^{(1)}x + aW_2^{(1)}t + ab^{(1)})$ for P-FNN, $\sin(aW^{(1)}t + ab^{(1)})$ for GalNN, where a is a hyperparameter, and $\sin(W^{(1)}t)$ for periodic GalNN.

of the predicted time series as well, shown in Figs. 6(d) and 6(e). These observations align with the beating behavior. We overcome this issue by using the GalNN tailored for this specific purpose as described in Sec. IV B 4. The long-term predictions are shown in Fig. 6(c) and the power spectral density of the predicted signal in Fig. 6(f). The estimated θ^* by the periodic GalNN is 3.2615, which aligns with the true signal. The trained weights, which are activated as $(\sin(\mathbf{W}^{(1)}t), \cos(\mathbf{W}^{(1)}t))$ in the periodic GalNN, are shown in Fig. 7(c). Increasing the training time improves the extrapolation ability for both P-FNN and GalNN. However, with the constrained structure of the GalNN, we eliminate the beating behavior and achieve accurate long-term extrapolation without using more training data. In addition to limit cycles, thermoacoustic systems can also exhibit quasiperiodic and chaotic oscillations through bifurcations (e.g., [42]). With $x_f = 0.2$ and $\tau = 0.2$ and increasing β , the system takes a Ruelle-Takens-Newhouse route to chaos; at $\beta = 0.5$, it bifurcates from a fixed-point solution to a limit cycle, at $\beta = 5.8$ from a limit cycle to a quasiperiodic attractor, and at $\beta = 6.5$ from a quasiperiodic attractor to chaos [41]. We collect quasiperiodic data from the $\beta = 6$ regime and chaotic data from the $\beta = 7$ regime. For the quasiperiodic system, we train a GalNN with two sine layers of 96 neurons on data from the 80 time units long time series following a hyperparameter tuning with a grid search. The predictions are shown in Figs. 8(a)–8(c) with the PSD. The Galerkin network can capture the quasiperiodic behavior and the frequency spectrum correctly. For the chaotic system, we train a GalNN with four sine layers of 256 neurons on data from the 200 time units long time series. The predictions are shown in Figs. 8(d)–8(f) along with the PSD. To train on chaotic data, we required a longer time series and a deeper and wider network because the nonlinear dynamics is richer than the quasiperiodic solution. The network’s prediction correctly captures the dominant frequencies and the amplitude of the true signal.

C. Learning from partial state measurements

Full-state measurements are difficult to perform in thermoacoustics; i.e., only pressure measurements are taken with microphones. To emulate such a scenario, we perform the training using only pressure data, discarding the velocity data. Therefore, in the formulation of the data-driven loss [Eq. (4)], the measurement matrix becomes $\mathbf{M} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. Our aim is to learn a model that can also output velocity information without ever being exposed to measurements from it during training. Figure 9 shows the predictions of the P-FNN and the GalNN trained with and without physics information. When the loss is computed solely from the data, only the pressure can be learned. On the other hand, velocity predictions are made possible by enforcing the physics via the physics-informed loss term. For the purely data-driven networks, the output weights related to velocity are not updated; hence the velocity predictions are made with the initial weight values which are small. When both momentum and energy equations are imposed, both physics-informed networks can accurately reconstruct the velocity.

VI. RESULTS ON HIGHER-FIDELITY DATA

A. Higher-fidelity data

In order to assess how well the GalNN approach generalizes to higher-fidelity data, which is not generated by the Galerkin method, we employ a higher-fidelity model that includes the effects of the mean flow and has a kinematic flame model for the flame [33] (Fig. 3). The solution of the PDEs (7) can be obtained by a traveling-wave approach in conjunction with the jump conditions at the flame. This model can also account for nonideal open boundary conditions via acoustic reflection coefficients R_{in} and R_{out} at the inlet and outlet. We simulate a system with ideal boundary conditions, i.e., $R_{\text{in}} = R_{\text{out}} = -1$, that results in limit-cycle oscillations using the mean-flow parameters for the inlet velocity, $\tilde{u}_1 = 4$ m/s; for the inlet temperature, $\tilde{T}_1 = 293$ K; for the inlet and outlet pressure, $\tilde{P}_1 = \tilde{P}_2 = 101\,300$ Pa; and for the heat-release rate, $\tilde{Q} = 2000$ W. This is a low-Mach-number configuration with $\tilde{M}_1 = 0.0117$ and $\tilde{M}_2 = 0.0107$. The cross-section area changes at the flame location

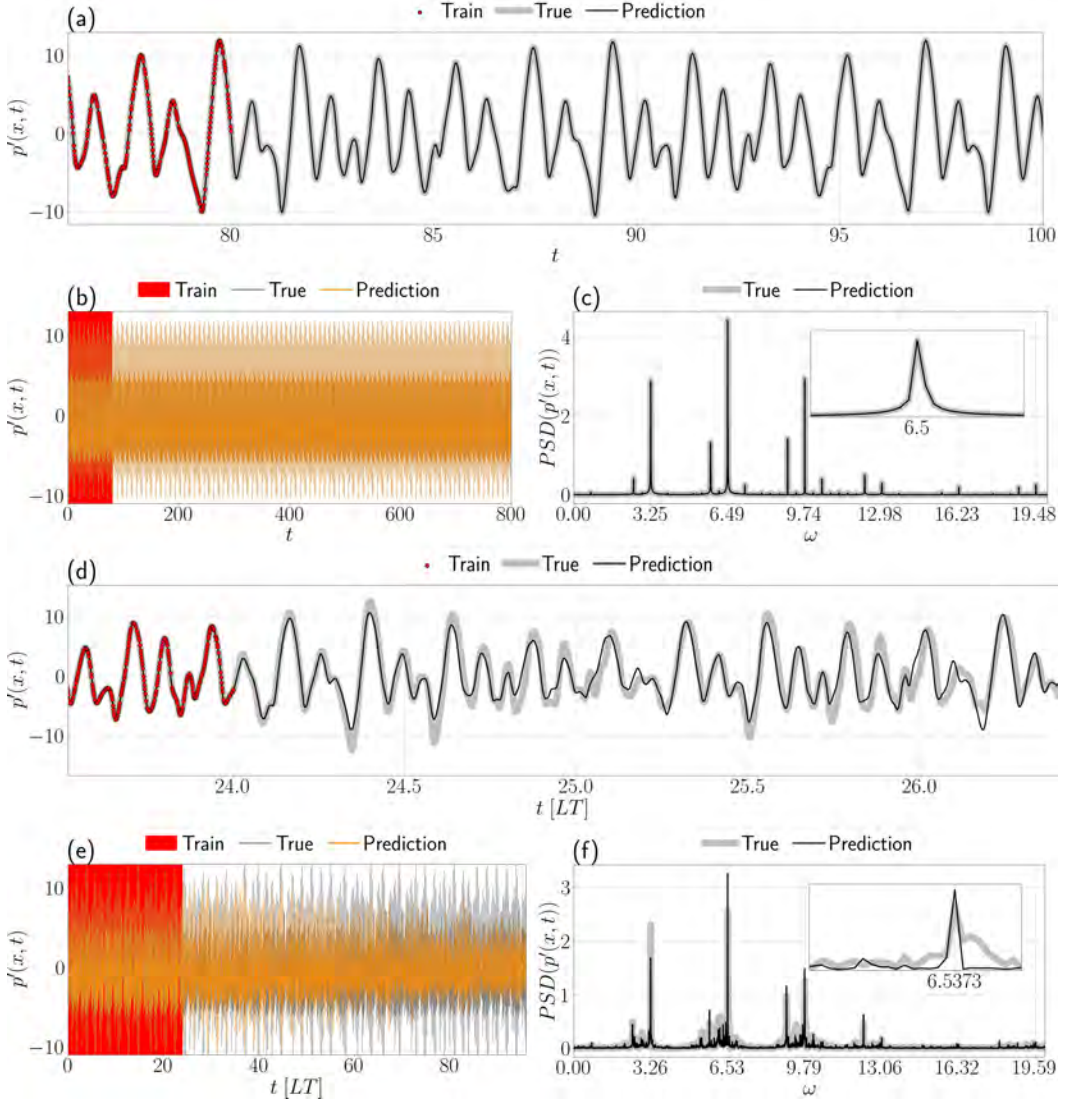


FIG. 8. Extrapolation in time for nonperiodic solutions. [(a)–(c)] Quasiperiodic solution and [(d)–(f)] chaotic solution. For the chaotic time series, the time axis is normalized by the Lyapunov time (LT) of the system, where $LT \approx 8.33$ time units in this case. Training data shown within the red vertical line [80 time units for the quasiperiodic case, 200 time units (24 LT) for the chaotic case].

with a ratio of $A_1/A_2 = 0.75$. In the preprocessing of the data, we nondimensionalize the variables as $x = \tilde{x}/\tilde{L}$, $t = \tilde{t}\tilde{c}_{\text{ref}}/\tilde{L}$, $u = \tilde{u}/\tilde{c}_{\text{ref}}$; $\rho = \tilde{\rho}/\tilde{\rho}_{\text{ref}}$; $p = \tilde{p}/(\tilde{\rho}_{\text{ref}}\tilde{c}_{\text{ref}}^2)$; and $q = \tilde{q}(\gamma - 1)/(\tilde{\rho}_{\text{ref}}\tilde{c}_{\text{ref}}^3)$, where $(\tilde{\cdot})_{\text{ref}} = (\tilde{L}_1(\tilde{\cdot})_1 + \tilde{L}_2(\tilde{\cdot})_2)/\tilde{L}$ is chosen as the weighted average of the mean values of variable (\cdot) across the flame. The nondimensional flame location is at $x_f = 0.61$, and the nondimensional densities in the duct segments before and after the flame are $\tilde{\rho}_1 = 1.145$ and $\tilde{\rho}_2 = 0.769$. These quantities will be used to determine the acoustic mode shapes for the Galerkin neural network (13). While creating the training and validation sets, we consider two types of validation: (i) validation of the interpolation capability and (ii) validation of the extrapolation capability, i.e., when the data are sampled from outside the time range of training. A good validation of the




FIG. 9. Reconstruction of acoustic velocity from pressure measurements only. Predictions on acoustic [(a), (b)] pressure and [(c), (d)] velocity from only 13 pressure observations with different physics-constrained networks. Top row: Spatial shapes at $t = 1.5$. Bottom row: Time series at $x = 0.25$. The acoustic velocity is accurately reconstructed from pressure measurements with the physics-informed periodic feedforward and Galerkin neural networks.

interpolation loss prevents overfitting. Validation of extrapolation loss indicates how well the model generalizes beyond the time range of training. After discarding a transient, we split the first four nondimensional time units long time series (approximately two periods) of the limit cycle into training and validation of the interpolation sets, assigning an 80% to 20% ratio between them. The validation data for the extrapolation is then taken as the next four nondimensional time units. The physics-informed loss of the higher-fidelity system is calculated by evaluating the left-hand side of PDEs (7) after their being nondimensionalized as described above. The momentum residual, \mathcal{F}_M , is defined as the left-hand side of Eq. (7a). In contrast to the twin experiments with the Rijke tube, for the higher-fidelity data, we assume that we do not have access to the heat-release model. Hence, the energy residual, \mathcal{F}_E , is defined as the left-hand side of Eq. (7b) excluding the heat-release term, \dot{q} . This is acceptable because the heat-release term only acts at the flame location due to the compact assumption and the Dirac delta.

B. Comparison of networks performance for reconstruction

We analyze the performance of the different networks that we have built in Sec. IV B on the reconstruction of the higher-fidelity data. We compare three configurations of the feedforward neural network with (i) ReLU activations in all hidden layers, (ii) sine activation in all hidden layers, and (iii) sine activation in the first layer, ReLU activation in the rest of the hidden layers, and three configurations of the periodic Galerkin neural network with one hidden trainable frequency layer, (i) with $\omega_j = j\pi$ as the acoustic angular frequencies as before, i.e., when $\bar{\rho}_1 = \bar{\rho}_2 = 1$, (ii) with ω_j determined from the dispersion relationship using the real mean-flow density values of the system,

TABLE III. Summary of neural network model properties trained on higher-fidelity data. Feedforward neural networks (FNNs) with different activations: ReLU, sine, and sine-ReLU (sine in the first layer, ReLU in the rest). Periodic Galerkin neural networks (P-GalNNs) with different choices of spatial bases: (i) natural acoustic modes (13) for $\bar{\rho}_1 = \bar{\rho}_2 = 1$, (ii) natural acoustic modes for $\bar{\rho}_1 = 1.145$ and $\bar{\rho}_2 = 0.769$, and (iii) case (ii) with the addition of step functions in the velocity modes to handle the jump discontinuity (14).

	ReLU FNN	sin FNN	sin-ReLU FNN	P-GalNN (i)	P-GalNN (ii)	P-GalNN (iii)
No. hidden layers	6	3	4	1	1	1
No. Galerkin modes				20	20	20
Activations	ReLU	sine	sine-ReLU	harmonics	harmonics	harmonics
No. neurons	64	64	32	20	20	20
Learning rate	0.01	0.01	0.01	0.0004	0.0004	0.0004
Optimizer				Adam		
Batch size				32		

i.e., when $\bar{\rho}_1 = 1.145$ and $\bar{\rho}_2 = 0.769$, and (iii) the Galerkin modes chosen as (ii) with the addition of step functions in the velocity modes to handle the jump discontinuity (14).

For each configuration of activations, the hyperparameters of the architecture and training (number of layers, number of neurons, and learning rate) have been tuned with a grid search, in which we chose the set of hyperparameters that resulted in the smallest validation losses. These properties are provided in Table III for each network. The training and validation loss histories of all models are shown in Fig. 10. The ReLU network has higher training and interpolation losses, cannot fit the data as well as the sine and sine-ReLU networks, and cannot extrapolate. For all the other networks [sine FNN, sine-ReLU FNN, P-GalNN (i), P-GalNN (ii), and P-GalNN (iii)], the validation losses are close to the training loss. The P-GalNN (iii) scores the lowest losses. The periodic Galerkin networks can have similar or better prediction performance as the other deeper feedforward neural networks even with one hidden layer following a physical choice of spatial basis. Figure 11 visualizes the velocity predictions and their errors between the predictions and the true data obtained with the networks. The P-GalNN (i) exhibits an error pattern that resembles an offset from the true solution, which indicates a biased model. The prediction is also oscillatory and the error is especially high at the boundaries. The more physically motivated choice of P-GalNN (ii) eliminates these errors; however, the discontinuity jump still cannot be captured, which leads to high-frequency oscillations in the prediction around the flame location. A similar phenomenon is also observed for the sine-ReLU network, where the error is the highest at the flame location. This problem is overcome by P-GalNN (iii), which results in the smallest training and validation errors overall. This network finds the nondimensional dominant angular frequency as $\theta^* = 3.3164$ and, thus, the period as $T^* = 1.8946$. From the power spectral density of the true pressure signal, we read $\theta^* = 3.3166$ and $T^* = 1.8945$. The periodic Galerkin network successfully learns the correct period.

C. Robustness to sparsity and noise

In this section, we demonstrate the robustness of the acoustic neural networks to sparsity and noise in the training data. We select P-GalNN (iii) from the previous section because of its more generalizable performance in comparison to the other architectures and investigate the regularization methods that can be used in order to prevent overfitting in the case of corrupted or scarce data. We compare a physics-based regularization, i.e., physics-informed loss, with ℓ_1 - and ℓ_2 -norm regularizations, which act on the network weights as $\|\mathbf{W}\|_1$ and $\|\mathbf{W}\|_2$, respectively, in the loss term. We add zero-mean Gaussian noise on the nondimensional training data with a standard deviation equal to 10% of the standard deviation of the true solution function over the whole domain (20 dB signal-to-noise ratio). Using the described regularization methods, we train networks with

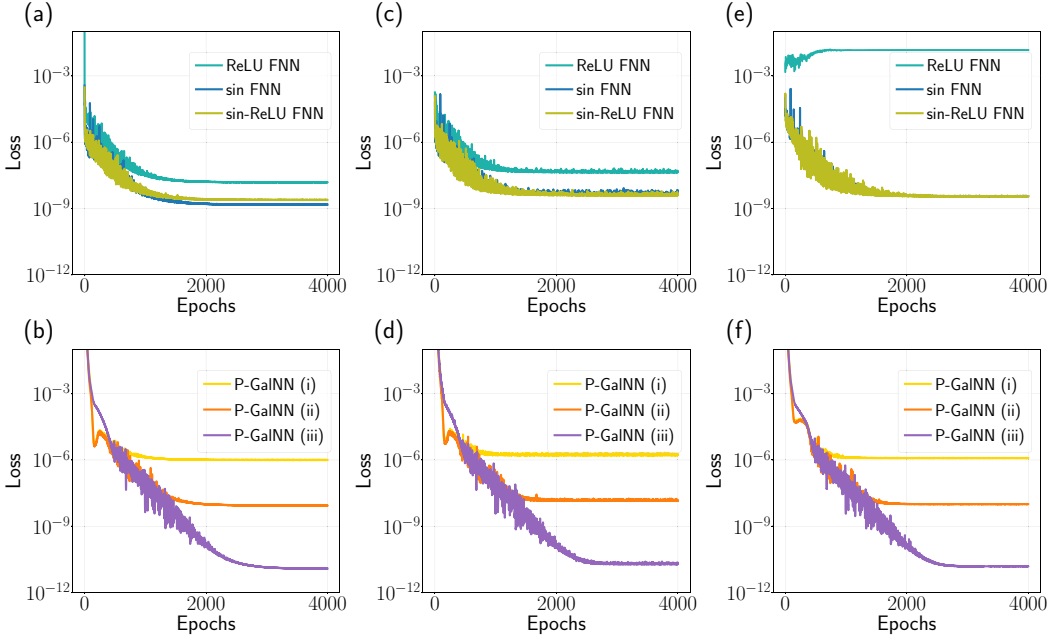


FIG. 10. Data from higher-fidelity model for a periodic acoustic solution. Loss-function histories. [(a), (b)] Training, [(c), (d)] validation of interpolation, and [(e), (f)] validation of extrapolation. Feedforward neural networks (FNNs) with different activations: ReLU, sine, and sine-ReLU (sine in the first layer, ReLU in the rest). Period Galerkin neural networks (P-GalNNs) with different choices of spatial bases: (i) natural acoustic modes (13) for $\bar{\rho}_1 = \bar{\rho}_2 = 1$, (ii) natural acoustic modes for $\bar{\rho}_1 = 1.145$ and $\bar{\rho}_2 = 0.769$, and (iii) case (ii) with the addition of step functions in the velocity modes to handle the jump discontinuity (14).

ten Galerkin modes on this new data set, this time given only on a coarse grid. The comparison of the obtained predictions when using (i) no regularization, (ii) ℓ_2 -norm, (iii) ℓ_1 -norm, and (iv) physics-based regularization is shown in Fig. 12. We train multiple networks with varying values of regularization coefficients and show only the results for those with the smallest validation losses. Without regularization, the model can fit the training data perfectly; however, the interpolation points show spikes, which is a clear demonstration of overfitting. This case of overfitting emerges when there is a discrepancy between the number of sensors and the number of Galerkin modes dictated by Nyquist-Shannon sampling theorem. The highest number of Galerkin modes is given as

$$N_g^* = \arg \max_j \{j \mid \omega_j \sqrt{\bar{\rho}_1} \leq k_{Nyquist}, \omega_j \sqrt{\bar{\rho}_2} \leq k_{Nyquist}\}, \quad (19)$$

where $k_{Nyquist} = \frac{\pi}{\Delta x}$ with Δx being the spacing between sensor locations. The physics information significantly eliminates overfitting and outperforms ℓ_1 - and ℓ_2 -norm regularizations. The ℓ_2 -norm regularization does not work as well because the weights are not selectively regularized, i.e., dominant modes or frequencies will be regularized in the same way as nondominant ones. The ℓ_1 -norm regularization performs better since modes corresponding to higher wavenumbers are discarded as the ℓ_1 norm promotes sparsity and in the physical basis, the energy is contained in the low-frequency modes. We perform a quantitative study to investigate the robustness of the prediction performance to sparsity and noise. First, we vary the number of sensors along the tube. Figure 13 shows the relative ℓ_2 error achieved by Galerkin networks with different number of modes when trained on noise-free data given on spatial grids of 2, 5, 11, and 23 points. The relative ℓ_2 error is calculated from the prediction on the fine grid of 49 points over the training range by dividing the ℓ_2 norm of the error by the ℓ_2 norm of the ground truth. The networks are trained with and without

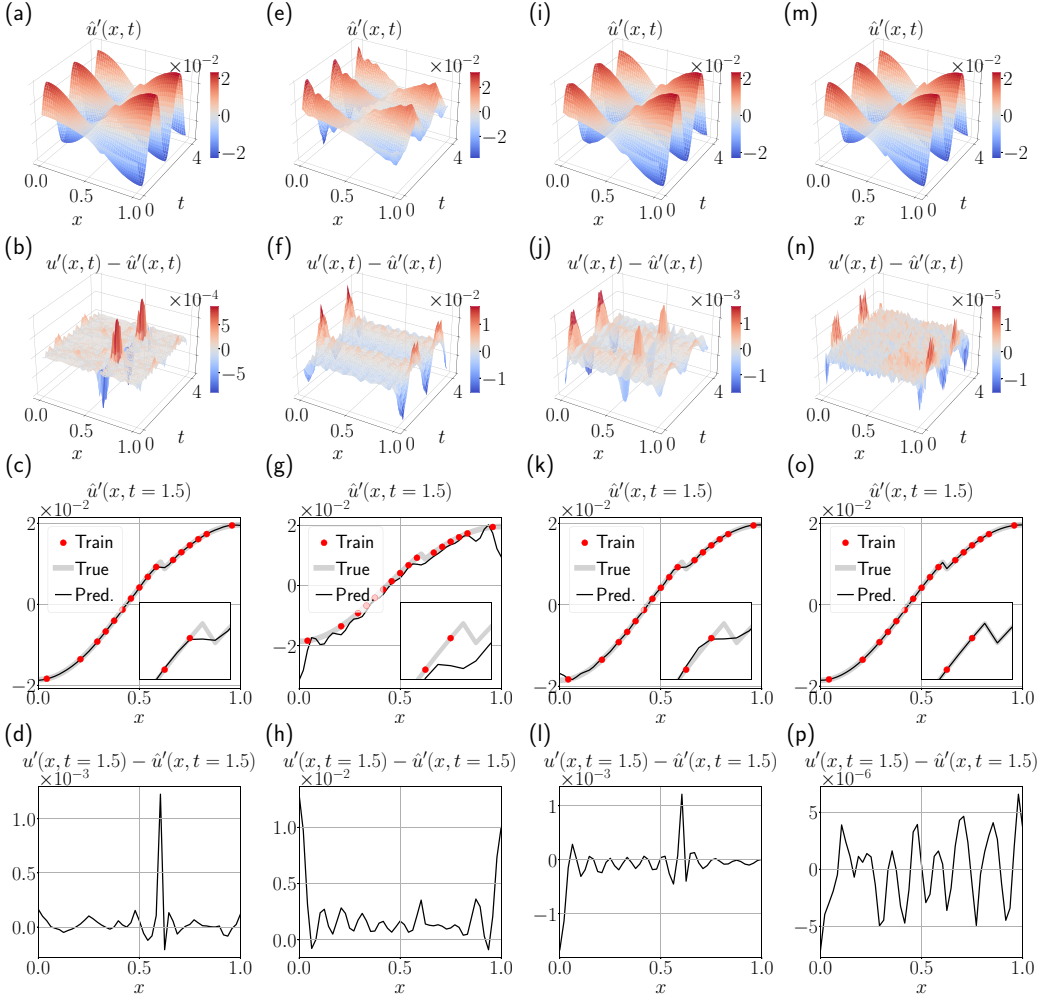


FIG. 11. Comparison of acoustic neural networks in terms of reconstruction performance from higher-fidelity data. From left to right, [(a)–(d)] sine-ReLU feedforward neural network (FNN), i.e., sine in the first layer, ReLU in the rest; [(e)–(h)] periodic Galerkin neural network (P-GalNN) (i) with natural acoustic modes (13) for $\bar{\rho}_1 = \bar{\rho}_2 = 1$; [(i)–(l)] P-GalNN (ii) with natural acoustic modes for $\bar{\rho}_1 = 1.145$ and $\bar{\rho}_2 = 0.769$; [(m)–(p)] P-GalNN (ii) with the addition of step functions in the velocity modes to handle the jump discontinuity (14). From top to bottom, velocity prediction in time and space, error between ground truth and prediction in time and space, velocity prediction and error at a fixed time instance, $t = 1.5$.

physics information in the loss function. We fix the weighting of the data-driven loss to $\lambda_D = 1$, and the physics loss to $\lambda_M = \lambda_E = 0.0001$ for physics-informed and to zero otherwise. As discussed above, a high number of Galerkin modes leads to overfitting for coarse grids, when the Nyquist wavenumber is not sufficient to resolve the high wavenumbers, e.g., five sensors and ten Galerkin modes. Furthermore, we observe large prediction error even when the Nyquist condition is satisfied, e.g., five sensors and five Galerkin modes. We find that this type of error is mostly concentrated at the boundaries of the velocity, as we have not provided any boundary data and, therefore, the model overfits to the rest of the data, not generalizing well over the boundaries. For pressure, this is not an issue, because the boundary conditions are Dirichlet and already encapsulated within the provided Galerkin basis. On the other hand, the low number of Galerkin modes may not be enough

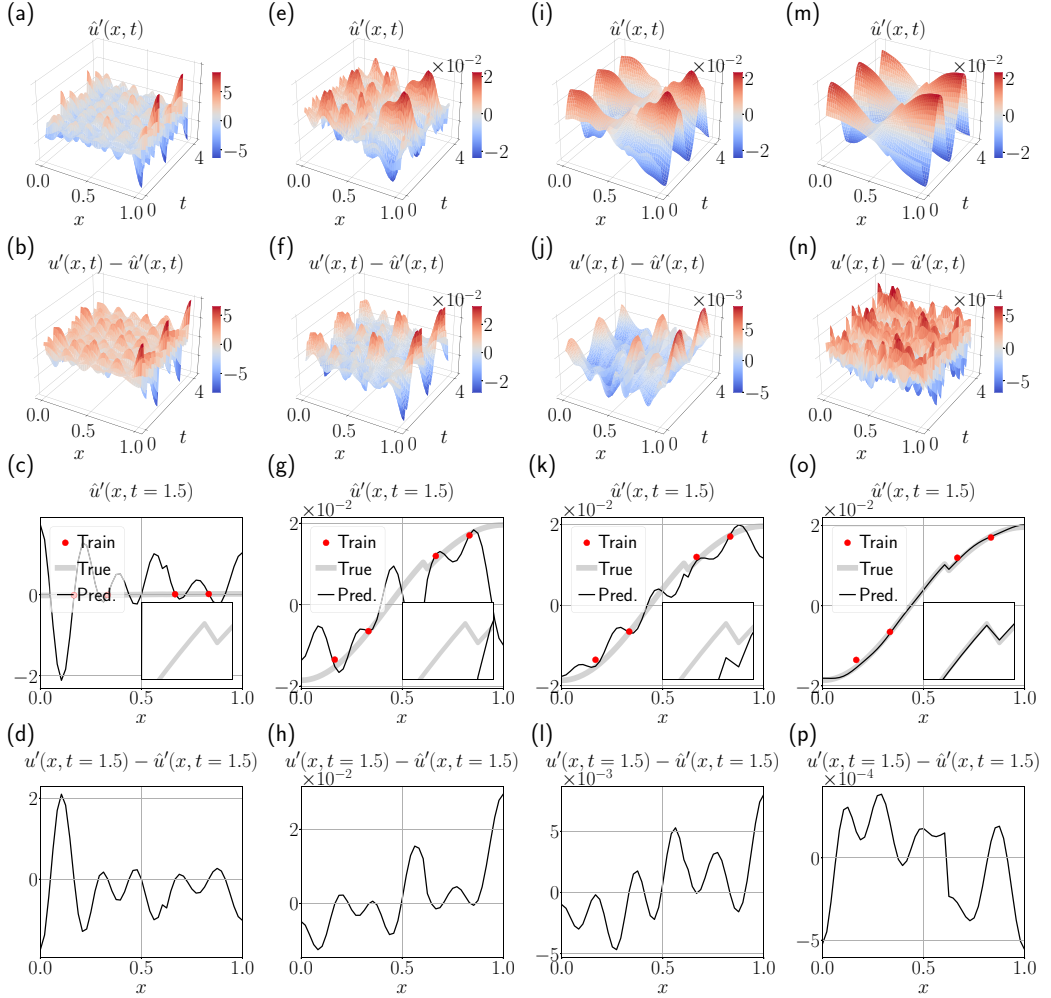


FIG. 12. In reference to Fig. 11, effect of the regularization hyperparameter on periodic Galerkin neural networks (P-GalNNs) in the case of sparse and noisy measurements. From left to right, [(a)–(d)] no regularization, [(e)–(h)] ℓ_2 -norm regularization, [(i)–(l)] ℓ_1 -norm regularization, and [(m)–(p)] physics-based regularization. From top to bottom: acoustic velocity, error between ground truth and prediction, velocity prediction at a fixed time instance, and error at a fixed time instance, $t = 1.5$.

to approximate the pressure and velocity shapes in fine detail, e.g., 11 sensors and 5 Galerkin modes. The addition of the physics information shows a marked difference and helps in overcoming these limitations. In Appendix E, we show the reconstruction errors associated with physics-informed training with varying numbers of pressure sensors.

Next, we choose the network with ten modes and add increasing levels of noise to the data from five sensor measurements. We take five different realizations of noise and show the prediction relative ℓ_2 error as a box plot in Fig. 14 for noise levels of 5%, 10%, 20%, and 40% (26, 20, 14, and 8 dB signal-to-noise ratios). Noting that, even in the noise-free case, the non-physics-informed network performs poorly, we focus on physics-informed training with varying weightings of the physics loss. As the noise level increases, the optimum weighting of the physics information increases as well, since the data become less reliable. We obtain good predictions of pressure and




FIG. 13. Effect of number of pressure sensors in data collection. Effect of number of modes in Galerkin neural networks, with and without physics-based loss. (The regularization hyperparameter of the data-driven loss is $\lambda_D = 1$.)

velocity over the entire domain, even though the data are sparse and noisy. The network is also robust to different realizations of noise.

We conclude by demonstrating how the physics-informed training performs in the absence of velocity data. The prediction results of a physics-informed periodic Galerkin neural network, PI-P-GalNN, are compared with a sine-ReLU physics-informed feedforward neural network, PI-P-FNN, in Fig. 15. For the PI-P-GalNN, we report the following relative ℓ_2 errors: 2.63% for pressure and 2.75% for velocity in the training range ($t = 0-4$ time units), and 2.84% for pressure and 2.80% for velocity in the extrapolation range ($t = 4-8$ time units). In comparison, for the sine-ReLU feedforward neural network, we report the following relative ℓ_2 errors: 4.81% for pressure and 164.22% for velocity in the training range ($t = 0-4$ time units), and 9.77% for pressure and 160.57% for velocity in the extrapolation range ($t = 4-8$ time units). With the PI-P-GalNN, we can reconstruct the velocity as well as pressure. The PI-P-FNN fits to the noisy data more compared to




FIG. 14. Robustness to noise. Galerkin neural networks with different regularization hyperparameters for physics-based losses with data corrupted with noise. (The regularization hyperparameter of the data-driven loss is $\lambda_D = 1$.)




FIG. 15. Reconstruction of acoustic velocity from noisy pressure measurements only. Predictions on acoustic [(a), (b)] pressure and [(c), (d)] velocity from only five pressure observations with physics-based loss using the periodically activated feedforward neural network (P-FNN) and periodic Galerkin neural network (P-GalNN). Top row: Pressure and velocity shapes along the tube at a fixed time instance, $t = 1.5$. Bottom row: Pressure and velocity time series at a fixed point in the tube, $x = 0.33$.

the PI-P-GalNN, and can only recover the velocity with an offset, completely discarding the velocity jump. A constant offset of the velocity, $\hat{u}'(x, t) = u'(x, t) + c$, satisfies the governing equations (7). Because the partial derivatives of the constant offset will be zero, the residual of the governing equations will not be affected. If the velocity fluctuations are known to have a zero mean, one strategy could be to remove the mean of the predictions in the postprocessing, and a second strategy could be to remove the biases from the output layer of the neural network. However, even when the offset is removed, the PI-P-FNN cannot capture the jump of the velocity at the flame location.

D. Nonideal boundary conditions

So far, we have restricted ourselves to ideal boundary conditions. However, in real experiments, we deal with nonideal conditions, in which the acoustic waves are not fully reflected at the tube ends (e.g., [65]). As a result, in contrast to the ideal case and the acoustic modes we have obtained under that assumption, the pressure fluctuations at the tube ends are not zero, i.e., $p'(x = 0, t) \neq 0$, $p'(x = 1, t) \neq 0$. In order to tackle nonideal boundary conditions, we add a linear term in the pressure modes (13a),

$$\Pi_{N_g+1}^{(1)}(x) = \Pi_{N_g+1}^{(2)}(x) = x, \quad (20a)$$

$$\Pi_{N_g+2}^{(1)}(x) = \Pi_{N_g+2}^{(2)}(x) = 1, \quad (20b)$$

such that the summation (12a) runs from 1 to $N_g + 2$. The modes Π_{N_g+1} and Π_{N_g+2} are weighted by independent coefficients, μ_{N_g+1} and μ_{N_g+2} . This way, the linear term provides the offset to the

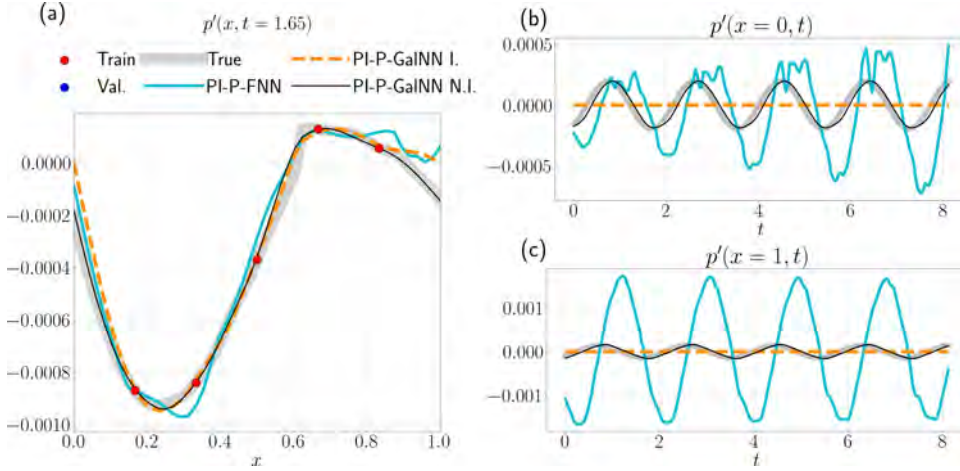


FIG. 16. Prediction in the case of nonideal ($R_{\text{in}} = R_{\text{out}} = -0.985$) boundary conditions for a training case of only five noise-free pressure sensors. Comparison of a standard physics-informed periodic Galerkin neural (PI-P-GalNN I.), physics-informed periodic feedforward network (PI-P-FNN) with a PI-P-GalNN for nonideal boundary conditions (PI-P-GalNN N.I.). (a) Pressure shape at a fixed time instance, $t = 1.65$, and [(b), (c)] pressure fluctuations at $x = 0$ and $x = 1$, respectively. The linear term can capture the nonzero pressure fluctuations at the boundaries, when they are close to ideal.

predicted pressure fluctuations at the tube ends, i.e., $\mu_{N_g+1}(t)x + \mu_{N_g+2}(t)$. Once trained, we expect to estimate $\mu_{N_g+2}(t) = p'(x = 0, t)$ and $\mu_{N_g+1}(t) = p'(x = 1, t) - p'(x = 0, t)$.

We consider three cases and generate data from the higher-fidelity model using the following parameters: (1) close to ideal, $\tilde{u}_1 = 8$ m/s, $\tilde{Q} = 10\,000$ W, $R_{\text{in}} = R_{\text{out}} = -0.985$; (2) less




FIG. 17. Prediction in the case of nonideal ($R_{\text{in}} = R_{\text{out}} = -0.85$) boundary conditions for a training case of only five noise-free pressure sensors. Comparison of a standard physics-informed periodic Galerkin neural (PI-P-GalNN I.), physics-informed periodic feedforward network (PI-P-FNN) with a PI-P-GalNN for nonideal boundary conditions (PI-P-GalNN N.I.). (a) Pressure shape at a fixed time instance, $t = 1.90$, and [(b), (c)] pressure fluctuations at $x = 0$ and $x = 1$, respectively. The linear term can capture the nonzero pressure fluctuations at the boundaries, when they are less close to ideal.

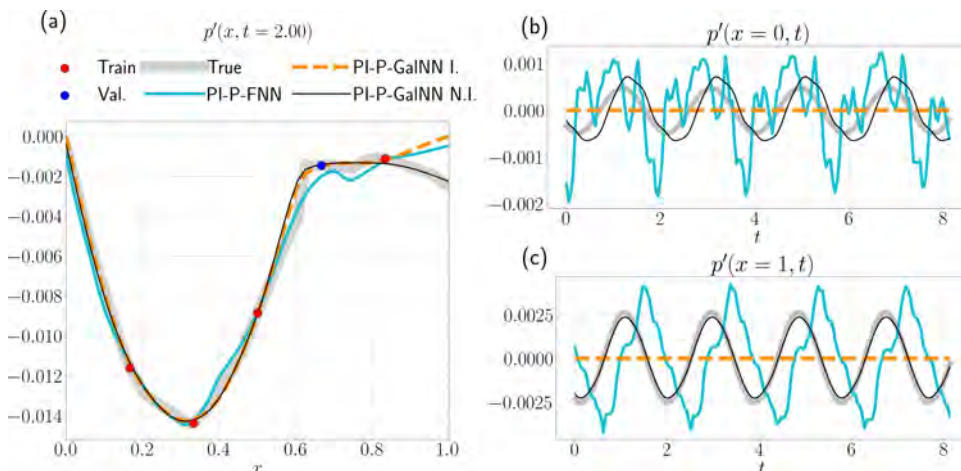


FIG. 18. Prediction in the case of nonideal boundary conditions ($R_{\text{in}} = -0.985$, $R_{\text{out}} = -0.85$) for a training case of only five noise-free pressure sensors. Comparison of a standard physics-informed periodic Galerkin neural (PI-P-GalNN I.), physics-informed periodic feedforward network (PI-P-FNN) with a PI-P-GalNN for nonideal boundary conditions (PI-P-GalNN N.I.). (a) Pressure shape at a fixed time instance, $t = 2.00$, and [(b), (c)] pressure fluctuations at $x = 0$ and $x = 1$, respectively. The linear term can capture the nonzero pressure fluctuations at the boundaries, even when the reflection coefficients are different.

close to ideal, $\tilde{u}_1 = 16$ m/s, $\tilde{Q} = 60\,000$ W, $R_{\text{in}} = R_{\text{out}} = -0.85$; and (3) different at both ends, $\tilde{u}_1 = 16$ m/s, $\tilde{Q} = 40\,000$ W, $R_{\text{in}} = -0.985$, $R_{\text{out}} = -0.85$.

The nondimensional densities in the duct segments before and after the flame are, for case (1), $\bar{\rho}_1 = 1.269$, $\bar{\rho}_2 = 0.571$, for case (2), $\bar{\rho}_1 = 1.434$, $\bar{\rho}_2 = 0.307$, and for case (3), $\bar{\rho}_1 = 1.376$, $\bar{\rho}_2 = 0.400$.

Figures 16–18 compare the predictions of a PI-P-GalNN (15 Galerkin modes) with a linear term against the standard PI-P-GalNN and the PI-P-FNN for the above cases, where the training data consist of only five pressure sensors. With the addition of the linear term, the PI-P-GalNN is capable of capturing the nonzero pressure fluctuations at the boundaries for near-ideal conditions even from sparse measurements given at other locations. While PI-P-FNN is a flexible tool that can handle various boundary conditions, it requires the boundary data [50], whereas the PI-P-GalNN does not. (Training computational costs are provided in Appendix D.)

VII. CONCLUSIONS

In this work, we model acoustic and thermoacoustic pressure and velocity oscillations from synthetic data. The synthetic data capture the rich nonlinear behavior of thermoacoustic oscillations observed in propulsion and power generation. We develop acoustic neural networks to tackle the tasks of (i) extrapolation in space and time and (ii) reconstruction of full acoustic state from partial observations by exploiting prior knowledge on the physics of acoustics. The prior knowledge is embedded in the network as both soft and hard constraints. First, as acoustic and thermoacoustic systems are dominated by sinusoidal eigenfunctions, we promote a physically motivated choice of sinusoidal activation functions. Unlike standard feedforward neural network architectures that employ ReLU or tanh activations, periodically activated networks can extrapolate in time, and the trained weights hold frequency information. This means that spatiotemporal patterns can be learned more efficiently from fewer data because the network is more expressive as well as robust. Second, we inform the training with the acoustic conservation laws with penalty terms in the loss function. This term regularizes the predictions in the case of noisy and scarce data, and enables reconstruction

of unobserved states. Typically, in thermoacoustics experiments, only pressure measurements are available. Third, inspired by Galerkin decomposition, we design a neural network with temporal and spatial branches (Galerkin neural network), which spans a physical function space via the choice of acoustic eigenmodes as the spatial basis. In order to account for nonideal boundary conditions that occur when the acoustic waves are not fully reflected at the boundaries, we add a linear mode in the pressure modes, which captures the nonzero pressure fluctuations at the inlet and outlet for near-ideal conditions without requiring data from the boundaries. We consider two test cases: (i) twin experiments on synthetic data from a Rijke tube with a nonlinear, time-delayed heat-release model with Galerkin discretization of the PDEs, and (ii) higher-fidelity data generated by a thermoacoustic network model with a kinematic flame model that takes into account mean-flow effects. The first case shows that the standard feedforward neural networks fail at extrapolation, while periodically activated networks can extrapolate and reconstruct accurately from partial measurements. On this test case, we also demonstrate the long-term extrapolation capability of Galerkin neural networks on periodic and quasiperiodic solutions, while for chaotic solutions, we recover the dominant components in the frequency spectrum. In the second test case, we show the generalizability and robustness of the Galerkin neural networks to higher-fidelity data, which can be noisy and contain only pressure measurements. Hard-constraining neural networks with prior knowledge reduce the required network size, search space of hyperparameters, as well as the amount of data required for training. This work opens up possibilities to learn the nonlinear dynamics of thermoacoustics using physics-aware data-driven methods. Future research directions include testing the method for geometries with cross-section changes, in which the predetermined spatial branch of the Galerkin neural network can be extended with more trainable parameters.


ACKNOWLEDGMENT

This research has received financial support from the ERC Starting Grant No. PhyCo 949388. We gratefully acknowledge an anonymous reviewer for their spot-on comments, which helped improve the paper.

APPENDIX A: CALCULATION OF MODAL DAMPING AND HEAT-RELEASE TERMS IN THE PHYSICS-INFORMED LOSS

1. Modal damping

The modal damping is defined as a multiplication in the frequency domain and a convolution in the spatial domain with the pressure. Using an FNN, we predict the pressure at a given time in the spatial domain. In order to compute the contribution of the damping in the physics-informed loss, at each training step, we will transform the pressure predicted in the spatial domain to the frequency domain via fast Fourier transform (FFT), compute the damping in the frequency domain, and then transform the result back to the spatial domain. When sampling from signals, the distinct Fourier frequencies are given as $\omega_k = k \frac{2\pi}{KT_s}$, $k = 0, 1, \dots, K/2$, where K is total number of samples, T_s is the sampling time, and KT_s gives the length of the sampled signal. During simulation, we considered N_g Galerkin modes. So, in our case, we set $K \geq 2N_g$. Notice that the spatial domain is $[0,1]$, while the wavenumbers of the Galerkin modes have a resolution of π , which is only half period in this domain. Consequently, the resolution of the Fourier frequencies must be set to $\frac{2\pi}{KT_s} = \pi$ and, hence, $KT_s = 2$, which means that the FFT must be taken over a domain of $[0,2]$, such that the samples are at $x_k = k \frac{2}{K}$, $k = 0, 1, \dots, K - 1$, and we have one period of discrete samples. One could predict the pressure over the $[0,2]$ domain, but this could lead to errors as there are no training data in this region. Instead, we first predict over the original $[0,1]$ domain and then stack this prediction with its negative symmetric with respect $x = 1$, since we know that pressure is given as a sum of sines,


 FIG. 19. Effect of hyperparameter a in the sine activation $\frac{1}{a} \sin(az)$.

which is an odd function. So, we have

$$\hat{p}' = \left\{ \hat{p}'(x_k, t) \Big|_{x_k = k \frac{2}{K}}, k = 0, 1, \dots, K-1, K \geq 2N_g \right\}, \quad (\text{A1})$$

and its Fourier transform $\mathcal{F}(\hat{p}')$. Now, we take only the part of $\mathcal{F}(\hat{p}')$ that corresponds to the positive frequencies and elementwise multiply it with the damping modes ζ ,

$$\xi_j = \mathcal{F}_j^+(\hat{p}') \zeta_j, \quad j = 1, 2, \dots, N_g. \quad (\text{A2})$$

In the next step, we take the inverse Fourier transform of the convolution. As a matter of fact, the training is done in batches, so the x locations we are interested in calculating the effect of damping at may not necessarily collide with the spatial grid that we have previously constructed. So, we will do the inverse Fourier transform for these locations in the training batch. Ultimately, the damping term is found as

$$\zeta p'(x) = \frac{1}{2K} \sum_{j=1}^{N_g} \xi_j e^{ij\pi x} + \xi_j^* e^{-ij\pi x}, \quad (\text{A3})$$

where ξ_j^* denotes the complex conjugate of ξ_j . We use the contribution of damping in the residual of the energy equation (11b) when computing the physics-informed loss (10).

2. Heat release

Although the heat-release term acts as a Dirac delta in the spatial domain, the solver implements Galerkin decomposition, which projects this term onto a truncated set of modes. Plugging back the decomposition (12) in the left-hand side of the energy equation (11b), the remaining heat-release

TABLE IV. Computation time of twin experiments with 13 pressure sensors. Training with a batch size of 32 for 2000 epochs (on Intel Xeon(R) Gold 5218R CPU at 2.10 GHz \times 80).

Network name	Description	Training time
PI-P-GalNN	10 Galerkin modes, 1 harmonic layer with 40 neurons	1 h, 50 min
PI-GalNN	10 Galerkin modes, 2 layers with 16 neurons	1 h, 50 min
PI-P-FNN	3 layers with 32 neurons	2 h, 40 min

term is equal to $\sum_{j=1}^{N_g} 2q'_j \sin(j\pi x_f) \sin(j\pi x)$, which is the effective term in the simulations. As the number of Galerkin modes approaches to infinity, this approaches the Dirac delta. However, from a practical point of view, the number of modes in the simulation is finite, thus we use this expression to find the contribution of heat release in the residual of the energy equation (11b) when computing the physics-informed loss (10).

APPENDIX B: EFFECT OF HYPERPARAMETER IN THE SINE ACTIVATION

We observe that varying hyperparameter a in the sine activation formulated as $\frac{1}{a} \sin(az)$ affects the frequency of the learned functions. Hence, it is a hyperparameter that requires tuning. Figure 19 illustrates this effect for the Rijke tube data discussed in Sec. VB for $a = 1, 10, 20$. Low a , $a = 1$, results in a low-frequency model, whereas high, $a = 20$, results in a high-frequency model, which can especially be observed in the pressure and velocity shapes in the spatial domain. We found the optimum value for this data set as $a = 10$. Since the sampling frequency of the training data is high enough, we do not observe the effect of high a in the time domain. Ziyin *et al.* [55] reported similar findings in their studies when using the $1 + \sin^2(z)$ activation.

APPENDIX C: TRAVELING-WAVE SOLUTION TO THE HIGHER FIDELITY MODEL

In this model, the pressure and velocity are expressed as functions of two acoustic traveling waves, which propagate up- and downstream of the tube. These waves are derived by applying the method of characteristics to the acoustic wave equation and defined as f and g , with propagation velocities $\tilde{c}_1 \pm \tilde{u}_1$ in the upstream region $\tilde{x} \leq \tilde{x}_f$, and h and j , with propagation velocities $\tilde{c}_2 \pm \tilde{u}_2$ in the downstream region $\tilde{x} \geq \tilde{x}_f$. This model can also account for nonideal open boundary conditions via reflection coefficients R_{in} and R_{out} at the inlet and outlet, $f(\tilde{t}) = R_{\text{in}}g(\tilde{t} - \tilde{\tau}_u)$ and $j(\tilde{t}) = R_{\text{out}}h(\tilde{t} - \tilde{\tau}_d)$, where $\tilde{\tau}_u$ and $\tilde{\tau}_d$ are the traveling times of the waves from the flame to the up- and downstream boundaries, respectively. If the boundary conditions are ideal, i.e., fully reflective, then $R_{\text{in}} = R_{\text{out}} = -1$. The full set of equations that describe the dynamics of the waves is given by

$$\mathbf{X} \begin{bmatrix} g(\tilde{t}) \\ h(\tilde{t}) \end{bmatrix} = \mathbf{Y} \begin{bmatrix} g(\tilde{t} - \tilde{\tau}_u) \\ h(\tilde{t} - \tilde{\tau}_d) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{\tilde{q} - \tilde{q}'}{A_1 \tilde{c}_1} \end{bmatrix}, \quad (\text{C1})$$

where the matrices \mathbf{X} and \mathbf{Y} are functions of the mean-flow variables and are obtained from the jump conditions. We generate the data with the code implementation of the kinematic flame model from [54] (code available at [66]).

TABLE V. Computation time of nonideal boundary condition higher-fidelity experiments with five pressure sensors. Training with a batch size of 32 for 10 000 epochs (on Intel Core(TM) i7-9750H CPU at 2.60 GHz \times 6).

Network name	Description	Training time
PI-P-GalNN	15 Galerkin modes, 1 harmonic layer with 10 neurons	20 min
PI-P-FNN	4 layers with 32 neurons	35 min




FIG. 20. Reconstruction from only pressure measurements with physics-informed periodic GalNN. From left to right: [(a)–(d)] 2 sensors, 5 Galerkin modes; [(e)–(h)] 5 sensors, 10 Galerkin modes; [(i)–(l)] 11 sensors, 20 Galerkin modes. From top to bottom: Acoustic pressure, error between ground truth and prediction, pressure prediction at a fixed time instance, and error at a fixed time instance, $t = 1.5$.

APPENDIX D: TRAINING COSTS

The training costs associated with some of the characteristic cases are provided in Tables IV and V.

APPENDIX E: RECONSTRUCTION FROM PARTIAL MEASUREMENTS WITH VARYING NUMBER OF SENSORS

Reconstruction from only pressure measurements with physics-informed periodic GalNN is shown in Figs. 20 and 21.

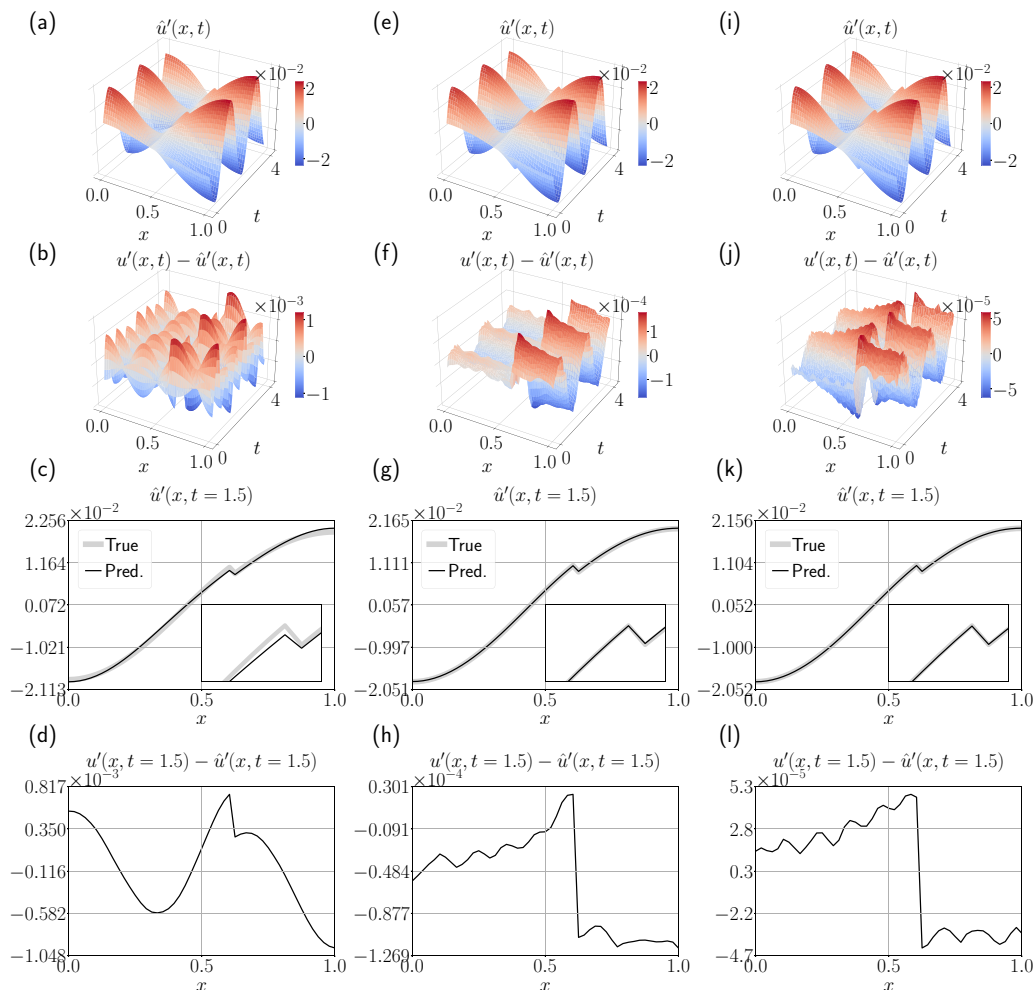


FIG. 21. Reconstruction from only pressure measurements with physics-informed periodic GalNN. From left to right: [(a)–(d)] 2 sensors, 5 Galerkin modes; [(e)–(h)] 5 sensors, 10 Galerkin modes; [(i)–(l)] 11 sensors, 20 Galerkin modes. From top to bottom: Acoustic velocity, error between ground truth and prediction, velocity prediction at a fixed time instance, and error at a fixed time instance, $t = 1.5$.

- [1] K. Duraisamy, G. Iaccarino, and H. Xiao, Turbulence modeling in the age of data, *Annu. Rev. Fluid Mech.* **51**, 357 (2019).
- [2] S. L. Brunton, B. R. Noack, and P. Koumoutsakos, Machine learning for fluid mechanics, *Annu. Rev. Fluid Mech.* **52**, 477 (2020).
- [3] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, Physics-informed machine learning, *Nat. Rev. Phys.* **3**, 422 (2021).
- [4] N. Doan, W. Polifke, and L. Magri, Short-and long-term predictions of chaotic flows and extreme events: A physics-constrained reservoir computing approach, *Proc. R. Soc. A* **477**, 20210135 (2021).
- [5] I. E. Lagaris, A. Likas, and D. I. Fotiadis, Artificial neural networks for solving ordinary and partial differential equations, *IEEE Trans. Neural Networks* **9**, 987 (1998).

- [6] M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378**, 686 (2019).
- [7] M. Raissi, Z. Wang, M. S. Triantafyllou, and G. E. Karniadakis, Deep learning of vortex-induced vibrations, *J. Fluid Mech.* **861**, 119 (2019).
- [8] M. Raissi, A. Yazdani, and G. E. Karniadakis, Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations, *Science* **367**, 1026 (2020).
- [9] M. P. Sitte and N. A. K. Doan, Velocity reconstruction in puffing pool fires with physics-informed neural networks, *Phys. Fluids* **34**, 087124 (2022).
- [10] S. Cai, Z. Wang, F. Fuest, Y. J. Jeon, C. Gray, and G. E. Karniadakis, Flow over an espresso cup: Inferring 3-D velocity and pressure fields from tomographic background oriented Schlieren via physics-informed neural networks, *J. Fluid Mech.* **915**, A102 (2021).
- [11] G. Kissas, Y. Yang, E. Hwuang, W. R. Witschey, J. A. Detre, and P. Perdikaris, Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks, *Comput. Methods Appl. Mech. Eng.* **358**, 112623 (2020).
- [12] M. F. Fathi, I. Perez-Raya, A. Baghaie, P. Berg, G. Janiga, A. Arzani, and R. M. D Souza, Super-resolution and denoising of 4D-flow MRI using physics-informed deep neural nets, *Comput. Methods Programs Biomed.* **197**, 105729 (2020).
- [13] H. Wang, Y. Liu, and S. Wang, Dense velocity reconstruction from particle image velocimetry/particle tracking velocimetry using a physics-informed neural network, *Phys. Fluids* **34**, 017116 (2022).
- [14] H. Eivazi and R. Vinuesa, Physics-informed deep-learning applications to experimental fluid mechanics, [arXiv:2203.15402](https://arxiv.org/abs/2203.15402)
- [15] H. Gao, L. Sun, and J.-X. Wang, Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels, *Phys. Fluids* **33**, 073603 (2021).
- [16] D. Kelshaw, G. Rigas, and L. Magri, Physics-informed CNNs for super-resolution of sparse observations on dynamical systems, [arXiv:2210.17319](https://arxiv.org/abs/2210.17319).
- [17] A. Goyal and Y. Bengio, Inductive biases for deep learning of higher-level cognition, *Proc. R. Soc. A* **478**, 20210068 (2022).
- [18] Y. Chen, D. Huang, D. Zhang, J. Zeng, N. Wang, H. Zhang, and J. Yan, Theory-guided hard constraint projection (HCP): A knowledge-based data-driven scientific machine learning method, *J. Comput. Phys.* **445**, 110624 (2021).
- [19] K. Xu and E. Darve, Physics constrained learning for data-driven inverse modeling from sparse observations, *J. Comput. Phys.* **453**, 110938 (2022).
- [20] A. T. Mohan, N. Lubbers, M. Chertkov, and D. Livescu, Embedding hard physical constraints in neural network coarse-graining of three-dimensional turbulence, *Phys. Rev. Fluids* **8**, 014604 (2023).
- [21] J. Ling, A. Kurzawski, and J. Templeton, Reynolds averaged turbulence modelling using deep neural networks with embedded invariance, *J. Fluid Mech.* **807**, 155 (2016).
- [22] D. Zhang, L. Guo, and G. E. Karniadakis, Learning in modal space: Solving time-dependent stochastic PDEs using physics-informed neural networks, *SIAM J. Sci. Comput.* **42**, A639 (2020).
- [23] S. Dong and N. Ni, A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks, *J. Comput. Phys.* **435**, 110242 (2021).
- [24] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, Galerkin projection, in *Turbulence, Coherent Structures, Dynamical Systems and Symmetry* (Cambridge University Press, Cambridge, UK, 2012), pp. 106–129.
- [25] Rayleigh, The explanation of certain acoustical phenomena, *Nature (London)* **18**, 319 (1878).
- [26] A. P. Dowling and A. S. Morgans, Feedback control of combustion oscillations, *Annu. Rev. Fluid Mech.* **37**, 151 (2005).
- [27] M. P. Juniper and R. I. Sujith, Sensitivity and nonlinearity of thermoacoustic oscillations, *Annu. Rev. Fluid Mech.* **50**, 661 (2018).
- [28] L. Magri, Adjoint methods as design tools in thermoacoustics, *Appl. Mech. Rev.* **71**, 020801 (2019).
- [29] L. Magri, P. J. Schmid, and J. P. Moeck, Linear flow analysis inspired by mathematical methods from quantum mechanics, *Annu. Rev. Fluid Mech.* **55**, 541 (2023).

- [30] T. C. Lieuwen and V. Yang, *Combustion Instabilities in Gas Turbine Engines* (American Institute of Aeronautics and Astronautics, Reston, VA, 2006).
- [31] T. Poinsot, Prediction and control of combustion instabilities in real engines, *Proc. Combust. Inst.* **36**, 1 (2017).
- [32] F. Culick and P. Kuentzmann, Unsteady motions in combustion chambers for propulsion systems, NATO RTO-AG-AVT-039, AGARDograph (2006).
- [33] A. P. Dowling, A kinematic model of a ducted flame, *J. Fluid Mech.* **394**, 51 (1999).
- [34] A. P. Dowling and S. R. Stow, Acoustic analysis of gas turbine combustors, *J. Propul. Power* **19**, 751 (2003).
- [35] J. Li and A. S. Morgans, Time domain simulations of nonlinear thermoacoustic behaviour in a simple combustor using a wave-based approach, *J. Sound Vib.* **346**, 345 (2015).
- [36] M. Bauerheim, J.-F. Parmentier, P. Salas, F. Nicoud, and T. Poinsot, An analytical model for azimuthal thermoacoustic modes in an annular chamber fed by an annular plenum, *Combust. Flame* **161**, 1374 (2014).
- [37] A. Orchini, G. A. Mensah, and J. P. Moeck, Effects of nonlinear modal interactions on the thermoacoustic stability of annular combustors, *J. Eng. Gas Turbines Power* **141**, 021002 (2019).
- [38] B. T. Zinn and M. E. Lores, Application of the Galerkin method in the solution of non-linear axial combustion instability problems in liquid rockets, *Combust. Sci. Technol.* **4**, 269 (1971).
- [39] K. Balasubramanian and R. I. Sujith, Thermoacoustic instability in a Rijke tube: Non-normality and nonlinearity, *Phys. Fluids* **20**, 044103 (2008).
- [40] T. Sayadi, V. L. Chenadec, P. J. Schmid, F. Richecoeur, and M. Massot, Thermoacoustic instability—A dynamical system and time domain analysis, *J. Fluid Mech.* **753**, 448 (2014).
- [41] F. Huhn and L. Magri, Stability, sensitivity and optimisation of chaotic acoustic oscillations, *J. Fluid Mech.* **882**, A24 (2020).
- [42] L. Kabiraj, R. I. Sujith, and P. Wahi, Bifurcations of self-excited ducted laminar premixed flames, *J. Eng. Gas Turbines Power* **134**, 031502 (2012).
- [43] K. I. Matveev, *Thermoacoustic Instabilities in the Rijke Tube: Experiments and Modeling* (California Institute of Technology, Pasadena, CA, 2003).
- [44] A. Novoa and L. Magri, Real-time thermoacoustic data assimilation, *J. Fluid Mech.* **948**, A35 (2022).
- [45] A. Novoa, A. Racca, and L. Magri, Bias-aware thermoacoustic data assimilation, in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings* (Institute of Noise Control Engineering, Glasgow, Scotland, 2023), Vol. 265, pp. 1924–1931.
- [46] F. Selimefendigil and W. Polifke, A nonlinear frequency domain model for limit cycles in thermoacoustic systems with modal coupling, *Int. J. Spray Combust. Dyn.* **3**, 303 (2011).
- [47] S. Jaensch and W. Polifke, Uncertainty encountered when modelling self-excited thermoacoustic oscillations with artificial neural networks, *Int. J. Spray Combust. Dyn.* **9**, 367 (2017).
- [48] N. Tathawadekar, N. A. K. Doan, C. F. Silva, and N. Thuerey, Modeling of the nonlinear flame response of a Bunsen-type flame via multi-layer perceptron, *Proc. Combust. Inst.* **38**, 6261 (2021).
- [49] D. E. Ozan and L. Magri, Physics-aware learning of nonlinear limit cycles and adjoint limit cycles, in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings* (Institute of Noise Control Engineering, Glasgow, Scotland, 2023), Vol. 265, pp. 1191–1199.
- [50] K. Niebler, P. Bonnaire, N. A. K. Doan, and C. F. Silva, Towards reconstruction of acoustic fields via physics-informed neural networks, in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings* (Institute of Noise Control Engineering, Glasgow, Scotland, 2023), Vol. 265, pp. 4773–4782.
- [51] K. Hornik, M. Stinchcombe, and H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* **2**, 359 (1989).
- [52] L. Magri, Introduction to neural networks for engineering and computational science, Zenodo: <https://doi.org/10.5281/zenodo.7655872>.
- [53] P. J. Olver, *Introduction to Partial Differential Equations*, 1st ed., Undergraduate Texts in Mathematics (Springer, Cham, Switzerland, 2020).
- [54] J. G. Aguilar, Sensitivity analysis and optimization in low order thermoacoustic models, Doctoral thesis, University of Cambridge (2019).

- [55] L. Ziyin, T. Hartwig, and M. Ueda, Neural networks fail to learn periodic functions and how to fix it, in *Advances in Neural Information Processing Systems* (Curran Associates, 2020), pp. 1583–1594.
- [56] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from [tensorflow.org](https://www.tensorflow.org)
- [57] L. Magri and M. P. Juniper, Global modes, receptivity, and sensitivity analysis of diffusion flames coupled with duct acoustics, *J. Fluid Mech.* **752**, 237 (2014).
- [58] L. Magri and M. P. Juniper, Sensitivity analysis of a time-delayed thermo-acoustic system via an adjoint-based approach, *J. Fluid Mech.* **719**, 183 (2013).
- [59] M. P. Juniper, Triggering in the horizontal Rijke tube: Non-normality, transient growth and bypass transition, *J. Fluid Mech.* **667**, 272 (2011).
- [60] L. D. Landau and E. M. Lifschitz, *Fluid Mechanics*, 2nd ed. (Pergamon Press, New York, 1987).
- [61] M. A. Heckl, Nonlinear acoustic effects in the Rijke tube, *Acustica* **72**, 63 (1990).
- [62] L. N. Trefethen, *Spectral Methods in MATLAB* (Society for Industrial and Applied Mathematics, Philadelphia, 2000).
- [63] K. He, X. Zhang, S. Ren, and J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in *Proceedings of the IEEE International Conference on Computer Vision* (IEEE, Piscataway, NJ, 2015), pp. 1026–1034.
- [64] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (PMLR, Sardinia, Italy, 2010), pp. 249–256.
- [65] H. Levine and J. Schwinger, On the radiation of sound from an unflanged circular pipe, *Phys. Rev.* **73**, 383 (1948).
- [66] J. G. Aguilar, Flame transfer function tutorial, <https://xoeg.github.io/Flame-Transfer-Function-Tutorial/>