

Data-driven computation of adjoint sensitivities without adjoint solvers: An application to thermoacoustics

Original

Data-driven computation of adjoint sensitivities without adjoint solvers: An application to thermoacoustics / Ozan Defne, E.; Magri, Luca. - In: PHYSICAL REVIEW FLUIDS. - ISSN 2469-990X. - 9:(2024), pp. 1-28.
[10.1103/PhysRevFluids.9.103902]

Availability:

This version is available at: 11583/2995090 since: 2024-12-23T22:47:03Z

Publisher:

American Physical Society - APS

Published

DOI:10.1103/PhysRevFluids.9.103902



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)

Data-driven computation of adjoint sensitivities without adjoint solvers: An application to thermoacoustics

Defne E. Ozan ¹ and Luca Magri ^{1,2,3}¹*Imperial College London, Department of Aeronautics, Exhibition Road, London SW7 2BX, United Kingdom*²*The Alan Turing Institute, London NW1 2DB, United Kingdom*³*DIMEAS, Politecnico di Torino, Corso Duca degli Abruzzi, 24, 10129 Torino, Italy*

(Received 17 April 2024; accepted 17 September 2024; published 25 October 2024)

Adjoint methods have been the pillar of gradient-based optimization for decades. They enable the accurate computation of a gradient (sensitivity) of a quantity of interest with respect to all system parameters in one calculation. When the gradient is embedded in an optimization routine, the quantity of interest can be optimized for the system to have the desired behavior. Adjoint methods, however, require the system's governing equations and their Jacobian. In this paper, we propose a computational strategy to infer the adjoint sensitivities from data when the governing equations might be unknown (or partly unknown), and noise might be present. The key component of this strategy is an echo state network, which learns the dynamics of nonlinear regimes with varying parameters, and it evolves dynamically via a hidden state. Although the framework does not make assumptions on the dynamical system, we focus on thermoacoustics, which are governed by nonlinear and time-delayed systems. First, we show that a parameter-aware echo state network (ESN) infers the parametrized dynamics. Second, we derive the adjoint of the ESN to compute two types of sensitivity: (i) parameter sensitivity, which is the gradient of a time-averaged cost functional with respect to physical or design parameters of the system, and (ii) initial condition sensitivity, which is the gradient of a cost functional of the final state with respect to the initial condition. Third, we propose the thermoacoustic echo state network (T-ESN), which embeds the physical knowledge in the network architecture for improved generalization. Fourth, we apply the framework to a variety of nonlinear thermoacoustic regimes of a prototypical system. We show that the T-ESN accurately infers the correct adjoint sensitivities of the acoustic energy with respect to the flame parameters and initial conditions. The results are robust to noisy data, from periodic, through quasiperiodic, to chaotic regimes. The inferred adjoint sensitivities are employed to suppress an instability via steepest descent. We show that a single network predicts the nonlinear bifurcations on unseen scenarios, which allows it to converge to the minimum of the acoustic energy. This work opens new possibilities for gradient-based data-driven design optimization.

DOI: [10.1103/PhysRevFluids.9.103902](https://doi.org/10.1103/PhysRevFluids.9.103902)

I. INTRODUCTION

Engineers and scientists strive to design flow and acoustic systems that behave in an optimal way (e.g., Refs. [1–5]). Finding the global optimum is challenging because the temporal dynamics can be nonlinear, time delayed, and chaotic, which makes the objective function nonconvex

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.

(e.g., Refs. [6,7]). A practical approach is to solve the optimization problem by gradient descent, in which the objective function is optimized along the gradient direction until a local optimum is found. The gradient (also known as sensitivity) offers information about how the quantity of interest changes with the system's parameters. Central to gradient-based optimization methods are adjoint methods. Adjoint methods enable the accurate computation of the gradient with a computational cost that is not affected by the number of design parameters. For example, if the system has N parameters, adjoint methods compute the gradient of an objective function with only one calculation. This is in contrast with tangent linear or finite difference methods, which would require N computations (e.g., Refs. [2,3,5,8–10]).

The optimization of a quantity of interest of a nonlinear system with adjoint methods requires, in general, six ingredients:

- (i) The system's nonlinear governing equations are integrated forward in time to obtain the baseline (direct) solution.
- (ii) The nonlinear equations are linearized to obtain the Jacobian of the system.
- (iii) The adjoint equations are derived from the Jacobian.
- (iv) The adjoint equations are initialized at the end of the direct solution, and are integrated backwards in time. (Because in nonlinear systems the Jacobian depends on the direct solution, strategies such as checkpointing become often necessary to decrease the memory requirements [11].)
- (v) At the end of the adjoint integration, the gradient of the objective function to all the design parameters of interest can be evaluated by an inner product.
- (vi) The gradient is embedded in a gradient-based optimization scheme to iteratively compute the optimum of the objective function.

This process is also known as direct-adjoint looping. On the one hand, in finite difference or tangent linear approaches, the computational cost increases linearly with the number of parameters because the sensitivity calculation needs to be performed for each parameter. On the other hand, the adjoint method computes the gradient of the objective with respect to all parameters with a single operation, which makes it computationally cheaper when the number of parameters is larger than the number of objective functionals (e.g., Refs. [8,12–15]). For this reason, adjoint methods have found extensive application in design optimization, and in both linear and nonlinear sensitivity analysis of fluid dynamics applications.

To bypass the need of the system's equations, a data-driven model can be employed. The key idea of this paper is to learn an approximation of the dynamical systems from observables, and then compute the adjoint of it. This approach is useful when working with experimental data, or data whose equations are not fully known. To infer the governing dynamics from observables (data), we employ a type of reservoir computer, the echo state network (ESN) [16], which is a universal approximator and a type of recurrent neural network (RNN) [17]. The ESN's training is straightforward because it does not require back-propagation. Its Jacobian can also be computed in a straightforward way due to the shallow architecture [18]. ESNs have been shown to successfully (i) make time-accurate short-term forecasts [19–21], (ii) reconstruct unmeasured variables [22,23], (iii) learn ergodic properties, such as Lyapunov exponents, and long-term statistics of chaotic flows [7,18,24], and (iv) infer model errors in thermoacoustics [25]. In recent advances, low-order representations of turbulence were learned from high-dimensional data, and the dynamics of this low-dimensional system was propagated with the ESN [26]. In this paper, we take an additional step. In order to perform design optimization of a system's parameters, we require a parameter-aware extension of the echo state network [27], which learns the parametrized dynamics and can operate across parameter regimes. A recent study [28] demonstrated the proposed data-driven sensitivity framework on a prototypical chaotic system, and obtained the adjoint sensitivities of long-time averages to the system's parameters.

In this paper, we develop a data-driven sensitivity framework, with an application on nonlinear and time-delayed thermoacoustic oscillations. These are phenomena that can be detrimental in the manufacturing of gas turbines [29] and rocket engines [30]. In thermoacoustic systems, if the heat

release and acoustic pressure are sufficiently in phase, self-sustained oscillations can occur [31,32]. Therefore, their prediction, control, and suppression are important both during design and during operation. This is especially challenging because thermoacoustics exhibit nonlinear behavior and bifurcations to periodic, quasiperiodic, and chaotic regimes [33], which motivates the development of methods for nonlinear analysis. Preliminary design of thermoacoustic systems is performed by linear stability [29,34]. Linear stability analysis determines the growth rates of infinitesimal perturbations to the mean flow, i.e., the eigenvalues of the linearized system of equations around the mean flow [35]. To name a few examples, adjoint-based linear stability analysis enabled calculation of sensitivity to parameters and passive control mechanisms [36], optimal placement and tuning of acoustic dampers [37], and stabilization of thermoacoustic oscillations by geometry changes [38]. An in-depth review of the development of adjoint methods for thermoacoustics can be found in Refs. [5,35]. Beyond linear stability analysis, adjoint methods were used for weakly nonlinear analysis to predict subcritical Hopf bifurcations [39], and Floquet analysis was used to calculate the stability and the period of limit-cycle oscillations [5]. The stability of chaotic thermoacoustics was studied in Ref. [6] using covariant Lyapunov analysis, where it was analytically shown that this is equivalent to eigenvalue and Floquet analysis for fixed-point and limit-cycle solutions, respectively. As a result, Ref. [6] achieved gradient-based minimization of acoustic energy in chaotic thermoacoustics using shadowing methods. The application of this method was challenged by the discontinuities observed in the acoustic energy when bifurcations occur, the computational cost of shadowing methods, and the nonhyperbolicity of thermoacoustics for some regimes. To circumvent these issues, a gradient-free optimization method based on Bayesian optimization was proposed [7]. Among applications of data-driven methods in thermoacoustics, the ESN was previously used to predict the nonlinear dynamics [7] and calculate the Lyapunov exponents [40] of a chaotic regime, and infer the model error of a low-order model [25].

The overarching objective of this paper is to address the challenge of developing code-specific adjoint solvers, which requires the system's equations and the Jacobian to be known. We propose data-driven nonlinear direct-adjoint looping to compute sensitivities to the system's parameters and initial conditions, and apply this to a prototypical thermoacoustic system. The specific objectives are to introduce a robust parameter-aware ESN architecture for this purpose, derive its adjoint, and demonstrate the data-driven inference of parametrized dynamics and associated adjoint sensitivities. We refer to this architecture as the thermoacoustic ESN and show that (i) it learns the parametrized dynamics of nonlinear, time-delayed thermoacoustics, (ii) predicts bifurcations, (iii) accurately obtains the adjoint sensitivities, (iv) the acoustic energy is minimized in a gradient-based optimization framework, and (v) demanding scenarios such as training with noisy or chaotic data can be handled.

This paper is organized as follows. In Sec. II, we give a background on the mathematical framework for sensitivity analysis and the parameter-aware ESN. In Sec. III, we derive the adjoint of the ESN. In Sec. IV, we present the nonlinear and time-delayed thermoacoustic model, namely, the Rijke tube. In Sec. V, we introduce the time-delayed ESN designed for thermoacoustics and derive its adjoint. In Sec. VI, we show the performance of the ESN in terms of prediction and accuracy of the learned adjoint sensitivities, and employ it for the gradient-based design optimization of thermoacoustics in Sec. VII. We address training with noisy data in Sec. VIII, and learning chaotic regimes in Sec. IX. We show the sensitivity to initial conditions in Sec. X. A discussion is offered in Sec. XI. Section XII concludes the paper.

II. BACKGROUND

A. Sensitivity analysis

We consider a dynamical system in the form of

$$\dot{\mathbf{x}} - \mathbf{F}(\mathbf{x}, \mathbf{p}) = 0, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^{n_x}$ is the state vector, $\dot{\mathbf{x}}$ is its time derivative, $\mathbf{p} \in \mathbb{R}^{n_p}$ is the parameters vector, and $\mathbf{F} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_p} \rightarrow \mathbb{R}^{n_x}$ is the operator that describes the system's evolution. We assume that \mathbf{F} is

smooth so that the Jacobian of the system, i.e., $\partial \mathbf{F} / \partial \mathbf{x}$, is defined everywhere, which is a requirement for gradient-based sensitivity analysis. We wish to optimize an objective functional, \mathcal{J} , which is a time average of $\tilde{\mathcal{J}}$ [5],

$$\mathcal{J} = \frac{1}{T} \int_0^T \tilde{\mathcal{J}}(\mathbf{x}(t), \mathbf{p}) dt, \quad (2)$$

subjected to the system dynamics (1). A straightforward approach to obtain sensitivity information is to apply a small perturbation, h , to the parameters of interest, $p_i \in \mathbf{p}$, and numerically compute the sensitivity using finite differences:

$$\left. \frac{d\mathcal{J}}{d\mathbf{p}} \right|_{\mathbf{p}=\bar{\mathbf{p}}} = \left. \frac{d\mathcal{J}}{dp_i} \right|_{p_i=\bar{p}_i} \approx \frac{\mathcal{J}|_{p_i=\bar{p}_i+h} - \mathcal{J}|_{p_i=\bar{p}_i}}{h} + O(h), \quad i = 1, 2, \dots, n_p, \quad (3)$$

which needs to be repeated for each $p_i \in \mathbf{p}$ to determine $d\mathcal{J}/d\mathbf{p}$. This means that the computational cost of the sensitivity increases linearly with the number of parameters. A more efficient approach is to use the adjoint equation. Analytically, we can express the gradient $d\mathcal{J}/d\mathbf{p}$ as

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \frac{1}{T} \int_0^T \left(\frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{p}} + \frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{x}} \mathbf{Q} \right) dt, \quad (4a)$$

$$\dot{\mathbf{Q}} = \frac{\partial \mathbf{F}}{\partial \mathbf{p}} + \frac{\partial \mathbf{F}}{\partial \mathbf{x}} \mathbf{Q}, \quad (4b)$$

$$\mathbf{Q}(0) = \mathbf{0}, \quad (4c)$$

where we define $\mathbf{Q} \equiv d\mathbf{x}/d\mathbf{p}$, $\mathbf{Q} \in \mathbb{R}^{n_x \times n_p}$. We refer to solving the system of equations (4) as the tangent linear problem. In this case, the sensitivity can be computed by formulating the tangent linear system around a base trajectory and solving for \mathbf{Q} in time. Although this approach eliminates the finite difference numerical error, it still suffers from a computational cost that is linear with the number of parameters as the dimension of \mathbf{Q} grows with the dimension of \mathbf{p} . Adjoint methods circumvent this issue. By introducing the Lagrangian of objective (2) and the adjoint variables $\mathbf{q}^+ \in \mathbb{R}^{n_x}$ as the Lagrange multipliers of constraint (1), the gradient $d\mathcal{J}/d\mathbf{p}$ can be rewritten as

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \frac{1}{T} \int_0^T \left(\frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{p}} + \mathbf{q}^{+\top} \frac{\partial \mathbf{F}}{\partial \mathbf{p}} \right) dt, \quad (5a)$$

$$\dot{\mathbf{q}}^+ = -\frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{x}} - \frac{\partial \mathbf{F}}{\partial \mathbf{x}} \mathbf{q}^+, \quad (5b)$$

$$\mathbf{q}^+(T) = \mathbf{0}, \quad (5c)$$

where the adjoint variables \mathbf{q}^+ need to be solved by integrating the adjoint system backwards in time. The details of the derivation of the adjoint problem from a Lagrangian perspective for a continuous-time system can be found in Ref. [5]. In this paper, we derive the tangent linear and adjoint equations for the proposed data-driven algorithm in Appendix A.

B. Parameter-aware echo state networks

Given the evolution equations of a dynamical system (1), we can compute its adjoint (5). Developing the adjoint solver, however, requires the system Jacobian, $\partial \mathbf{F} / \partial \mathbf{x}$. This Jacobian is code specific. The key idea of this paper is to eliminate code-specific adjoint solvers by having a machine that learns the dynamics from data, and is itself a dynamical system, the Jacobian of which is readily available independent of the system under investigation. RNNs are well suited for this task, as they are neural network architectures that can process sequential data and make forecasts of time series from dynamical systems. Although the dynamical systems we consider in Eq. (1) are continuous in time, RNNs accept data as a sequence in discrete time. RNNs possess memory through a hidden

state and thus they can learn temporal relationships. The hidden state enables the RNNs to be framed as discrete dynamical systems [18]. In this work, we utilize the ESN [41], which is a type of RNN that is straightforward to train because it does not require back-propagation during training. The ESN is a universal approximator [17] and it follows this dynamic equation:

$$\mathbf{r}(i+1) = (1 - \alpha)\mathbf{r}(i) + \alpha \tanh(\mathbf{W}_{\text{in}}\mathbf{y}_{\text{in}}(i) + \mathbf{W}\mathbf{r}(i)), \quad (6a)$$

where $\mathbf{y}_{\text{in}} \in \mathbb{R}^{n_y}$ is the input vector, $\mathbf{r}(i) \in \mathbb{R}^{n_r}$ is the reservoir state, $\mathbf{W}_{\text{in}} \in \mathbb{R}^{n_r \times n_y}$ is the input matrix, and $\mathbf{W} \in \mathbb{R}^{n_r \times n_r}$ is the state matrix. Equation (6a) maps the reservoir state at time t_i to the reservoir state at time $t_{i+1} = t_i + \Delta t$. The prediction of the output is a linear combination of the reservoir state,

$$\hat{\mathbf{y}}(i+1) = \mathbf{W}_{\text{out}}\mathbf{r}(i+1), \quad (6b)$$

where $\mathbf{y} \in \mathbb{R}^{n_y}$ is the target output vector, the dynamics of which we aim to model; $\hat{\mathbf{y}} \in \mathbb{R}^{n_y}$ is the predicted output vector; and $\mathbf{W}_{\text{out}} \in \mathbb{R}^{n_y \times n_r}$ is the output matrix that maps the reservoir state to the output vector. The matrices \mathbf{W}_{in} and \mathbf{W} are sparse, randomly generated, and not trained, while \mathbf{W}_{out} is trained via ridge regression [16]. The input matrix is scaled as $\mathbf{W}_{\text{in}} = \sigma_{\text{in}}\tilde{\mathbf{W}}_{\text{in}}$, where $\tilde{\mathbf{W}}_{\text{in}}$ is constructed such that each row contains only one nonzero element drawn from a uniform distribution $\sim \mathcal{U}(-1, 1)$; i.e., the inputs are activated individually. The state matrix \mathbf{W} describes the reservoir connections, and is scaled as $\mathbf{W} = \rho\tilde{\mathbf{W}}$, where $\tilde{\mathbf{W}}$ is an Erdős-Renyi matrix with a given average number of nonzero elements per row (connectivity) drawn from a uniform distribution $\sim \mathcal{U}(-1, 1)$, which is scaled to have a unity spectral radius. So, the scalar ρ is set as the spectral radius of $\tilde{\mathbf{W}}$. In this architecture, the reservoir state variables “echo” the input, and so the reservoir holds memory. During training \mathbf{y}_{in} is fed into the network in open loop and the evolution of the reservoir state is saved. The reservoir state is initialized to a zero vector. An initial time window is used for the washout stage, which eliminates the effects of the initial condition and is discarded in training. The training is performed by solving for \mathbf{W}_{out} [16],

$$\mathbf{W}_{\text{out}}^* = \arg \min_{\mathbf{W}_{\text{out}}} \frac{1}{n_y} \sum_{j=1}^{n_y} \left(\sum_{i=1}^{N_t} (y_j(i) - \hat{y}_j(i))^2 + \lambda \|\mathbf{w}_{\text{out},j}\|^2 \right), \quad (7)$$

where N_t is the number of training steps, $\mathbf{w}_{\text{out},j}$ is the j th row of \mathbf{W}_{out} , and λ is the Tikhonov coefficient. This is a ridge regression problem, which minimizes both the error between network predictions for the next step and the norm of the output weights for regularization. The solution is [16]

$$\mathbf{W}_{\text{out}}^* = \mathbf{Y}\mathbf{R}^T(\mathbf{R}\mathbf{R}^T + \lambda\mathbf{I})^{-1}, \quad (8)$$

where $\mathbf{R} = [\mathbf{r}(1) \ \mathbf{r}(2) \ \dots \ \mathbf{r}(N_t)]$, $\mathbf{R} \in \mathbb{R}^{n_r \times N_t}$, $\mathbf{Y} = [\mathbf{y}(1) \ \mathbf{y}(2) \ \dots \ \mathbf{y}(N_t)]$, $\mathbf{Y} \in \mathbb{R}^{n_y \times N_t}$, and \mathbf{I} is the identity matrix. Tuning an ESN involves searching for optimal hyperparameters that consist of input matrix scaling σ_{in} , the spectral radius of the state matrix ρ (typically ≤ 1 to satisfy the echo state property), the leak rate α , and the Tikhonov coefficient λ . These hyperparameters are determined by validation. A detailed explanation of training and robust validation of echo state networks can be found in Ref. [21]. After training, the ESN can run autonomously in closed loop by setting $\mathbf{y}_{\text{in}}(i) = \hat{\mathbf{y}}(i)$.

In formulation (6a), the ESN is unaware of the system parameters; i.e., it can only operate in one regime. For our purposes, we require an ESN to learn the parameter-dependent dynamics of a system such that we can capture the variations in different regimes. A parameter-aware extension of the ESN has been employed to predict when the system bifurcates to fixed-point solutions, a phenomenon known as amplitude death [27], and bifurcations in multistable regimes including chaos [42]. It can be formulated as [27]

$$\mathbf{r}(i+1) = (1 - \alpha)\mathbf{r}(i) + \alpha \tanh(\mathbf{W}_{\text{in}}[\mathbf{y}_{\text{in}}(i); \text{diag}(\boldsymbol{\sigma}_p)(\mathbf{p} - \mathbf{k}_p)] + \mathbf{W}\mathbf{r}(i)), \quad (9)$$

where $(;)$ denotes vertical concatenation, and $\mathbf{p} \in \mathbb{R}^{n_p}$ denotes the system's parameters, which are fed to the network through the input channel, and we assume to be constant in time for each system. The input matrix $\mathbf{W}_{\text{in}} = [\mathbf{W}_{\text{in}}^y \ \mathbf{W}_{\text{in}}^p]$, $\mathbf{W}_{\text{in}}^y \in \mathbb{R}^{n_r \times n_y}$, $\mathbf{W}_{\text{in}}^p \in \mathbb{R}^{n_r \times n_p}$ is constructed such that the parameters are fully connected with all the reservoir nodes; i.e., the weights associated with the parameters \mathbf{W}_{in}^p are nonzero for all rows. The parameter-aware ESN adds two additional hyperparameters per physical parameter p_i . A scalar k_{p_i} shifts the parameter p_i as $p_i - k_{p_i}$, and a scalar σ_{p_i} scales it as $\sigma_{p_i}(p_i - k_{p_i})$. Hence, we apply the input scaling σ_{in} on \mathbf{W}_{in}^y only. The training of the parameter-aware ESN is performed as the standard ESN (8). The data set contains data from different parameter regimes, which means that we need to save the evolution of the reservoir state for each training regime, and concatenate as $\mathbf{R} = [\mathbf{R}^1 \ \mathbf{R}^2 \ \dots \ \mathbf{R}^{N_{tr}}]$ (likewise, the outputs are concatenated as $\mathbf{Y} = [\mathbf{Y}^1 \ \mathbf{Y}^2 \ \dots \ \mathbf{Y}^{N_{tr}}]$), where N_{tr} is the number of training regimes.

III. ADJOINT OF THE ECHO STATE NETWORK

In this section, we derive the adjoint of the parameter-aware ESN. In closed loop, the ESN is governed by

$$\mathbf{r}(i+1) = (1 - \alpha)\mathbf{r}(i) + \alpha \tanh(\mathbf{W}_{\text{in}}[\mathbf{W}_{\text{out}}\mathbf{r}(i); \text{diag}(\sigma_{p_i})(\mathbf{p} - \mathbf{k}_p)] + \mathbf{W}\mathbf{r}(i)), \quad (10)$$

which is an autonomous dynamical system. This is the configuration that we use for the computation of the adjoint. We first tackle sensitivity to the system's parameters, and then to the system's initial conditions.

A. Sensitivity to the system's parameters

We consider the time-averaged objective functional (2), which we wish to optimize. In discrete time, we reformulate it as a sum over a discrete number of time steps, denoted by N :

$$\mathcal{J} = \frac{1}{N} \sum_{i=1}^N \tilde{\mathcal{J}}(\mathbf{r}(i)). \quad (11)$$

Without loss of generality, we assume that the objective functional does not explicitly depend on the parameters, i.e., $\partial \tilde{\mathcal{J}} / \partial \mathbf{p} = 0$. For the parameter-aware ESN, the tangent linear problem is

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \frac{d\tilde{\mathcal{J}}(\mathbf{r}(i))}{d\mathbf{r}(i)} \mathbf{Q}(i), \quad (12a)$$

$$\mathbf{Q}(i+1) = \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{p}} + \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i)} \mathbf{Q}(i), \quad (12b)$$

$$\mathbf{Q}(0) = \mathbf{0}, \quad (12c)$$

where $\mathbf{Q}(i) \equiv d\mathbf{r}(i)/d\mathbf{p}$, $\mathbf{Q}(i) \in \mathbb{R}^{n_r \times n_p}$. Physically, Eq. (12c) means that the initial condition is independent of the parameter. Because the dimension of the tangent linear system, $\mathbf{Q}(i)$, grows with the number of parameters, adjoint methods enable a computationally cheap and accurate way to obtain sensitivities [5]. Using the Lagrangian formulation of the objective and constraints given by the evolution equations (Appendix A), the adjoint problem is formulated as

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \sum_{i=1}^N \mathbf{q}^{+\top}(i) \frac{\partial \mathbf{r}(i)}{\partial \mathbf{p}}, \quad (13a)$$

$$\mathbf{q}^+(i) = \frac{1}{N} \frac{d\tilde{\mathcal{J}}(\mathbf{r}(i))}{d\mathbf{r}(i)}^\top + \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i)}^\top \mathbf{q}^+(i+1), \quad (13b)$$

$$\mathbf{q}^+(N) = \frac{1}{N} \frac{d\tilde{\mathcal{J}}(\mathbf{r}(N))}{d\mathbf{r}(N)}^\top. \quad (13c)$$

Equations (13) are in a form similar to the continuous time equations (5); however, we have replaced the system's Jacobian with the ESN's Jacobian and thus the adjoint variables are the adjoint of the ESN's reservoir state. Algorithmically, we first obtain the direct solution, $\mathbf{r}(i)$, as a baseline trajectory by running the ESN in closed loop for a time window. Then, the adjoint equations (13b) are solved backwards in time starting from the terminal condition $\mathbf{q}^+(N)$ [Eq. (13c)]. This procedure requires the computation of the Jacobian, $\partial\mathbf{r}(i+1)/\partial\mathbf{r}(i)$, and the partial derivative with respect to the parameters $\partial\mathbf{r}(i)/\partial\mathbf{p}$, associated with the ESN. These gradients are provided in Appendix B.

The objective functional we are interested in depends on the output vector, e.g.,

$$\mathcal{J} = \frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{y}}(i)\|_2^2 = \frac{1}{N} \sum_{i=1}^N \|\mathbf{W}_{\text{out}}\mathbf{r}(i)\|_2^2. \quad (14)$$

Thus, the gradient of the objective with respect to the reservoir state is

$$\frac{d\tilde{\mathcal{J}}(\mathbf{r}(i))}{d\mathbf{r}(i)} = 2\mathbf{W}_{\text{out}}^\top \mathbf{W}_{\text{out}}\mathbf{r}(i). \quad (15)$$

Minimization of Eq. (14) defines the optimization problem where the ESN models the physical dynamical system.

B. Sensitivity to the system's initial conditions

We consider the sensitivity of an objective functional of the final reservoir state after N time steps, i.e., $\mathcal{J} = \mathcal{J}(\mathbf{r}(N))$, to the initial conditions. In the ESN's Eqs. (9), the system's initial conditions $\mathbf{y}(0)$ appear in the input as $\mathbf{y}_{\text{in}}(0)$. When computing the sensitivity to the initial input vector $\mathbf{y}_{\text{in}}(0)$, we need to make a distinction between whether the first step of the propagation of the ESN's equations is performed in open loop or closed loop. Because of the recurrence in the ESN, the full state of the ESN is described by the reservoir state. Therefore, for the ESN's estimate of the sensitivity to match the original system, we first run a washout so that the output of the reservoir synchronizes with the original system and then perform the first step in closed loop. (For a use case of the resulting open-loop Jacobian, see an application in data assimilation in Ref. [25].) The sensitivity to the initial condition is

$$\frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{y}_{\text{in}}(0)} = \frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{r}(N)} \frac{d\mathbf{r}(N)}{d\mathbf{r}(1)} \left(\frac{\partial\mathbf{r}(1)}{\partial\mathbf{y}_{\text{in}}(0)} + \frac{\partial\mathbf{r}(1)}{\partial\mathbf{r}(0)} \frac{d\mathbf{r}(0)}{d\mathbf{y}_{\text{in}}(0)} \right), \quad (16)$$

where the gradients $\partial\mathbf{r}(1)/\partial\mathbf{y}_{\text{in}}(0)$ and $\partial\mathbf{r}(1)/\partial\mathbf{r}(0)$ are found by differentiating the open-loop Eq. (9) and the gradients $\partial\mathbf{r}(i+1)/\partial\mathbf{r}(i)$, $i = 1, 2, \dots, N-1$, are found by differentiating the closed-loop Eq. (10). From the output of the ESN [Eq. (6b)], we find the gradient $d\mathbf{r}(0)/d\mathbf{y}_{\text{in}}(0) = \mathbf{W}_{\text{out}}^\dagger$, where $(\cdot)^\dagger$ indicates pseudoinverse. Alternatively, we compute $d\mathcal{J}(\mathbf{r}(N))/d\mathbf{r}(0)$ in closed-loop configuration for all time steps and then determine by the chain rule

$$\frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{y}_{\text{in}}(0)} = \frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{r}(0)} \mathbf{W}_{\text{out}}^\dagger. \quad (17)$$

The tangent linear equations are given by

$$\frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{r}(0)} = \frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{r}(N)} \mathbf{Q}(N), \quad (18a)$$

$$\mathbf{Q}(i+1) = \frac{\partial\mathbf{r}(i+1)}{\partial\mathbf{r}(i)} \mathbf{Q}(i), \quad (18b)$$

$$\mathbf{Q}(0) = \mathbf{I}, \quad (18c)$$

where $\mathbf{Q}(i) \equiv d\mathbf{r}(i)/d\mathbf{r}(0)$, $\mathbf{Q}(i) \in \mathbb{R}^{n_r \times n_r}$, and \mathbf{I} is the identity matrix. The adjoint equations are

$$\frac{d\mathcal{J}(\mathbf{r}(N))}{d\mathbf{r}(0)} = \mathbf{q}^{+\top}(0), \quad (19a)$$

$$\mathbf{q}^+(i) = \frac{\partial \mathbf{r}(i+1)^\top}{\partial \mathbf{r}(i)} \mathbf{q}^+(i+1), \quad (19b)$$

$$\mathbf{q}^+(N) = \frac{d\mathcal{J}(\mathbf{r}(N))^\top}{d\mathbf{r}(N)}. \quad (19c)$$

In this case, the number of ‘‘parameters’’ is equal to the dimension of the reservoir state vector, which is typically large, and therefore the adjoint approach significantly reduces the computational cost as compared to finite differences or the tangent linear approach.

IV. THERMOACOUSTICS

We consider a prototypical thermoacoustic system known as the Rijke tube, which consists of a duct with a heat source located inside. The Rijke tube is a setup used to propose fundamental methods of thermoacoustic instabilities because it captures the qualitative nonlinear dynamics and bifurcations observed in real-life applications (e.g., Refs. [6,33,43]). In this model, we make the following assumptions: (i) the acoustics are one dimensional, (ii) the mean flow has a low Mach number and the acoustic variables are modeled as perturbations on top of a mean flow with uniform density across the tube, (iii) the heat release is compact, and (iv) the boundary conditions are ideal, i.e., fully reflective. The governing partial differential equations (PDEs) are derived from conservation of momentum, energy, and mass (e.g., Ref. [5]),

$$\frac{\partial u}{\partial t} + \frac{\partial p}{\partial x} = 0, \quad (20a)$$

$$\frac{\partial p}{\partial t} + \frac{\partial u}{\partial x} + \zeta p - \dot{q}\delta(x - x_f) = 0, \quad (20b)$$

where u and p are the acoustic velocity and pressure fluctuations, respectively, ζ is the frequency-dependent damping, which accounts for the acoustic dissipation, \dot{q} is the nonlinear heat release model, and $\delta(x - x_f)$ is a Dirac delta at the heat source location x_f . The solution to the governing PDEs is expressed by a Galerkin decomposition on n_g acoustic eigenfunctions as [44]

$$u(x, t) = \sum_{j=1}^{n_g} \eta_j(t) \cos(j\pi x), \quad p(x, t) = - \sum_{j=1}^{n_g} \mu_j(t) \sin(j\pi x), \quad (21)$$

where $j\pi$ are the nondimensional angular frequencies of the purely acoustic modes [45]. We denote the vectors of velocity and pressure Galerkin amplitudes by $\boldsymbol{\eta} = [\eta_1; \eta_2; \dots; \eta_{n_g}]$ and $\boldsymbol{\mu} = [\mu_1; \mu_2; \dots; \mu_{n_g}]$, respectively. By substituting pressure and velocity variables in the PDEs (20) with their Galerkin decompositions in Eq. (21) 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 (ODEs)

$$\dot{\eta}_j - \mu_j j\pi = 0, \quad (22a)$$

$$\dot{\mu}_j + \eta_j j\pi + \zeta_j \mu_j + 2\dot{q} \sin(j\pi x_f) = 0, \quad (22b)$$

where the modal damping is modeled by $\zeta_j = c_1 j^2 + c_2 j^{1/2}$, and \dot{q} is the heat release rate described by a nonlinear time-delay model [46],

$$\dot{q} = \beta(\sqrt{|1 + u(x_f, t - \tau)|} - 1), \quad (23)$$

where β is the heat release strength, and τ is the time delay between the velocity at the heat source and heat release. The system exhibits different nonlinear solutions, e.g., bifurcation from fixed point

to limit cycle, when the parameters x_f , β , and τ are varied. In the remainder of this paper, we fix $x_f = 0.2$, which causes thermoacoustic instabilities to occur, and analyze the dynamics as functions β and τ . The derivative of the nonlinear heat-release model (23) with respect to $u(x_f, t - \tau)$ is not defined for $u(x_f, t - \tau) = -1$, which means that the Jacobian and consequently the sensitivity of the system cannot be calculated around this value. Therefore, the heat-release model (23) is modified to have smooth dynamics by approximating a small neighborhood of $u(x_f, t - \tau) = -1$ with a fourth-degree polynomial [6]. Further, the time-delayed problem is reformulated as an advection equation via the dummy variable v to make the problem Markovian [6]:

$$\frac{\partial v}{\partial t} + \frac{1}{\tau} \frac{\partial v}{\partial X} = 0, \quad 0 \leq X \leq 1, \quad v(X = 0, t) = u(x_f, t). \quad (24)$$

PDE (24) is discretized using a Chebyshev spectral method with n_c points such that the dimension of the system is $2n_g + n_c$ and, thus, can be solved with a standard time-marching scheme. For the thermoacoustic system, the objective functional \mathcal{J} is the time-averaged acoustic energy, which relates to the average energy of the undesired oscillations, given by (e.g., Ref. [6])

$$E_{ac}(t) = \int_0^1 \frac{1}{2} (u^2(x, t) + p^2(x, t)) dx = \frac{1}{4} \sum_{j=1}^{n_g} (\eta_j^2(t) + \mu_j^2(t)). \quad (25)$$

V. THERMOACOUSTIC ECHO STATE NETWORK

We propose a parameter-aware ESN with hard-constrained physics for thermoacoustics [47]. In the context of machine learning, hard constraining means embedding prior knowledge into the architecture of the neural network, as opposed to soft constraining, which refers to imposing prior knowledge as a penalty during training [20]. This approach makes neural networks more robust and generalizable on unseen data [47]. First of all, our aim is to develop a parameter-aware ESN that learns the dynamics of the Galerkin amplitudes (21) from the thermoacoustic system introduced in Sec. IV. Initially, we consider two cases without any hard constraining: (i) the Galerkin amplitudes as the input and output vectors $\mathbf{y}_{in}(i) = \hat{\mathbf{y}}(i) = [\hat{\boldsymbol{\eta}}(i); \hat{\boldsymbol{\mu}}(i)]$, and (ii) the Galerkin amplitudes and the advection variables as the input and output vectors $\mathbf{y}_{in}(i) = \hat{\mathbf{y}}(i) = [\hat{\boldsymbol{\eta}}(i); \hat{\boldsymbol{\mu}}(i); \hat{\mathbf{v}}(i)]$. In these cases, we feed both β and τ as parameters through the input channel, $\mathbf{p} = [\beta; \tau]$. However, neither of these options is inherently time delayed. Therefore, we design an ESN for time-delayed systems and refer to it as thermoacoustic ESN (T-ESN). During closed-loop prediction, we compute the delayed velocity at the flame location from past predictions of the velocity Galerkin amplitudes. We then feed the delayed velocity into the network as an input for the prediction of the next time step. One step of T-ESN is

$$u_f(i - N_\tau) = \sum_{j=1}^{n_g} \hat{\eta}_j(i - N_\tau) \cos(j\pi x_f), \quad (26a)$$

$$\mathbf{r}(i + 1) = (1 - \alpha)\mathbf{r}(i) + \alpha \tanh(\mathbf{W}_{in}[\mathbf{y}_{in}(i); u_f(i - N_\tau); \sigma_\beta(\beta - k_\beta)]) + \mathbf{W}\mathbf{r}(i), \quad (26b)$$

$$\hat{\mathbf{y}}(i + 1) = [\hat{\boldsymbol{\eta}}(i + 1); \hat{\boldsymbol{\mu}}(i + 1)] = \mathbf{W}_{out}\mathbf{r}(i + 1), \quad (26c)$$

where $N_\tau = \text{int}(\tau/\Delta t)$, which implies that the ESN handles delays that are integer multiples of the time step, as it is a discrete-time map. (To extend this, one can use interpolation to approximate the velocity at the desired time delay; however, this is beyond the scope of this paper.) A schematic of the T-ESN in closed-loop configuration is shown in Fig. 1.




FIG. 1. Schematic of the proposed thermoacoustic echo state network (T-ESN) in closed-loop configuration. T-ESN hard-constrains the physical knowledge about the nonlinear and time-delayed thermoacoustics.

The tangent linear problem for computing the gradient with respect to the parameters becomes

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \frac{d\tilde{\mathcal{J}}(\mathbf{r}(i))}{d\mathbf{r}(i)} \mathbf{Q}(i), \quad (27a)$$

$$\mathbf{Q}(i+1) = \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{p}} + \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i)} \mathbf{Q}(i) + \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i-N_\tau)} \mathbf{Q}(i-N_\tau), \quad (27b)$$

$$\mathbf{Q}(\leq 0) = \mathbf{0}. \quad (27c)$$

The adjoint of the time-delayed system is time advanced:

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \sum_{i=1}^N \mathbf{q}^{+\top}(i) \frac{\partial \mathbf{r}(i)}{\partial \mathbf{p}}, \quad (28a)$$

$$\mathbf{q}^+(i) = \frac{1}{N} \frac{d\tilde{\mathcal{J}}(\mathbf{r}(i))}{d\mathbf{r}(i)}^\top + \frac{\partial \mathbf{r}(i+1)^\top}{\partial \mathbf{r}(i)} \mathbf{q}^+(i+1) + \frac{\partial \mathbf{r}(i+1+N_\tau)^\top}{\partial \mathbf{r}(i)} \mathbf{q}^+(i+1+N_\tau), \quad (28b)$$

$$\mathbf{q}^+(N) = \frac{1}{N} \frac{d\tilde{\mathcal{J}}(\mathbf{r}(N))}{d\mathbf{r}(N)}^\top, \quad (28c)$$

$$\mathbf{q}^+(\geq N) = \mathbf{0}. \quad (28d)$$

The derivation of the time-delayed adjoint is provided in Appendix A. The time-delayed Jacobian is provided in Appendix B. (The tangent linear and adjoint equations for the initial condition sensitivity are found similarly by including the time-delayed or time-advanced Jacobian in Eqs. (18) and (19) in Sec. III, respectively.)

To obtain the sensitivity to time delay, we need

$$\frac{\partial \mathbf{r}(i+1)}{\partial \tau} = \frac{\partial \mathbf{r}(i+1)}{\partial u_f(i-\tau)} \frac{\partial u_f(i-\tau)}{\partial \tau}. \quad (29)$$

In the delayed formulation of the T-ESN, the delayed velocity $u_f(t-\tau)$ [Eq. (26a)] does not continuously depend on τ . We approximate the gradient $\partial u_f(t-\tau)/\partial \tau$ with a central finite

difference,

$$\frac{\partial u_f(i - \tau)}{\partial \tau} \approx \frac{u_f(i - (N_\tau + 1)) - u_f(i - (N_\tau - 1))}{2\Delta t}. \quad (30)$$

Motivated by thermoacoustic physics, we design the input weights matrix based on the following observations: (i) the acoustics are linear and the modes are sparsely connected, (ii) the heat release is nonlinear, and (iii) the heat-release strength β is only coupled with $u_f(t - \tau)$:

$$\mathbf{W}_{\text{in}} = \begin{bmatrix} \mathbf{w}_{1,1} & 0 & 0 & 0 & 0 & 0 \\ 0 & \mathbf{w}_{2,2} & 0 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathbf{w}_{2n_g, 2n_g} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{w}_f & \mathbf{w}_\beta \end{bmatrix}, \quad (31)$$

where $\mathbf{w}_{j,j} \in \mathbb{R}^{\text{int}(n_r/(2n_g+1))}$, $j = 1, \dots, 2n_g$, $\mathbf{w}_f \in \mathbb{R}^{\text{int}(n_r/(2n_g+1))}$, $\mathbf{w}_\beta \in \mathbb{R}^{\text{int}(n_r/(2n_g+1))}$. The tanh activation acts approximately linearly near zero, so the terms that linearly contribute to the dynamics have a smaller scaling, whereas the nonlinear terms have a larger scaling. Therefore, we scale the input weights that multiply the linear acoustics with one scaling, i.e., $\mathbf{w}_{j,j} = \sigma_{\text{in}} \tilde{\mathbf{w}}_{j,j}$, and the input weights that multiply the delayed velocity that contributes to the nonlinear heat-release term with another scaling, i.e., $\mathbf{w}_f = \sigma_f \tilde{\mathbf{w}}_f$, because we expect that $\sigma_f > \sigma_{\text{in}}$ ($\tilde{\mathbf{w}}_{j,j}$, $\tilde{\mathbf{w}}_f$, and \mathbf{w}_β are drawn from the uniform distribution $\sim \mathcal{U}(-1, 1)$ as before). While doing so we do not impose the linearity, but rather allow the machine to discover this via the optimization of the scaling terms. Separately, the parameter β is shifted and scaled by $\sigma_\beta(\beta - k_\beta)$. The scalings are determined by the hyperparameter search. In summary, we hard-constrain the physical knowledge about the thermoacoustics in the network architecture in two ways: (i) by making the network explicitly time delayed and (ii) by designing the input weights matrix.

VI. LEARNING THERMOACOUSTICS AND PARAMETER SENSITIVITIES

In this section, we learn the dynamics of the Rijke tube parametrized by the heat-release strength, β , and the time delay, τ . First, we focus on limit-cycle regimes and find that the T-ESN outperforms other parameter-aware ESN architectures in terms of generalizability over unseen regimes. Then, using the trained T-ESN, we infer the adjoint sensitivities of the acoustic energy to these parameters. We compare the time-series predictions and inferred adjoint sensitivities to the ground truth. The ground truth is obtained by integrating the original system described by the ODEs (22), and its adjoint equations (5), which requires the computation of the Jacobian of Eq. (22).

A. Training and validation

For training and validation, we create a data set from time delays, $\tau = \{0.1, 0.15, 0.2, 0.25, 0.3\}$, and heat-release strengths, $\beta = \{1.0, 2.0, 3.0, 4.0, 5.0\}$, giving a total of 25 thermoacoustic regimes. These parameters generate limit-cycle solutions except for $(\beta = 5.0, \tau = 0.15)$, which produces a quasiperiodic solution. We fix the flame location $x_f = 0.2$, modal damping coefficients $c_1 = 0.1$ and $c_2 = 0.06$, the number of Galerkin modes, $n_g = 4$, and the number of Chebyshev discretization, $n_c = 10$ [6]. With four Galerkin modes, we observe the rich nonlinear behavior associated with thermoacoustics. We generate the data by time-marching ODEs (22) using the `scipy.odeint` tool with a time step of 0.001, which we downsample to 0.01 for the ESN. We remove a transient of 200 time units, select the next 4 time units as washout and 8 time units as training data. (We found that increasing the training length further does not improve the performance, whereas it increases the computational effort.) We randomly split this data set into 20 regimes for training and 5 regimes for validation of the hyperparameters. During validation, we evaluate the closed-loop performance of five random realizations of the ESN, i.e., realizations of the randomly generated input and state matrices, over five trajectories, or folds, of length 4 time units (approximately two periods) starting

TABLE I. Optimal hyperparameters of (i) standard parameter-aware ESN (9), (ii) advection-augmented parameter-aware ESN (9), and (iii) thermoacoustic ESN (26) trained on $\beta = \{1.0, 2.0, 3.0, 4.0, 5.0\}$ and $\tau = \{0.1, 0.15, 0.2, 0.25, 0.3\}$ regimes. Refer to Eqs. (9) and (26) for the hyperparameters.

	ρ	σ_{in}	σ_f	σ_β	k_β	σ_τ	k_τ	α	λ
	0.01–1.0	0.01–2.0	0.01–2.0	0.01–2.0	−10.0–10.0	0.01–2.0	−1.0–1.0	0.01–1.0	10^{-6} – 10^{-1}
ESN (i)	0.0363	0.3476		0.0152	−5.3759	0.7741	−0.2063	0.0597	10^{-5}
ESN (ii)	0.0367	0.3293		0.0162	−5.4293	0.7319	0.72734	0.0627	10^{-5}
ESN (iii)	0.0125	0.0790	1.0501	0.1092	−10.0			0.0363	10^{-6}

from random time steps. As an error metric, we compute relative ℓ_2 error between the prediction, $\hat{\mathbf{y}}$, and the ground truth, \mathbf{y} :

$$\epsilon = \frac{\sqrt{\sum_{i=1}^{N_v} \sum_{j=1}^{n_y} (y_j(i) - \hat{y}_j(i))^2}}{\sqrt{\sum_{i=1}^{N_v} \sum_{j=1}^{n_y} y_j^2(i)}}, \quad (32)$$

where N_v is the number of validation steps, which we average over the number of folds, number of validation regimes, and number of realizations. We employ a Bayesian optimization scheme using the `skopt.gp_minimize` tool to determine the set of hyperparameters that have the minimal validation error [21].

We consider three networks: (i) the standard parameter-aware ESN (9) with the Galerkin amplitudes as the input and output vectors $\mathbf{y}_{\text{in}}(i) = \hat{\mathbf{y}}(i) = [\hat{\boldsymbol{\eta}}(i); \hat{\boldsymbol{\mu}}(i)]$, (ii) input and output vectors that we augment with the advection variable $\mathbf{y}_{\text{in}}(i) = \hat{\mathbf{y}}(i) = [\hat{\boldsymbol{\eta}}(i); \hat{\boldsymbol{\mu}}(i); \hat{\mathbf{v}}(i)]$, and (iii) thermoacoustic ESN (T-ESN) (26). The standard and the advection-augmented parameter-aware ESNs have both β and τ as explicit parameters in the input channel, i.e., in Eq. (9) $\mathbf{p} = [\beta; \tau]$. The T-ESN has β as a parameter in the input channel and τ is modeled as a time delay in the delayed velocity term in the input channel $u_f(t - \tau)$ [Eq. (26a)]. For comparison, we fix the reservoir size, i.e., number of reservoir state variables, n_r , to 1200 and connectivity of the reservoir state variables, i.e., the average number of nonzero elements per row in the state matrix \mathbf{W} , to 20. The optimal hyperparameters are provided in Table I. The first row of Table I shows the ranges from which the hyperparameters are sampled during the validation; ρ , σ_{in} , σ_f , σ_β , σ_τ , α , and λ are sampled uniformly on a \log_{10} scale, while k_β and k_τ are sampled uniformly.

B. Prediction performance

The computed adjoint sensitivities are only as accurate as the model we have. Therefore, the first task of the ESN is to accurately emulate the dynamics of regimes both seen and unseen during training. This includes two subtasks: (i) time-accurate prediction from a given initial condition and (ii) convergence to the correct attractor. The latter is necessary for long-term performance because we are interested in the sensitivity of long-time-averaged quantities. Figure 2 shows the short-term prediction performance of the standard, advection-augmented, and thermoacoustic ESNs over a grid of β and τ values. For a given regime, after a washout stage, the ESN is run in closed loop for prediction. This means that the network is forecasting the dynamics. The relative ℓ_2 error (32) between prediction and ground truth of $[\boldsymbol{\eta}; \boldsymbol{\mu}]$ over 20 time units is averaged over five folds starting from random time steps, and five realizations of the ESN. The standard parameter-aware ESN is not capable of sustaining the limit-cycle solutions for approximately 10 periods even for the training regimes, whereas the augmentation of the advection variables \mathbf{v} in the input and output significantly improves the performance. The T-ESN generalizes to unseen regimes. Owing to the nonlinearity in the system, the prediction performance decreases as we move away from the training regimes.




FIG. 2. Short-term closed-loop prediction performance of (a) standard parameter-aware ESN (9), (b) advection-augmented parameter-aware ESN (9), and (c) thermoacoustic ESN (26). Relative ℓ_2 error between prediction and ground truth of $[\eta; \mu]$ over 20 time units on a \log_{10} scale.

In Fig. 3, we show the performance of the T-ESN on (i) a limit-cycle regime seen during training ($\beta = 2.0$, $\tau = 0.25$) and (ii) an unseen limit-cycle regime ($\beta = 4.5$, $\tau = 0.12$). Based on Fig. 2, we choose regime (ii) as an example of an unseen regime with higher relative ℓ_2 error compared to the other limit-cycle regimes. In these figures, we assess the performance by (from left to right) short-term time-series prediction for 20 time units after washout, the probability density function (PDF), i.e., statistics, and amplitude spectrum of the acoustic energy E_{ac} [Eq. (25)]. For the PDF and amplitude spectrum, we feed the initial condition of zeros except $\eta_1(t_0) = 1.0$ in the washout stage, let the ESN evolve in closed loop, discard a transient of 200 time units, and use the next 1000 time units. This procedure and results verify that the ESN has learned the correct limit-cycle attractor because it converges to this set of states from an initial condition that is not in the training set.

C. Computation of adjoint sensitivities

We compare the adjoint sensitivities computed by the T-ESN following the procedure described in Sec. III with the adjoint sensitivities obtained from the original system described by Eq. (22). We refer to the latter as the ground truth, which is computed by deriving the Jacobian of ODEs (22) and

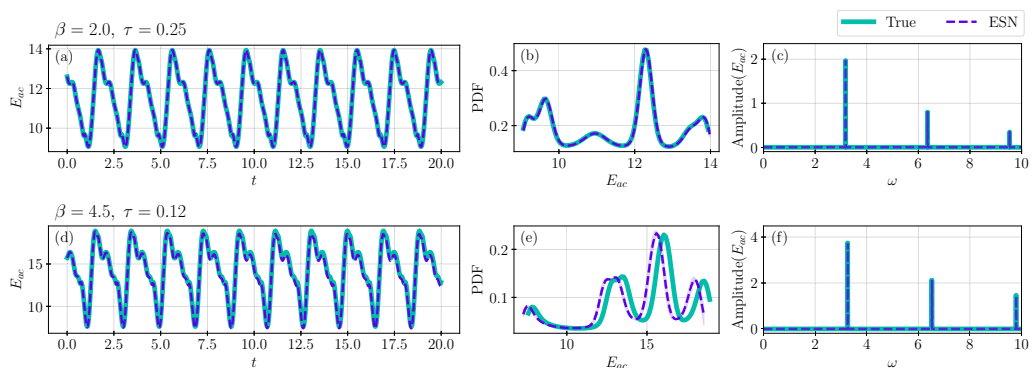


FIG. 3. Thermoacoustic echo state network learns limit-cycle regimes both seen and unseen during training. Short-term prediction, probability distribution function (PDF), and amplitude spectrum of the acoustic energy of [(a)–(c)] a limit-cycle regime seen during training ($\beta = 2.0$, $\tau = 0.25$) and [(d)–(f)] an unseen limit-cycle regime ($\beta = 4.5$, $\tau = 0.12$). In the PDFs, the mean of an ensemble of five realizations of the ESN is shown; one standard deviation from the mean is plotted in lighter color.




FIG. 4. Thermoacoustic echo state network infers sensitivity of the time-averaged acoustic energy, i.e., objective \mathcal{J} , to heat-release strength β . Shown is $d\mathcal{J}/d\beta$ at fixed (a) $\tau = 0.07$, (b) $\tau = 0.12$, (c) $\tau = 0.22$, and (d) $\tau = 0.32$. The mean of the sensitivity of an ensemble of five realizations of the ESN is shown; one standard deviation from the mean is plotted in lighter color.

integrating the adjoint Eqs. (5). The sensitivities are shown for varying β and τ values in Figs. 4 and 5, respectively. For ease of results, we show the sensitivity to β when β is varied and τ is fixed, and vice versa. In the washout stage, we feed the initial condition of zeros except $\eta_1(t_0) = 1.5$, assuming that no data are available from the original system for the washout. This value of $\eta_1(t_0)$ makes sure that the trajectories converge to the stable limit cycle and not the stable fixed point in the multistable regimes. After the washout and the transient, we solve the adjoint equations around a trajectory of 100 time units, for which we find that the adjoint sensitivity converges to a steady-state value. The learned adjoint sensitivities match closely with the ground truth, while the results are also robust to different realizations of the network. Even though the training data were generated from discrete values of (β, τ) on a grid (Fig. 2), we are able to obtain sensitivity information from a range of (β, τ) values that are unseen during the training using a single instance of the T-ESN. The smoothness of the sensitivity curve also indicates that the dynamics of the T-ESN varies smoothly with the parameters. The accuracy decreases as we move away from the training regimes, e.g., beyond $\beta = 5.0$ in the $d\mathcal{J}/d\beta$ plot at $\tau = 0.32$ [Fig. 4(d)], but the ESN captures the trend of the sensitivities.

Figures 4 and 5 also show that the T-ESN infers the existence of bifurcations to fixed-point solutions. The fixed point of the Rijke tube is the zero vector, i.e., the mean flow, which means that the sensitivity is zero for the fixed-point regimes, e.g., at fixed $\tau = 0.22$ [Fig. 4(c)] for $\beta < 0.45$. However, the parameter value at which the bifurcation occurs might have a small error compared to the true system. For example, at fixed $\tau = 0.22$ the Hopf bifurcation occurs at $\beta = 0.45$, whereas four out of five realizations of the T-ESN predict it as $\beta \approx 0.35$, and the remaining realization predicts it as $\beta \approx 0.30$.




FIG. 5. Thermoacoustic echo state network infers sensitivity of the time-averaged acoustic energy, i.e., objective \mathcal{J} , to time delay τ . Shown is $d\mathcal{J}/d\tau$ at fixed (a) $\beta = 1.25$, (b) $\beta = 2.5$, (c) $\beta = 3.75$, and (d) $\beta = 4.5$. The mean of the sensitivity of an ensemble of five realizations of the ESN is shown; one standard deviation from the mean is plotted in lighter color.




FIG. 6. Gradient-based optimization of the time-averaged acoustic energy starting from different initial (β, τ) using adjoint sensitivities computed by a single instance of the thermoacoustic echo state network. (a) Example optimization paths starting from different initial (β, τ) are plotted on top of the landscape of the objective functional \mathcal{J} . (b) The values \mathcal{J} takes during the iterations of optimization are plotted.

VII. GRADIENT-BASED OPTIMIZATION

We apply the inferred adjoint sensitivities to a gradient-based optimization framework. One realization of the T-ESN suffices for the computation of the gradients and optimization for all parameters. In this work, the goal of the optimization is to determine the set of parameters (β, τ) that minimizes the time-averaged acoustic energy (25). We restrict the domain to $\beta = [0.5, 5.5]$, $\tau = [0.05, 0.35]$, which prevents the parameters from taking negative (nonphysical) values, and the system to bifurcate to quasiperiodic or chaotic regimes. We wish to avoid this during gradient-based optimization because, first, the acoustic energy tends to display discontinuous jumps at the parameter values where the system bifurcates [6], which can lead to undefined gradients between different nonlinear regimes. Second, while the adjoint sensitivity can be computed for limit-cycle and quasiperiodic regimes for arbitrary integration times, for chaotic regimes the integration time is limited before the adjoint system becomes unstable, and the sensitivities diverge [6,48]. We deal with the sensitivity of the chaotic regimes separately in Sec. IX. Another consideration is the respective orders of magnitudes of the parameters and the sensitivities to them. From Figs. 4 and 5, β and τ as well as $d\mathcal{J}/d\beta$ and $d\mathcal{J}/d\tau$ have different orders of magnitude, i.e., $\beta \sim O(1)$, $\tau \sim O(10^{-1})$. The difference in the orders of magnitude requires the parameters and their respective gradients to be normalized, such that $\check{\beta} = \beta/\bar{\beta}$, $\check{\tau} = \tau/\bar{\tau}$, $d\mathcal{J}/d\check{\beta} = \bar{\beta}d\mathcal{J}/d\beta$, and $d\mathcal{J}/d\check{\tau} = \bar{\tau}d\mathcal{J}/d\tau$, where $\bar{(\cdot)}$ denotes the mean of the parameter in the domain considered [6]. The optimization is performed with the steepest-descent method using the normalized gradients with a step size of 0.2. (Since τ can only take discrete values, we round it according to the time-step size of the ESN before the next step of the optimization.) The optimization stops when (i) the acoustic energy is near zero, $\mathcal{J}_i < 10^{-4}$, or (ii) the relative change in acoustic energy is below a threshold $(\mathcal{J}_i - \mathcal{J}_{i-1})/\mathcal{J}_{i-1} < 10^{-4}$ where subscript i denotes the i th iteration step. Iterations of optimization starting from different initial (β, τ) are shown in Fig. 6. The optimal (β, τ) values are found to be in the fixed-point regimes, i.e., when the instability is suppressed. For the Rijke tube, the fixed point is the zero vector (mean flow), and as such the acoustic energy, i.e., the acoustic energy of the fluctuations from the mean flow, is equal to zero at these values. This result shows the significance of the ability of the ESN to infer fixed points. The process of optimization is illustrated for path 3 in Fig. 7. The optimization follows a path where the amplitude of the oscillations and so the acoustic energy decreases, for which the ESN accurately predicts the changing dynamics. Similar behavior is observed for other paths. An exceptional case is path 4, which passes through a saddle in the landscape of the objective




FIG. 7. Gradient-based optimization of time-averaged acoustic energy marked as path 3 in Fig. 6. [(a)–(c)] Iteration 0, ($\beta = 4.0$, $\tau = 0.25$); [(d)–(f)] iteration 2, ($\beta = 3.57$, $\tau = 0.17$); [(g)–(i)] iteration 4, ($\beta = 2.85$, $\tau = 0.1$); and [(j)–(l)] iteration 8, ($\beta = 1.64$, $\tau = 0.07$). From left to right: the first two pressure modes in the phase space, short-term time-series prediction of the acoustic energy, and long-term amplitude spectrum of the acoustic energy. The optimization follows a path along which the amplitude of the oscillations and the acoustic energy decrease, while the thermoacoustic echo state network accurately predicts the changing dynamics.

function near ($\beta = 4.5$, $\tau = 0.15$). The change in the sign of the gradient, as shown in Fig. 5(d) in the sensitivity plot for $\beta = 4.5$ around $\tau = 0.13$, causes the “zigzagging” path in this region.

VIII. ROBUSTNESS TO NOISE

If the data are acquired from real-life experiments, the measurements may be corrupted by noise. This is the case when the system equations are unknown and we wish to learn sensitivities from experimental data. To test the performance of the T-ESN under such a scenario, we consider zero-mean additive Gaussian noise with a standard deviation that is equal to 2%, 5%, and 10% of the standard deviation of the true data (these noise levels correspond to 33, 26, and 20 dB signal-to-noise ratios, respectively). We repeat the training on the noisy data using the hyperparameters chosen before (Table I) and investigate the effect of regularization. The regularization parameter, λ , can be determined by evaluating the validation error during hyperparameter selection. Figure 8 shows the statistics (mean and one standard deviation) of the validation error of the T-ESN for different levels of noise. The validation error is given by the mean relative ℓ_2 error (32) as described in Sec. VI A. (In the noisy case, we increase the number of evaluated trajectories to 20 as this helps with the convergence of the error statistics.) To obtain the error statistics, we repeat this calculation for 20 different realizations of noise applied to the training and validation data. Because the validation data




FIG. 8. Selection of the regularization parameter when the data are noisy. Shown is the validation error of the thermoacoustic echo state network on data with different levels of noise (indicated in the legend) when the training is performed with increasing regularization. The error statistics for 20 different realizations of noise are plotted.

contain noise, the irreducible error is the noise level. We find that there is a range of λ values that results in a similarly low validation error, and this optimal range of regularization increases with increasing levels of noise.

For the 5% noise level case, we show the performance of closed-loop prediction on a parameter regime ($\beta = 2.0$, $\tau = 0.25$), which is in the training set, in Fig. 9. The T-ESN successfully learns the limit-cycle attractor from noisy data and does not become unstable when the training data are noisy. We compute the adjoint sensitivities for the same test cases in Sec. VIC; these are shown in Figs. 10 and 11 for the heat-release strength β and time delay τ , respectively. The sensitivity computations are affected by the noisy training especially in small- τ regimes ($\tau \leq 0.1$) that are unseen during training, e.g., Figs. 10(a) and 10(b). These errors are alleviated by increasing the regularization.




FIG. 9. Thermoacoustic echo state network learns the true limit-cycle attractors of the original system from noisy data. Short-term closed-loop prediction on the training data from the ($\beta = 2.0$, $\tau = 0.25$) regime is shown (a) for the first pressure mode as time series and (b) for the first and second pressure modes in the phase space. (c) The amplitude spectrum is computed for the same time window as the training data (8 time units) for consistency in frequency resolution and shown in logarithmic scale. (d) Short-term prediction, (e) probability distribution function (PDF), and (f) amplitude spectrum of the acoustic energy are shown as in Fig. 3.




FIG. 10. Regularization with an optimal Tikhonov coefficient $\lambda = 10^{-3}$ improves inference of sensitivity to heat-release strength in the case of noise level 5% (see Fig. 8). Shown is $d\mathcal{J}/d\beta$ at fixed (a) $\tau = 0.07$, (b) $\tau = 0.12$, (c) $\tau = 0.22$, and (d) $\tau = 0.32$. The mean of the sensitivity of an ensemble of five realizations of the ESN is shown for different regularization coefficients; one standard deviation from the mean is plotted in lighter color.

However, using a very large regularization might decrease the performance, for example, in higher- τ regimes ($\tau \geq 0.25$) as shown in Figs. 9 and 11(b).

IX. BIFURCATIONS TO NONLINEAR REGIMES

In Sec. VI, we discuss that the ESN predicts bifurcations from limit-cycle oscillations to fixed points when the heat-release strength β is decreased. On the other hand, if β is increased, the thermoacoustic system displays bifurcations from limit-cycle oscillations to chaos via period doubling and quasiperiodic routes [6]. In this section, we investigate whether the ESN can also predict these bifurcations and attractors. We train and validate a separate T-ESN on data generated by larger values $\beta = \{6.0, 6.5, 7.0, 7.5, 8.0\}$ (for brevity, we refer to this data set as $\beta_{\text{train}} = [6.0-8.0]$). The reservoir size and connectivity are 1200 and 20, respectively, as before. The optimal hyperparameters found for this case are provided in Table II along with the hyperparameters for the previous case when the training was performed with data from $\beta = \{1.0, 2.0, 3.0, 4.0, 5.0\}$ (for brevity, we refer to this data set as $\beta_{\text{train}} = [1.0-5.0]$). The faster changing dynamics, e.g., chaotic behavior, in the larger- β regimes might be responsible for the increase in the spectral radius and leak rate.

We examine the T-ESN's prediction performance across different nonlinear regimes. First, in Figs. 12(a)–12(c) we show a limit-cycle regime in which period doubling occurs ($\beta = 7.5$, $\tau = 0.3$), and in Figs. 12(d)–12(f), a quasiperiodic regime ($\beta = 6.1$, $\tau = 0.2$). Second, we show two chaotic thermoacoustic regimes with different Lyapunov times: in Figs. 13(a)–13(f) a regime ($\beta = 7.6$, $\tau = 0.22$) with a Lyapunov time of 8.5 time units and in Figs. 13(g)–13(i) a regime ($\beta = 8.7$, $\tau = 0.23$) with a Lyapunov time of 3.9 time units. Lyapunov time (LT) is the characteristic




FIG. 11. Regularization with an optimal Tikhonov coefficient $\lambda = 10^{-3}$ improves inference of sensitivity to time delay τ in the case of noise level 5% (see Fig. 8). Shown is $d\mathcal{J}/d\tau$ at fixed (a) $\beta = 1.25$, (b) $\beta = 2.5$, (c) $\beta = 3.75$, and (d) $\beta = 4.5$. The mean of the sensitivity of an ensemble of five realizations of the ESN is shown for different regularization coefficients; one standard deviation from the mean is plotted in lighter color.

TABLE II. Optimal hyperparameters of the thermoacoustic ESN (26) trained on data from $\beta = \{1.0, 2.0, 3.0, 4.0, 5.0\}$ and $\beta = \{6.0, 6.5, 7.0, 7.5, 8.0\}$ regimes ($\tau = \{0.1, 0.15, 0.2, 0.25, 0.3\}$). Refer to Eqs. (26) for the hyperparameters.

	ρ	σ_{in}	σ_f	σ_β	k_β	α	λ
ESN $\beta_{\text{train}} = [1, 5]$	0.01–1.0	0.01–2.0	0.01–2.0	0.01–2.0	–10.0–10.0	0.01–1.0	10^{-6} – 10^{-1}
ESN $\beta_{\text{train}} = [6, 8]$	0.0125	0.0790	1.0501	0.1092	–10.0	0.0363	10^{-6}
ESN $\beta_{\text{train}} = [6, 8]$	0.1602	0.0122	1.296	0.3800	–4.008	0.6369	10^{-6}

timescale by which two nearby trajectories diverge exponentially in a chaotic system and is an indicator of the predictability of a system. The ESN trained on $\beta_{\text{train}} = [1.0\text{--}5.0]$ can qualitatively predict the bifurcations and capture the statistics and dominant frequencies correctly, especially for the nonchaotic regimes. Training the T-ESN on $\beta_{\text{train}} = [6.0\text{--}8.0]$, which encompasses the regimes that we are testing, increases accuracy of both short- and long-term predictions.

The ability of the T-ESNs trained on different β regimes to predict bifurcations is further shown in Fig. 14 in the bifurcation diagrams for fixed $\tau = 0.22$ [Figs. 14(a)–14(c)] and $\beta = 7.6$ [Figs. 14(d)–14(f)] cases. This ability of the parameter-aware ESN to predict bifurcations to quasiperiodic or chaotic regimes at unseen parameter values is in itself not an obvious result as the network only had access to limit-cycle regimes at lower β values for training and validation. When we increase β and a bifurcation occurs, the acoustic energy increases with a discontinuous jump, which indicates that the amplitude of the oscillations abruptly changes. Therefore, the prediction of these regimes becomes an extrapolation problem because the range of data we are interested in is different from the range of data the network was trained on. In this case, incorporating physical knowledge via the T-ESN architecture enables extrapolation performance. The bifurcation diagrams confirm that in the higher- β regimes, the T-ESN trained on $\beta_{\text{train}} = [1.0\text{--}5.0]$ can qualitatively capture the bifurcations to quasiperiodic and chaotic regimes; however, this is not as quantitatively accurate as the T-ESN trained on $\beta_{\text{train}} = [6.0\text{--}8.0]$. (The reverse is true for lower- β regimes.) When β is further increased beyond $\beta = 10.0$, the quantitatively accurate performance of the T-ESN trained on $\beta_{\text{train}} = [6.0\text{--}8.0]$ also starts diminishing.

In chaotic systems, due to the positive Lyapunov exponent, the adjoint system is unstable and the adjoint sensitivities diverge [6,48]. The ESN learns and replicates the underlying chaotic dynamics of the system under investigation (e.g., Ref. [18]). Because the ESN is a dynamical




FIG. 12. Thermoacoustic echo state network captures bifurcations to period doubling and quasiperiodic regimes. Short-term prediction, statistics, and amplitude spectrum of [(a)–(f)] an unseen limit-cycle regime where period doubling occurs ($\beta = 7.5$, $\tau = 0.3$) and [(g)–(l)] an unseen quasiperiodic regime ($\beta = 6.1$, $\tau = 0.2$).

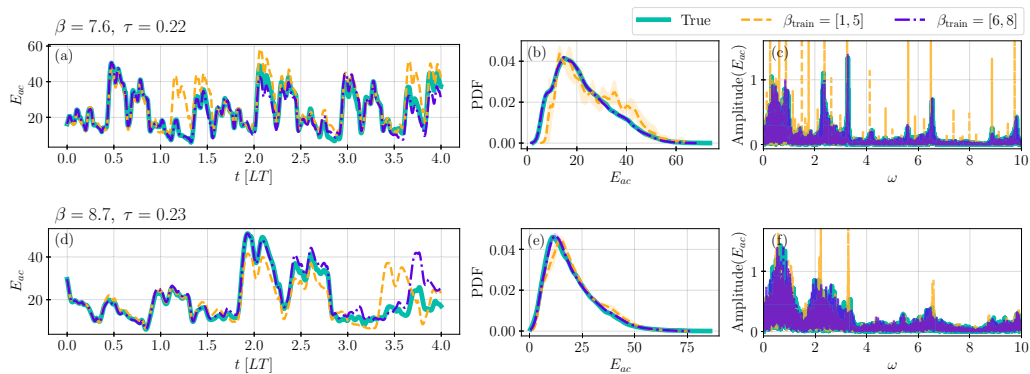


FIG. 13. Thermoacoustic echo state network captures bifurcations to chaotic regimes. Prediction, statistics, and amplitude spectrum of [(a)–(f)] an unseen chaotic regime ($\beta = 7.6$, $\tau = 0.22$) with a Lyapunov time of 8.5, and [(g)–(l)] an unseen chaotic regime ($\beta = 8.7$, $\tau = 0.23$) with a Lyapunov time of 3.9.

system [Eq. (10)], the adjoint sensitivity in chaotic regimes is unstable in the long term [6]. This means that we can only integrate adjoint equations for short time windows. Using the T-ESN trained on $\beta_{\text{train}} = [6.0\text{--}8.0]$, we compute the adjoint sensitivities of short-time trajectories for the ($\beta = 7.6$, $\tau = 0.22$) regime [shown in Figs. 13(a)–13(c)]. Figures 15(a) and 15(b) show the evolution of the adjoint sensitivities with integration time for one trajectory for an ensemble of five T-ESNs. For short integration times, the T-ESN predicts the chaotic sensitivities close to the original system, but diverges as the integration time increases due to the chaotic nature of the system. One approach to estimate the sensitivity of chaotic systems is by taking the mean of an ensemble of trajectories [48]. In Figs. 15(c) and 15(d) and Figs. 15(e) and 15(f), we show the probability distributions and the mean of the adjoint sensitivities for 2000 trajectories for integration times of 0.5 and 1 LT, respectively. In order to obtain the PDFs, we let the T-ESN run autonomously without any washout data from the original system. For an integration time of 0.5 LT, we can recover the PDF as well as the mean, while for 1.0 LT, low statistical moments are matched.




FIG. 14. Bifurcation diagrams obtained from the true system (top row), and from the thermoacoustic echo state network trained on $\beta_{\text{train}} = [1, 5]$ (middle row) and trained on $\beta_{\text{train}} = [6, 8]$ (bottom row). The peaks of the acoustic energy are plotted. [(a)–(c)] Varying β for fixed $\tau = 0.22$ and [(d)–(f)] varying τ for fixed $\beta = 7.6$.




FIG. 15. Computation of chaotic adjoint sensitivities using thermoacoustic echo state network. [(a), (b)] The adjoint sensitivities for a single trajectory at ($\beta = 7.6$, $\tau = 0.22$) are shown on the left. The probability distribution of the sensitivity is computed from 2000 short trajectories with integration time of [(c), (d)] 0.5 Lyapunov time and [(e), (f)] 1.0 Lyapunov time.

X. COMPUTATION OF SENSITIVITIES TO INITIAL CONDITIONS

The adjoint method allows us to efficiently compute sensitivities to initial conditions. In this section, we infer the sensitivity to initial conditions, i.e., $\mathbf{y}(0) = [\boldsymbol{\eta}(0); \boldsymbol{\mu}(0)]$, using the T-ESNs trained in the previous sections. The objective functional is the acoustic energy of the final flow state. The sensitivities to perturbations to the first velocity and pressure modes, η_1 and μ_1 , are shown in Fig. 16 for a trajectory from exemplary limit-cycle ($\beta = 4.5$, $\tau = 0.12$), quasiperiodic ($\beta = 6.1$, $\tau = 0.2$), and chaotic ($\beta = 7.6$, $\tau = 0.22$) regimes. To assess the accuracy of the estimates, we repeat this procedure for multiple independent trajectories. Because in gradient-based optimization we are interested in the gradient direction, we compute the angle between the sensitivity gradients of the original system and the T-ESN estimates as an indicator of the accuracy. The mean of the computed angles over 20 trajectories and 5 realizations of T-ESN are provided in Table III for different integration times. We find that T-ESN achieves high accuracy for all




FIG. 16. Computation of adjoint sensitivities to initial conditions using thermoacoustic echo state network. The adjoint sensitivity for a single trajectory at unseen [(a), (b)] limit-cycle ($\beta = 4.5$, $\tau = 0.12$), [(c), (d)] quasiperiodic ($\beta = 6.1$, $\tau = 0.2$), and [(e), (f)] chaotic ($\beta = 7.6$, $\tau = 0.22$) regimes. The sensitivity of the acoustic energy of the final flow state to perturbation to the first velocity mode η_1 (top row) and to the first pressure mode μ_1 (bottom row) are shown for increasing integration times.

TABLE III. Angle between the sensitivity gradients of the original system to initial conditions and the estimates from the thermoacoustic echo state network (T-ESN). The mean angle over 20 independent trajectories and 5 realizations of T-ESN is given in degrees.

		Limit cycle ($\beta = 4.5, \tau = 0.12$)	Quasiperiodic ($\beta = 6.1, \tau = 0.2$)		Chaotic ($\beta = 7.6, \tau = 0.22$)
	2.0	2.2	6.8	0.5 LT	8.0
Integration time	4.0	2.3	12.8	1.0 LT	15.5
	10.0	3.2	16.2	2.0 LT	34.3

integration times in the limit-cycle regimes and for short integration times in quasiperiodic and chaotic regimes.

XI. DISCUSSION

In this section, we offer a comparison of the proposed ESN-based adjoint solver with traditional, i.e., physics-based, adjoint solvers for gradient-based optimization. With the physics-based approach, we need (i) governing equations (plus boundary conditions and initial conditions), (ii) a numerical scheme to solve them, (iii) the linearization (tangent model), and (iv) its dual with respect to a chosen inner product (adjoint code). We need to derive or modify the adjoint code any time that any of the above is modified. On the other hand, in the ESN-based approach the adjoint code is derived once and does not change; i.e., the adjoint equations of the ESN described by Eqs. (13) and (19) are independent of the data and system under investigation. Similarly, the adjoint equations of the time-delayed ESN described by Eq. (28) are applicable to any system with a constant time delay. The ESN-based approach is also particularly attractive when we have experimental data for which the equations might be unknown (or partly unknown) and the data contain noise, because it learns the dynamics from observables.

As a data-driven approach, the ESN-based method relies on (i) having sufficient data that cover the solution space of the problem and (ii) training and validation of the ESN. When utilizing a parameter-aware ESN for the computation of parameter sensitivities, as the number of system parameters increases, more training regimes are required. This is in contrast with physics-based adjoint solvers, which are derived from the system’s equations that generalize over the parameter space. Thus, if the ESN-based adjoint solver is sought to replace a physics-based adjoint solver, the advantage can vary from case to case depending on the cost of data acquisition and simulation time. Because the network learns a continuous parametrization of the dynamics even if it is trained on data from a discrete set of parameters, the adjoint can be computed around any set of parameters. (With a continuous parametrization of the objective functional, a future research direction is to combine Bayesian optimization (e.g., Ref. [7]), which relies on sampling via the ESN, and adjoint-based ESN gradients.) The physics-based adjoint solver depends on the modeling assumptions, whereas the accuracy of the ESN-based adjoint solver depends on data and the hyperparameters. In high-dimensional systems, the ESN-based adjoint solver can be combined with a data-driven model reduction technique, e.g., autoencoding as in Ref. [26].

Finally, regarding computational costs, the training of the ESN is faster compared to other neural networks that are trained with stochastic gradient descent and back-propagation, and the T-ESN runs faster than the original solver when simulating the thermoacoustic system considered in this paper. Although this can be case dependent, in general the propagation of the ESN equations is fast thanks to the sparse weight matrices even though the reservoir state can be high dimensional (here, 1200). An optimization of the code for matrix multiplications on a GPU, e.g., using JAX [49], is scope for future work.

XII. CONCLUSION

Adjoint methods are established tools for gradient-based optimization. Although adjoint methods offer an accurate and efficient gradient computation, they are problem specific. In this work, we propose a strategy to infer adjoint sensitivities from data (observables). We showcase the method on a prototypical nonlinear time-delayed thermoacoustic system. This system is described by a nonlinear wave equation. The parameters of interest are the heat release strength and the time delay at the heat source location. By varying these parameters, the system displays fixed-point, limit-cycle, quasiperiodic, and chaotic solutions. The cost functional is the time-averaged acoustic energy, which we wish to minimize to suppress an instability.

First, we propose a parameter-aware echo state network (ESN). This network learns the parametrized dynamics; i.e., it infers how the dynamics change with varying parameters, thereby making predictions for parameters that are not used during training. Second, we derive the adjoint parameter-aware echo state network, which is used to compute the adjoint sensitivities of the physical system. Although the proposed framework is general and has been applied to prototypical low-dimensional dynamical systems without any modification of the standard architecture, we find that for the thermoacoustic system of interest embedding physical knowledge improves generalizability and requires less data. Therefore, third, we design a parameter-aware ESN motivated by the physics of thermoacoustics (thermoacoustic ESN, or T-ESN in short) by (i) making the ESN explicitly time delayed and (ii) designing the input matrix in consideration of acoustic physics. We extend the derivation of the adjoint to the time-delayed case. We find that the T-ESN accurately predicts the dynamics of different nonlinear regimes, and successfully infers the adjoint sensitivities. These results are validated against the ground truth generated from integration of the direct and adjoint equations of the original system. We optimize the acoustic energy with steepest gradient descent, in which the gradients are computed by only one T-ESN. Once the T-ESN has been trained, the algorithm runs autonomously; i.e., we do not run any additional simulations of the original model. As the optimization proceeds with the iterations, the acoustic energy and the amplitude of the oscillations decrease, finally reaching a fixed-point solution with zero acoustic energy. Fourth, we investigate the robustness to noisy data and chaotic regimes. The T-ESN can qualitatively predict bifurcations to nonlinear regimes beyond the training range of parameter values. Quantitative accuracy increases when the range of training parameters covers the regime in which we are interested. Finally, in addition to parameter sensitivities, we infer sensitivities to initial conditions. This work opens up possibilities for data-driven gradient-based optimization without developing code-specific adjoint solvers.

The code is available at [50].

ACKNOWLEDGMENTS

This research has received financial support from the ERC Starting Grant No. PhyCo 949388 and UKRI AI for Net Zero Grant No. EP/Y005619/1. L.M. is also grateful for the support from the grant EU-PNRR YoungResearcher TWIN ERC-PI_0000005.

APPENDIX A: ADJOINTS OF DISCRETE-TIME DYNAMICAL SYSTEMS WITH AND WITHOUT TIME DELAY

1. Adjoint of a discrete-time system

The following derivation of the adjoint sensitivity applies to any discrete map of choice that can be expressed as

$$\mathbf{x}(i) = \mathbf{f}(\mathbf{x}(i-1), \mathbf{p}), \tag{A1}$$

where $\mathbf{x}(i) \in \mathbb{R}^{n_x}$ is the state vector at time step i . Equation (A1) is rewritten as a constraint

$$\mathbf{F}(\mathbf{x}(i), \mathbf{x}(i-1), \mathbf{p}) := \mathbf{x}(i) - \mathbf{f}(\mathbf{x}(i-1), \mathbf{p}) = 0. \quad (\text{A2})$$

We consider a time-averaged quantity as the objective functional

$$\mathcal{J}(\mathbf{x}, \mathbf{p}) := \frac{1}{N} \sum_{i=1}^N \tilde{\mathcal{J}}(\mathbf{x}(i), \mathbf{p}). \quad (\text{A3})$$

Following Ref. [5], the problem of minimizing the objective functional (A3) under the constraints (A2) can be formulated as a Lagrangian optimization problem,

$$\mathcal{L} := \mathcal{J}(\mathbf{x}, \mathbf{p}) - \langle \mathbf{q}^+, \mathbf{F} \rangle, \quad (\text{A4})$$

where \mathbf{q}^+ are the Lagrange multipliers. The inner product $\langle \mathbf{q}^+, \mathbf{F} \rangle$ is given by

$$\langle \mathbf{q}^+, \mathbf{F} \rangle := \sum_{i=1}^N \mathbf{q}^{+\top}(i) \mathbf{F}(i), \quad (\text{A5})$$

where we denote $\mathbf{F}(i) = \mathbf{F}(\mathbf{x}(i), \mathbf{x}(i-1), \mathbf{p})$ for brevity. Because $d\mathcal{J}/d\mathbf{p} = d\mathcal{L}/d\mathbf{p}$, we seek the gradient $d\mathcal{L}/d\mathbf{p}$:

$$\frac{d\mathcal{L}}{d\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{p}} + \frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{x}(i)} \frac{d\mathbf{x}(i)}{d\mathbf{p}} - \sum_{i=1}^N \mathbf{q}^{+\top}(i) \left(\frac{\partial \mathbf{F}(i)}{\partial \mathbf{p}} + \frac{\partial \mathbf{F}(i)}{\partial \mathbf{x}(i)} \frac{d\mathbf{x}(i)}{d\mathbf{p}} + \frac{\partial \mathbf{F}(i)}{\partial \mathbf{x}(i-1)} \frac{d\mathbf{x}(i-1)}{d\mathbf{p}} \right). \quad (\text{A6})$$

Without loss of generality, we assume that the objective function does not explicitly depend on \mathbf{p} and therefore $\partial \tilde{\mathcal{J}}/\partial \mathbf{p}$ is zero. Writing out the terms of the summation, we have

$$\begin{aligned} \frac{d\mathcal{L}}{d\mathbf{p}} &= \sum_{i=1}^N -\mathbf{q}^{+\top}(i) \frac{\partial \mathbf{F}(i)}{\partial \mathbf{p}} \\ &+ \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(1))}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} - \mathbf{q}^{+\top}(1) \frac{\partial \mathbf{F}(1)}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} - \mathbf{q}^{+\top}(1) \frac{\partial \mathbf{F}(1)}{\partial \mathbf{x}(0)} \frac{d\mathbf{x}(0)}{d\mathbf{p}} \\ &+ \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(2))}{\partial \mathbf{x}(2)} \frac{d\mathbf{x}(2)}{d\mathbf{p}} - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(2)} \frac{d\mathbf{x}(2)}{d\mathbf{p}} - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} \\ &\vdots \\ &+ \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(N))}{\partial \mathbf{x}(N)} \frac{d\mathbf{x}(N)}{d\mathbf{p}} - \mathbf{q}^{+\top}(N) \frac{\partial \mathbf{F}(N)}{\partial \mathbf{x}(N)} \frac{d\mathbf{x}(N)}{d\mathbf{p}} - \mathbf{q}^{+\top}(N) \frac{\partial \mathbf{F}(N)}{\partial \mathbf{x}(N-1)} \frac{d\mathbf{x}(N-1)}{d\mathbf{p}}. \quad (\text{A7}) \end{aligned}$$

Because of Eq. (A1), the terms $\partial \mathbf{F}(i)/\partial \mathbf{x}(i)$ are equal to the identity matrix, and $d\mathbf{x}(0)/d\mathbf{p}$ is equal to zero. We choose the Lagrange multipliers to eliminate the terms $d\mathbf{x}(i)/d\mathbf{p}$, the size of which grows with the dimension of \mathbf{p} . For this purpose, we gather the terms multiplying $d\mathbf{x}(i)/d\mathbf{p}$, thereby obtaining the evolution equations (13) for the Lagrange multipliers, i.e., the adjoint variables,

$$\begin{aligned} \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(1))}{\partial \mathbf{x}(1)} - \mathbf{q}^{+\top}(1) - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(1)} &= 0 \\ \vdots & \\ \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(N))}{\partial \mathbf{x}(N)} - \mathbf{q}^{+\top}(N) &= 0, \quad (\text{A8}) \end{aligned}$$

where from Eq. (A1), $\partial \mathbf{F}(i)/\partial \mathbf{x}(i-1) = -\partial \mathbf{f}(i)/\partial \mathbf{x}(i-1)$.

2. Adjoint of a time-delayed discrete-time system

The discrete map for the time-delayed system can be expressed as

$$\mathbf{F}(\mathbf{x}(i), \mathbf{x}(i-1), \mathbf{x}(i-1-N_\tau), \mathbf{p}) = 0, \quad (\text{A9})$$

where N_τ is the time delay in discrete time steps. The Lagrangian \mathcal{L} is formulated similarly to Eq. (A4). The gradient $d\mathcal{L}/d\mathbf{p}$ is given by

$$\begin{aligned} \frac{d\mathcal{L}}{d\mathbf{p}} = & \frac{1}{N} \sum_{i=1}^N \frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{p}} + \frac{\partial \tilde{\mathcal{J}}}{\partial \mathbf{x}(i)} \frac{d\mathbf{x}(i)}{d\mathbf{p}} - \sum_{i=1}^N \mathbf{q}^{+\top}(i) \left(\frac{\partial \mathbf{F}(i)}{\partial \mathbf{p}} + \frac{\partial \mathbf{F}(i)}{\partial \mathbf{x}(i)} \frac{d\mathbf{x}(i)}{d\mathbf{p}} + \frac{\partial \mathbf{F}(i)}{\partial \mathbf{x}(i-1)} \frac{d\mathbf{x}(i-1)}{d\mathbf{p}} \right. \\ & \left. + \frac{\partial \mathbf{F}(i)}{\partial \mathbf{x}(i-1-N_\tau)} \frac{d\mathbf{x}(i-1-N_\tau)}{d\mathbf{p}} \right). \end{aligned} \quad (\text{A10})$$

Writing out the terms of the summation, we have

$$\begin{aligned} \frac{d\mathcal{L}}{d\mathbf{p}} = & \sum_{i=1}^N -\mathbf{q}^{+\top}(i) \frac{\partial \mathbf{F}(i)}{\partial \mathbf{p}} \\ & + \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(1))}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} - \mathbf{q}^{+\top}(1) \frac{\partial \mathbf{F}(1)}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} - \mathbf{q}^{+\top}(1) \frac{\partial \mathbf{F}(1)}{\partial \mathbf{x}(0)} \frac{d\mathbf{x}(0)}{d\mathbf{p}} \\ & - \mathbf{q}^{+\top}(1) \frac{\partial \mathbf{F}(1)}{\partial \mathbf{x}(-N_\tau)} \frac{d\mathbf{x}(-N_\tau)}{d\mathbf{p}} \\ & + \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(2))}{\partial \mathbf{x}(2)} \frac{d\mathbf{x}(2)}{d\mathbf{p}} - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(2)} \frac{d\mathbf{x}(2)}{d\mathbf{p}} - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} \\ & - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(1-N_\tau)} \frac{d\mathbf{x}(1-N_\tau)}{d\mathbf{p}} \\ & \vdots \\ & + \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(2+N_\tau))}{\partial \mathbf{x}(2+N_\tau)} \frac{d\mathbf{x}(2+N_\tau)}{d\mathbf{p}} - \mathbf{q}^{+\top}(2+N_\tau) \frac{\partial \mathbf{F}(2+N_\tau)}{\partial \mathbf{x}(2+N_\tau)} \frac{d\mathbf{x}(2+N_\tau)}{d\mathbf{p}} \\ & - \mathbf{q}^{+\top}(2+N_\tau) \frac{\partial \mathbf{F}(2+N_\tau)}{\partial \mathbf{x}(1+N_\tau)} \frac{d\mathbf{x}(1+N_\tau)}{d\mathbf{p}} - \mathbf{q}^{+\top}(2+N_\tau) \frac{\partial \mathbf{F}(2+N_\tau)}{\partial \mathbf{x}(1)} \frac{d\mathbf{x}(1)}{d\mathbf{p}} \\ & \vdots \\ & + \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(N))}{\partial \mathbf{x}(N)} \frac{d\mathbf{x}(N)}{d\mathbf{p}} - \mathbf{q}^{+\top}(N) \frac{\partial \mathbf{F}(N)}{\partial \mathbf{x}(N)} \frac{d\mathbf{x}(N)}{d\mathbf{p}} - \mathbf{q}^{+\top}(N) \frac{\partial \mathbf{F}(N)}{\partial \mathbf{x}(N-1)} \frac{d\mathbf{x}(N-1)}{d\mathbf{p}} \\ & - \mathbf{q}^{+\top}(N) \frac{\partial \mathbf{F}(N)}{\partial \mathbf{x}(N-1-N_\tau)} \frac{d\mathbf{x}(N-1-N_\tau)}{d\mathbf{p}}. \end{aligned} \quad (\text{A11})$$

Because of Eq. (A1), the terms $\partial \mathbf{F}(i)/\partial \mathbf{x}(i)$ are equal to the identity matrix, and due to the initial conditions $d\mathbf{x}(i \leq 0)/d\mathbf{p}$ are equal to zero. Eliminating $d\mathbf{x}/d\mathbf{p}$ terms, the evolution of the adjoint variables is given by

$$\begin{aligned} \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(1))}{\partial \mathbf{x}(1)} - \mathbf{q}^{+\top}(1) - \mathbf{q}^{+\top}(2) \frac{\partial \mathbf{F}(2)}{\partial \mathbf{x}(1)} - \mathbf{q}^{+\top}(2+N_\tau) \frac{\partial \mathbf{F}(2+N_\tau)}{\partial \mathbf{x}(1)} &= 0, \\ \vdots \\ \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{x}(N))}{\partial \mathbf{x}(N)} - \mathbf{q}^{+\top}(N) &= 0. \end{aligned} \quad (\text{A12})$$

We set $\mathbf{q}^+(i > N) = 0$ because there are no terms with $d\mathbf{x}(i > N)/d\mathbf{p}$.

APPENDIX B: GRADIENTS OF THE ECHO STATE NETWORK

The gradients associated with the ESN are provided here. We first define

$$\tilde{\mathbf{r}}(i) = \frac{\mathbf{r}(i+1) - (1 - \alpha)\mathbf{r}(i)}{\alpha}, \quad (\text{B1})$$

which is used to compute the derivative of the tanh term in the ESN equations [e.g., Eq. (6a)]. The Jacobian and the partial derivative with respect to the system's parameters are

$$\frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i)} = (1 - \alpha)\mathbf{I}_{n_r \times n_r} + \alpha \text{diag}(1 - \tilde{\mathbf{r}}^2(i))(\mathbf{W}_{\text{in}}^y \mathbf{W}_{\text{out}} + \mathbf{W}), \quad (\text{B2})$$

$$\frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{p}} = \alpha \text{diag}(1 - \tilde{\mathbf{r}}^2(i))\mathbf{W}_{\text{in}}^p \text{diag}(\boldsymbol{\sigma}_p), \quad (\text{B3})$$

where \mathbf{I} denotes the identity matrix, and $\text{diag}(\cdot)$ denotes a diagonal matrix that has (\cdot) as its diagonal. For the thermoacoustic ESN, we have the following gradients:

$$\frac{\partial \mathbf{r}(i+1)}{\partial u_f(i - N_\tau)} = \alpha \text{diag}(1 - \tilde{\mathbf{r}}^2(i))\mathbf{w}_f \quad (\text{B4})$$

$$\frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i - N_\tau)} = \frac{\partial \mathbf{r}(i+1)}{\partial u_f(i - N_\tau)} \Psi \mathbf{W}_{\text{out}}, \quad (\text{B5})$$

where $\Psi = [\cos(\pi x_f); \cos(2\pi x_f); \dots; \cos(n_g \pi x_f); \mathbf{0}_{n_g}]^\top$ contains the Galerkin modes associated with velocity and masks the Galerkin modes associated with pressure.

-
- [1] J. E. Peter and R. P. Dwight, Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches, *Comput. Fluids* **39**, 373 (2010).
- [2] D. Sipp, O. Marquet, P. Meliga, and A. Barbagallo, Dynamics and control of global instabilities in open-flows: A linearized approach, *Appl. Mech. Rev.* **63**, 030801 (2010).
- [3] P. Luchini and A. Bottaro, Adjoint equations in stability analysis, *Annu. Rev. Fluid Mech.* **46**, 493 (2014).
- [4] S. Camarri, Flow control design inspired by linear stability analysis, *Acta Mech.* **226**, 979 (2015).
- [5] L. Magri, Adjoint methods as design tools in thermoacoustics, *Appl. Mech. Rev.* **71**, 020801 (2019).
- [6] F. Huhn and L. Magri, Stability, sensitivity and optimisation of chaotic acoustic oscillations, *J. Fluid Mech.* **882**, A24 (2020).
- [7] F. Huhn and L. Magri, Gradient-free optimization of chaotic acoustics with reservoir computing, *Phys. Rev. Fluids* **7**, 014402 (2022).
- [8] M. B. Giles and N. A. Pierce, An introduction to the adjoint approach to design, *Flow Turbul. Combust.* **65**, 393 (2000).
- [9] M. D. Gunzburger, *Perspectives in Flow Control and Optimization* (Society for Industrial and Applied Mathematics, Philadelphia, 2002).
- [10] F. Giannetti and P. Luchini, Structural sensitivity of the first instability of the cylinder wake, *J. Fluid Mech.* **581**, 167 (2007).
- [11] A. Griewank and A. Walther, Algorithm 799: Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation, *ACM Trans. Math. Software* **26**, 19 (2000).
- [12] O. Pironneau, On optimum design in fluid mechanics, *J. Fluid Mech.* **64**, 97 (1974).
- [13] A. Jameson, Aerodynamic design via control theory, *J. Sci. Comput.* **3**, 233 (1988).
- [14] A. Jameson, L. Martinelli, and N. Pierce, Optimum aerodynamic design using the Navier-Stokes equations, *Theor. Comput. Fluid Dyn.* **10**, 213 (1998).
- [15] A. Jameson, Re-engineering the design process through computation, *J. Aircr.* **36**, 36 (1999).

- [16] M. Lukoševičius, A practical guide to applying echo state networks, in *Neural Networks: Tricks of the Trade*, edited by G. Montavon, G. B. Orr, and K.-R. Müller (Springer, Berlin, 2012), Vol. 7700, pp. 659–686.
- [17] L. Grigoryeva and J.-P. Ortega, Echo state networks are universal, *Neural Networks* **108**, 495 (2018).
- [18] G. Margazoglou and L. Magri, Stability analysis of chaotic systems from data, *Nonlinear Dyn.* **111**, 8799 (2023).
- [19] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach, *Phys. Rev. Lett.* **120**, 024102 (2018).
- [20] N. Doan, W. Polifke, and L. Magri, Physics-informed echo state networks, *J. Comput. Sci.* **47**, 101237 (2020).
- [21] A. Racca and L. Magri, Robust optimization and validation of echo state networks for learning chaotic dynamics, *Neural Networks* **142**, 252 (2021).
- [22] Z. Lu, J. Pathak, B. Hunt, M. Girvan, R. Brockett, and E. Ott, Reservoir observers: Model-free inference of unmeasured variables in chaotic systems, *Chaos* **27**, 041102 (2017).
- [23] N. A. K. Doan, W. Polifke, and L. Magri, Learning hidden states in a chaotic system: A physics-informed echo state network approach, in *Computational Science—ICCS 2020*, edited by V. V. Krzhizhanovskaya, G. Závodszy, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, and J. Teixeira (Springer International Publishing, Cham, 2020), Vol. 12142, pp. 117–123.
- [24] J. Pathak, Z. Lu, B. R. Hunt, M. Girvan, and E. Ott, Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data, *Chaos* **27**, 121102 (2017).
- [25] A. Nóvoa, A. Racca, and L. Magri, Inferring unknown unknowns: Regularized bias-aware ensemble Kalman filter, *Comput. Methods Appl. Mech. Eng.* **418**, 116502 (2024).
- [26] A. Racca, N. A. K. Doan, and L. Magri, Predicting turbulent dynamics with the convolutional autoencoder echo state network, *J. Fluid Mech.* **975**, A2 (2023).
- [27] R. Xiao, L.-W. Kong, Z.-K. Sun, and Y.-C. Lai, Predicting amplitude death with machine learning, *Phys. Rev. E* **104**, 014205 (2021).
- [28] D. E. Ozan and L. Magri, Adjoint sensitivities of chaotic flows without adjoint solvers: A data-driven approach, in *Computational Science—ICCS 2024*, edited by L. Franco, C. De Mulatier, M. Paszynski, V. V. Krzhizhanovskaya, J. J. Dongarra, and P. M. A. Sloot (Springer Nature, Cham, 2024), Vol. 14836, pp. 345–352.
- [29] T. C. Lieuwen and V. Yang, *Combustion Instabilities In Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modeling* (American Institute of Aeronautics and Astronautics, Reston, VA, 2006).
- [30] F. Culick and P. Kuentzmann, Unsteady motions in combustion chambers for propulsion systems, NATO RTO-AG-AVT-039, AGARDograph (2006).
- [31] Rayleigh, The explanation of certain acoustical phenomena 1, *Nature (London)* **18**, 319 (1878).
- [32] L. Magri, M. P. Juniper, and J. P. Moeck, Sensitivity of the Rayleigh criterion in thermoacoustics, *J. Fluid Mech.* **882**, R1 (2020).
- [33] L. Kabiraj, R. I. Sujith, and P. Wahi, Bifurcations of self-excited ducted laminar premixed flames, *J. Eng. Gas Turbines Power* **134**, 031502 (2012).
- [34] M. P. Juniper and R. Sujith, Sensitivity and nonlinearity of thermoacoustic oscillations, *Annu. Rev. Fluid Mech.* **50**, 661 (2018).
- [35] 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).
- [36] 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).
- [37] G. A. Mensah and J. P. Moeck, Acoustic damper placement and tuning for annular combustors: An adjoint-based optimization study, *J. Eng. Gas Turbines Power* **139**, 061501 (2017).
- [38] J. G. Aguilar and M. P. Juniper, Adjoint methods for elimination of thermoacoustic oscillations in a model annular combustor via small geometry modifications, in *Volume 4A: Combustion, Fuels, and Emissions* (American Society of Mechanical Engineers, Oslo, Norway, 2018), p. V04AT04A054.

- [39] A. Orchini, G. Rigas, and M. P. Juniper, Weakly nonlinear analysis of thermoacoustic bifurcations in the Rijke tube, *J. Fluid Mech.* **805**, 523 (2016).
- [40] G. Margazoglou and L. Magri, Data-driven stability analysis of a chaotic time-delayed system, in *Computational Science–ICCS 2023*, edited by J. Mičkyška, C. De Mulatier, M. Paszynski, V. V. Krzhizhanovskaya, J. J. Dongarra, and P. M. Sloot (Springer Nature, Cham, 2023), Vol. 10476, pp. 406–413.
- [41] H. Jaeger and H. Haas, Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication, *Science* **304**, 78 (2004).
- [42] M. Roy, S. Mandal, C. Hens, A. Prasad, N. V. Kuznetsov, and M. Dev Shrimali, Model-free prediction of multistability using echo state network, *Chaos* **32**, 101104 (2022).
- [43] M. P. Juniper, Triggering in the horizontal Rijke tube: Non-normality, transient growth and bypass transition, *J. Fluid Mech.* **667**, 272 (2011).
- [44] 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).
- [45] 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).
- [46] M. A. Heckl, Nonlinear acoustic effects in the Rijke tube, *Acustica* **72**, 63 (1990).
- [47] D. E. Ozan and L. Magri, Hard-constrained neural networks for modeling nonlinear acoustics, *Phys. Rev. Fluids* **8**, 103201 (2023).
- [48] D. J. Lea, M. R. Allen, and T. W. N. Haine, Sensitivity analysis of the climate of a chaotic system, *Tellus A* **52**, 523 (2000).
- [49] D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer, Machine learning–accelerated computational fluid dynamics, *Proc. Natl. Acad. Sci. USA* **118**, e2101784118 (2021).
- [50] <https://github.com/MagriLab/Adjoint-ESN>.