Reduced order modeling of parametrized systems through autoencoders and SINDy approach: continuation of periodic solutions

Paolo Conti[a,], Giorgio Gobat[a], Stefania Fresca[b], Andrea Manzoni[b], Attilio Frangi[a]

[a]*Department of Civil Engineering, Politecnico di Milano*
[b]*MOX – Department of Mathematics, Politecnico di Milano*

Abstract

Highly accurate simulations of complex phenomena governed by partial differential equations (PDEs) typically require intrusive methods and entail expensive computational costs, which might become prohibitive when approximating steady-state solutions of PDEs for multiple combinations of control parameters and initial conditions. Therefore, constructing efficient reduced order models (ROMs) that enable accurate but fast predictions, while retaining the dynamical characteristics of the physical phenomenon as parameters vary, is of paramount importance. In this work, a data-driven, non-intrusive framework which combines ROM construction with reduced dynamics identification, is presented. Starting from a limited amount of full order solutions, the proposed approach leverages autoencoder neural networks with parametric sparse identification of nonlinear dynamics (SINDy) to construct a low-dimensional dynamical model. This model can be queried to efficiently compute full-time solutions at new parameter instances, as well as directly fed to continuation algorithms. These aim at tracking the evolution of periodic steady-state responses as functions of system parameters, avoiding the computation of the transient phase, and allowing to detect instabilities and bifurcations. Featuring an explicit and parametrized modeling of the reduced dynamics, the proposed data-driven framework presents remarkable capabilities to generalize with respect to both time and parameters. Applications to structural mechanics and fluid dynamics problems illustrate the effectiveness and accuracy of the proposed method.

Keywords: Nonlinear dynamics, reduced order modeling, data-driven methods, autoencoder neural networks, sparse identification of nonlinear dynamics

1. Introduction

The ever-increasing resources and computational power allow nowadays to simulate very complex physical phenomena more and more accurately, even in multi-scale and multi-physics scenarios. However, the solution of parametrized, time-dependent systems of partial differential equations (PDEs) by means of full order models (FOMs) – such as the finite element method – may clash with time and computational budget restrictions. Moreover, using FOMs to explore different scenarios with varying initial conditions and parameter combinations might be a computationally prohibitive task, or even infeasible in several practical applications. Differently from the problem of estimating output quantities of interest that depend on the solution of the differential problem, the computation of the whole solution field is intrinsically high-dimensional, with additional difficulties related to the nonlinear and time-dependent nature of the problem. All these reasons drive the search of efficient, but accurate, reduced order models (ROMs). Among these, the reduced basis method [56, 33, 4] is a very well-known approach, exploiting, e.g., proper orthogonal decomposition (POD) to build a reduced space, either global or local [1, 55], to approximate the solution of the problem. However, despite their accuracy and mathematical, these techniques are in general intrusive [25].

Among machine and deep learning techniques widely used to build surrogate models or emulators to the solution of parametrized, nonlinear, time-dependent system of PDEs, autoencoder (AE) neural networks [27] have recently become a popular strategy because they allow to non-intrusively reduce dimensionality and unveil latent features directly from data streams, without accessing the FOM operators [26, 43, 48, 37]. Their success is due to the expressiveness capacity of neural networks [11, 44, 34], which enables outstanding performances in nonlinear compression and great flexibility in identifying coordinate transformations [46].

Recently proposed techniques consider convolutional AEs, suitably coupled to deep feedforward neural networks in order to simultaneously perform dimensionality reduction and learn the parameter-to-solution map, as in the case of deep learning-based ROMs (DL-ROMs) [18, 17], and their enhanced version through POD (POD-DL-ROMs) [21]. These techniques outperform (in some cases, also by several orders of magnitude) classical ROMs in terms of both the dimension of the reduced problem and the query time for predicting the solution at unseen instances. The resulting ROMs are nonintrusive and low-dimensional, ultimately enabling real-time simulations in complex physical scenarios; see, e.g., [19] for a possible application to the real-time simulation of the mechanical behavior of micro-electro-mechanical systems (MEMS). We highlight that this approach relies on the idea that the system dynamics lies on a low dimensional invariant manifold, setting a clear parallel with the Direct Parametrization of Invariant Manifolds approach preconized, e.g., in [60, 52, 59]; see also [24] for further details.

Nonetheless, this approach relies on the underlying physical model only when computing FOM snapshots required for training the deep neural networks involved, but not when querying the parameter-to-solution map for unseen parameter values at testing stage. Moreover, although showing impressive efficiency during both the training and the testing phases, and providing reliable predictions on unseen scenarios in the time-parameter range where training data are generated, POD-DL-ROMs show some limitations in terms of interpretability and generalization capability outside the observed parametric-temporal domain.

A possible alternative approach to learn the reduced space dynamics is represented by the so-called Sparse Identification of Dynamical Systems (SINDy) method, a sparse regression technique initially introduced for learning dynamics from time-series data, and then further generalized to address several features [35, 5, 6, 29]. In the past few years, SINDy has been widely applied to identify models of fluid flows, convection phenomena, structural models, and many others. The appeal of the SINDy approach is the generation, from a dictionary of pre-defined (analytical) functions, an explicit ROM in the form of a system of ordinary differential equations (ODEs), that can improve the interpretability of the latent space dynamics thanks to suitable sparsity constraints. Regarding time integration, as well as the computation of derived output quantities of interest, the identified dynamical system can be then treated with standard numerical tools. For instance, when the transient dynamics is relevant, or when steady-state regimes are not required, the obvious choice is represented by time-marching techniques. However, this might not be the optimal choice when periodic solutions are needed. These latter arise in many practical applications, ranging from aircraft design and simulation [2], chemistry [16] to MEMS [51], just to mention some instances. Periodicity conditions can be enforced directly in the formulation of the dynamical system and ad-hoc numerical methods can be applied like, e.g., the Harmonic Balance [40] or collocation techniques [15]. While their computational burden is usually prohibitive for large-scale FOMs [13, 51], these techniques become much more appealing when used in conjunction with ROM strategies.

Even if the SINDy approach has been so far mainly applied to rather low-dimensional problems, Champion et al. in [8] have fostered its use in combination with AE neural networks, laying the foundation for a unified framework for dimensionality reduction and system identification. This approach with recent extensions [3, 28, 39, 7, 23], however, does not encompass the possibility to treat parameters and forcing dependencies. This inevitably precludes the possibility of drawing efficiently a complete portrait of the dynamics of the observed system in multiple scenarios. In this respect, the SINDy method offers a great flexibility in view of incorporating parameters and forcing terms in the identified ODE system [6]. The major benefit of identifying an explicit parameter-dependency in the latent system is represented by the possibility to use *continuation* algorithms to track the evolution of periodic responses as a function of system parameters avoiding the computation of transients, which represents a significant computational advantage for long transient systems or for real-time applications. Moreover, these algorithms allow to identify, in a straightforward way, instabilities and bifurcations which might occur in the latent dynamics and result in macroscopic changes of the steady-state behavior of the physical system. For example, variations of input parameters can result in the development of vorticity in fluids or resonance phenomena in mechanical structures. In practical applications, these behaviors can significantly affect the operation of the observed systems, therefore it is essential to estimate at which parameter configurations these regime transitions occur.

In the present work, we propose an extension of the AE+SINDy approach presented in [8] to efficiently construct reduced order approximations of parametrized PDE solutions, as well as to discover their underlying dynamics and parametric dependencies, starting from a limited number of FOM snapshots. Our method hinges upon a data-driven, non-intrusive, ROM construction employing POD and AEs – the one provided by

the so-called POD-DL-ROMs [21] – together with reduced dynamics identification through a parametrized SINDy approach. The proposed framework thus leverages recent SINDy extensions with the aforementioned POD-DL-ROM reduction. In contrast to recent works in the same direction [22, 36], the main appeal of our method is that, by representing the latent dynamics as a parametrized system of ODEs, we can naturally embed continuation algorithms in the overall framework, thus providing a comprehensive description of the dynamics of the physical system.

The proposed method is first applied to a structural mechanics problem dealing with a straight beam MEMS resonator, which is excited at different forcing amplitudes and frequencies. We show how the proposed strategy allows to construct a ROM which manages both to generalize with respect to parameters and to accurately forecast the solution over long term horizons, even reaching the *steady-state* response for which no measurement is provided. In addition, the nonlinear hardening behaviour of the beam is portrayed by means of continuation algorithms directly from the reduced identified system. Next, the problem of a fluid flow past a cylinder is considered, employing the proposed method to efficiently approximate the fluid velocity and pressure at different regimes as the Reynolds number varies, as well as to empirically identify, through continuation algorithms, the bifurcation value of the transition from laminar to unsteady behaviour.

Besides the considered examples, the AE+SINDy method can be applied straightforwardly to any time series which tends to a periodic regime. Moreover, although this work focuses on periodic orbits, the entire framework could be also extended to handle quasi-periodic solutions by employing continuation techniques which are suitable for these regimes [32], or, more in general, to any time series that tends towards a limit cycle – which is the only requirement for continuation algorithms to work [41].

The paper is structured as follows. In Sect. 2 we detail the structure of the method and we present how to use it for the efficient approximation of parametrized PDE solutions as well as for dynamic analysis using both time-marching schemes and continuation algorithms. In Sect. 3 we present and discuss its performance on the two aforementioned numerical examples, finally drawing some concluding remarks in Sect. 4.

2. Reduced order modeling with AE and SINDy

2.1. Problem setup

Let us consider the following general form of a dynamical system

$$\begin{cases} \dot{\mathbf{x}}(t;\boldsymbol{\beta}) = \mathbf{f}(t,\mathbf{x}(t;\boldsymbol{\beta});\boldsymbol{\beta}), & t \in (0,T), \\ \mathbf{x}(0;\boldsymbol{\beta}) = \mathbf{x}_0, \end{cases} \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^N$ is the state of the system, $\dot{\mathbf{x}}$ its time derivative, t the time, $\mathbf{x}_0$ the initial state, while $\mathbf{f}$ is the function that defines the dynamical evolution of the physical system. Here, $\boldsymbol{\beta} \in \mathbb{R}^p$ represents the vector collecting p (possibly) time dependent parameters and/or forcing terms. System (1) consists of a set of ordinary differential equations (ODEs), whose dimensionality N represents the number of degrees of freedom associated to the space discretization technique (such as, for instance, finite element, finite volume or isogeometric analysis methods) applied to a system of partial differential equations (PDEs) which defines the physical system. Therefore, the size N is typically extremely large and the system (1) is denoted as full order model (FOM).

The goal is to learn the dynamics $\mathbf{f}$ from a set of time histories of the state $\mathbf{x}$ and its time-derivative $\dot{\mathbf{x}}$ for different instances of $\boldsymbol{\beta}$. In general, $\dot{\mathbf{x}}$ can be either computed directly or approximated numerically from $\mathbf{x}$. Snapshots of the time histories are stacked in matrices

$$\mathbf{X} = \left[\ \mathbf{x}(t_1;\boldsymbol{\beta}_1) \mid \mathbf{x}(t_2;\boldsymbol{\beta}_1) \mid \ldots \mid \mathbf{x}(t_{N_t};\boldsymbol{\beta}_1) \mid \mathbf{x}(t_1;\boldsymbol{\beta}_2) \mid \ldots \mid \mathbf{x}(t_{N_t};\boldsymbol{\beta}_2) \mid \ldots \mid \mathbf{x}(t_{N_t};\boldsymbol{\beta}_{N_\beta})\ \right]^T,$$

$$\dot{\mathbf{X}} = \left[\ \dot{\mathbf{x}}(t_1;\boldsymbol{\beta}_1) \mid \dot{\mathbf{x}}(t_2;\boldsymbol{\beta}_1) \mid \ldots \mid \dot{\mathbf{x}}(t_{N_t};\boldsymbol{\beta}_1) \mid \dot{\mathbf{x}}(t_1;\boldsymbol{\beta}_2) \mid \ldots \mid \dot{\mathbf{x}}(t_{N_t};\boldsymbol{\beta}_2) \mid \ldots \mid \dot{\mathbf{x}}(t_{N_t};\boldsymbol{\beta}_{N_\beta})\ \right]^T,$$

respectively, such that $(\mathbf{X}, \dot{\mathbf{X}}) \in \mathbb{R}^{N_t N_\beta \times N}$, where N_t is the total number of time instants and N_β is the number of parameter/forcing instances considered. However, approximating such a high dimensional function $\mathbf{f}$ might be an extremely complicated and computationally demanding task.

3

2.2. Latent space construction

In order to significantly speed up the computationally expensive simulations, to have a more interpretable and parsimonious dynamic model, and to efficiently estimate the dynamical behaviour of the system at steady-state by continuation, we propose a reduction technique that shall simultaneously reduce the problem dimensionality and provide a new set of coordinates better suited to represent the dynamics of the problem. Starting from the assumption that the full order solution of the parametrized system (1) lies on a low-dimensional manifold embedded in high-dimensional (discrete) space [52], we efficiently reduce the dimensionality of the problem directly from the snapshot data matrix $\mathbf{X}$, without accessing the FOM operators appearing in (1). The reduced space dimensionality n can be set equal to the intrinsic dimension of the solution manifold whenever this is known a priori or, conversely, it can be tuned to discover the intrinsic dimension of an unknown system for which only state measurements are available.

Therefore, we consider a preliminary linear reduction by POD, projecting the full order states of the system $\mathbf{x}$ onto the N_{POD} dominant *singular vectors* denoted as $\tilde{\mathbf{U}} \in \mathbb{R}^{N \times N_{\text{POD}}}$, i.e., $\tilde{\mathbf{x}} = \mathbf{x}\tilde{\mathbf{U}}^1$. POD coordinates of FOM data matrices are then $\tilde{\mathbf{X}} = \mathbf{X}\tilde{\mathbf{U}}, \dot{\tilde{\mathbf{X}}} = \dot{\mathbf{X}}\tilde{\mathbf{U}}$. Next, a further nonlinear reduction is performed by means of an autoencoder neural network. In particular, we identify a new set of latent variables $\mathbf{z}$ such that

$$\mathbf{z}(t; \boldsymbol{\beta}) = \boldsymbol{\varphi}(\tilde{\mathbf{x}}(t; \boldsymbol{\beta})), \tag{2}$$

where $\boldsymbol{\varphi}(\cdot) = \boldsymbol{\varphi}(\cdot; \mathbf{W}_\varphi) : \mathbb{R}^{N_{\text{POD}}} \to \mathbb{R}^n$, $n \ll N_{\text{POD}} \ll N$, is an *encoder* consisting of a fully connected neural network and $\mathbf{W}_\varphi$ denotes the network parameters. Moreover, in order to go back to POD coordinates from intrinsic ones, we make use of a *decoder*, that is, another fully connected neural network $\boldsymbol{\psi}(\cdot; \mathbf{W}_\psi) : \mathbb{R}^n \to \mathbb{R}^{N_{\text{POD}}}$ such that

$$\boldsymbol{\psi}(\mathbf{z}(t; \boldsymbol{\beta})) = \boldsymbol{\psi}(\boldsymbol{\varphi}(\tilde{\mathbf{x}}(t; \boldsymbol{\beta}))) \approx \tilde{\mathbf{x}}(t; \boldsymbol{\beta}). \tag{3}$$

Finally, to reconstruct the FOM solution, we project back through POD modes the output of the decoder, namely $\boldsymbol{\psi}(\mathbf{z}(t; \boldsymbol{\beta}))\tilde{\mathbf{U}}^T \approx \mathbf{x}(t; \boldsymbol{\beta})$.

The pair $(\boldsymbol{\varphi}, \boldsymbol{\psi})$ constitutes the *autoencoder* (AE) neural network [27] which is trained at once by solving

$$\min_{\mathbf{W}_\varphi, \mathbf{W}_\psi} \left\| \tilde{\mathbf{X}} - \boldsymbol{\psi}\left(\boldsymbol{\varphi}\left(\tilde{\mathbf{X}}; \mathbf{W}_\varphi\right); \mathbf{W}_\psi\right) \right\|_F^2, \tag{4}$$

where $\|\cdot\|_F$ is the Frobenius norm and ψ, φ are intended to be applied row-wise to the input matrices. The preliminary POD reduction enables a more efficient encoding process [21], since the AE only needs to reduce the dimension from a moderate number N_{POD} of features to n, instead of starting from an extremely large amount N of degrees of freedom. It is worth stressing that POD, a linear reduction technique, presents several limitations when applied alone to nonlinear, time-dependent parametrized problems, thus motivating the inclusion of the neural network encoder.

In the new set of intrinsic coordinates, defined by (2), system (1) can be reformulated as

$$\begin{cases} \dot{\mathbf{z}}(t; \boldsymbol{\beta}) = \tilde{\mathbf{f}}(t, \mathbf{z}(t; \boldsymbol{\beta}); \boldsymbol{\beta}), & t \in (0, T), \\ \mathbf{z}(0; \boldsymbol{\beta}) = \mathbf{z}_0, \end{cases} \tag{5}$$

where $\tilde{\mathbf{f}}$ encodes the dynamics of the low-dimensional system, $\dot{\mathbf{z}} = \dot{\varphi}(\tilde{\mathbf{x}}) = \dot{\tilde{\mathbf{x}}}\nabla_{\tilde{\mathbf{x}}}^T \boldsymbol{\varphi}(\tilde{\mathbf{x}})$ and $\mathbf{z}_0 = \boldsymbol{\varphi}(\tilde{\mathbf{x}}_0)$, with $\tilde{\mathbf{x}}_0 = \mathbf{x}_0\tilde{\mathbf{U}}$.

2.3. Latent dynamics identification

The problem of estimating the high-dimension function $\mathbf{f} \in \mathbb{R}^N$ in (1) is now reduced to the estimation of $\tilde{\mathbf{f}} \in \mathbb{R}^n$ in (5), which describes the latent dynamics of the system in a low-dimensional space. We are interested in a model representation that should not only be accurate, but also interpretable and generalizable, that is,

---

[1]In the present work, snapshots are stacked as rows for consistency with how neural networks are trained, and in contrast with standard POD literature, where snapshots are listed in columns. For clarification, the reader must intend POD to be applied to $\mathbf{X}^T$, i.e. $\mathbf{X}^T \approx \tilde{\mathbf{U}}\tilde{\boldsymbol{\Sigma}}\tilde{\mathbf{V}}^T$. Consistently with the notation adopted, POD projections – usually indicated as $\tilde{\mathbf{x}} = \tilde{\mathbf{U}}^T\mathbf{x}$ – are rewritten here in terms of row vectors, namely $\tilde{\mathbf{x}} = \mathbf{x}\tilde{\mathbf{U}}$.

capable of forecasting in time and extrapolating with respect to parameters. For the latter reasons, a purely black-box approach is not optimal in this context.

Therefore, to identify the latent dynamics, we make use of the SINDy regression strategy [35, 5, 6], which assumes that $\tilde{\mathbf{f}}$ can be expressed as a sparse combination of a set of linear and nonlinear candidate basis functions of the latent state $\mathbf{z}$ and of $\boldsymbol{\beta}$. The choice of the library of candidate features is typically guided by a possible prior knowledge of the physical system and of the parameter/forcing dependency. A typical choice is polynomials as they frequently appear in most dynamical system models.

In practice, SINDy aims to approximate the latent dynamics $\tilde{\mathbf{f}}$ in (5) as

$$\dot{\mathbf{z}} = \tilde{\mathbf{f}}(\mathbf{z}; \boldsymbol{\beta}) \approx \boldsymbol{\Theta}(\mathbf{z}, \boldsymbol{\beta}) \boldsymbol{\Xi} \tag{6}$$

where $\boldsymbol{\Theta}(\mathbf{z}, \boldsymbol{\beta}) = [\, \theta_1(\mathbf{z}, \boldsymbol{\beta}), \ldots, \theta_r(\mathbf{z}, \boldsymbol{\beta}) \,] \in \mathbb{R}^r$ is the library of r candidate functions to describe the dynamics of the data, while $\boldsymbol{\Xi} \in \mathbb{R}^{r \times n}$ is the unknown matrix of coefficients that determine the active terms from $\boldsymbol{\Theta}$ in the dynamics $\tilde{\mathbf{f}}$. Matrix $\boldsymbol{\Xi}$ is estimated by sparse regression

$$\arg\min_{\boldsymbol{\Xi}} \left\| \dot{\mathbf{Z}} - \boldsymbol{\Theta}(\mathbf{Z}, \mathcal{B}) \boldsymbol{\Xi} \right\|_F^2 + \lambda \|\boldsymbol{\Xi}\|_1, \tag{7}$$

where $\mathbf{Z} = \left[\mathbf{z}(t_1, \boldsymbol{\beta}_1) \mid \ldots \mid \mathbf{z}(t_{N_t}, \boldsymbol{\beta}_{N_\beta}) \right]^T \in \mathbb{R}^{N_t N_\beta \times n}$ is the matrix of low-dimensional encoded snapshots, $\dot{\mathbf{Z}} = \dot{\tilde{\mathbf{X}}} \nabla_{\tilde{\mathbf{x}}}^T \boldsymbol{\varphi}(\tilde{\mathbf{X}})$ contains their time derivatives, $\mathcal{B} = \left[\boldsymbol{\beta}_1(t_1) \mid \ldots \mid \boldsymbol{\beta}_{N_\beta}(t_{N_t}) \right]^T \in \mathbb{R}^{N_t N_\beta \times p}$ collects the parameter/forcing data, while $\boldsymbol{\Theta}$ is intended to be applied row-wise to the input matrices. Here, $\|\boldsymbol{\Xi}\|_1$ is a regularization term chosen to promote sparsity in $\boldsymbol{\Xi}$ and $\lambda \in \mathbb{R}^+$ weights its contribution. More sophisticated techniques, such as, e.g., sequentially thresholded least-square algorithm, might be used to promote sparsity on the entries of $\boldsymbol{\Xi}$ more efficiently [8, 3]. Prior knowledge can be incorporated by imposing constraints on the entries of $\boldsymbol{\Xi}$, which reflect physical properties of the problem. Specific entries can be masked and constrained to assume values of given coefficients. As a result, they remain constant over the training and only unconstrained entries of $\boldsymbol{\Xi}$ are updated at each iteration. In the following, we discuss how to practically solve the sparse regression problem (7) together with (4).

2.4. Offline training

In the previous sections we introduced the overall framework which consists of *(i)* a dimensionality reduction and a change of coordinates from physical to intrinsic ones via POD and AE and *(ii)* the identification of the latent dynamics by sparse regression. Tasks *(i)* and *(ii)* are addressed by solving minimization problems (4) and (7), respectively. Although of different types, these tasks are intrinsically related [8, 3]. Indeed, the possibility of finding an accurate latent dynamical description with few meaningful and interpretable terms depends strongly on the choice of intrinsic coordinates, thus on the AE mapping.

Therefore, both tasks are performed by a single neural network architecture model denoted AE+SINDy, as illustrated in Fig. 1. This unified structure allows us to reformulate (4) and (7) as a single minimization problem by incorporating the latter in the former, namely the sparse regression terms are inserted as additional weighted components in the AE loss function. In this perspective, the unknown entries of $\boldsymbol{\Xi}$ are treated as network variables and they are estimated at the same time as AE parameters $\mathbf{W}_\varphi$ and $\mathbf{W}_\psi$ during the training of the neural network. Specifically, the optimization problem reads as

$$\underbrace{\left\| \tilde{\mathbf{X}} - \boldsymbol{\psi}\left(\boldsymbol{\varphi}\left(\tilde{\mathbf{X}}; \mathbf{W}_\varphi \right); \mathbf{W}_\psi \right) \right\|_F^2}_{\text{Autoencoder loss}} + \underbrace{\lambda_1 \left\| \dot{\mathbf{Z}} - \boldsymbol{\Theta}(\mathbf{Z}, \mathcal{B}) \boldsymbol{\Xi} \right\|_F^2 + \lambda_2 \|\boldsymbol{\Xi}\|_1}_{\text{Sparse regression loss}}$$
$$+ \underbrace{\lambda_3 \left\| \dot{\mathbf{X}} - \nabla_{\mathbf{z}} \boldsymbol{\psi}(\mathbf{Z}; \mathbf{W}_\psi) \boldsymbol{\Theta}(\mathbf{Z}, \mathcal{B}) \boldsymbol{\Xi} \right\|_F^2}_{\text{Consistency loss}} \to \min_{\mathbf{W}_\varphi, \mathbf{W}_\psi, \boldsymbol{\Xi}}, \tag{8}$$

where the latter term, denoted as *consistency loss*, ensures that the time derivative of the network output matches the input time derivatives $\dot{\mathbf{X}}$. This term is proposed in [8] and extended here to account for parametric/forcing dependencies. The coefficients $\lambda_1, \lambda_2, \lambda_3 \in \mathbb{R}^+$ are hyperparameters which weight the contribution of each term. As a rule of thumb, the AE loss represents the leading term, so the coefficient λ_1, relative to the sparse regression term, should be smaller than unity and kept two orders of magnitude larger
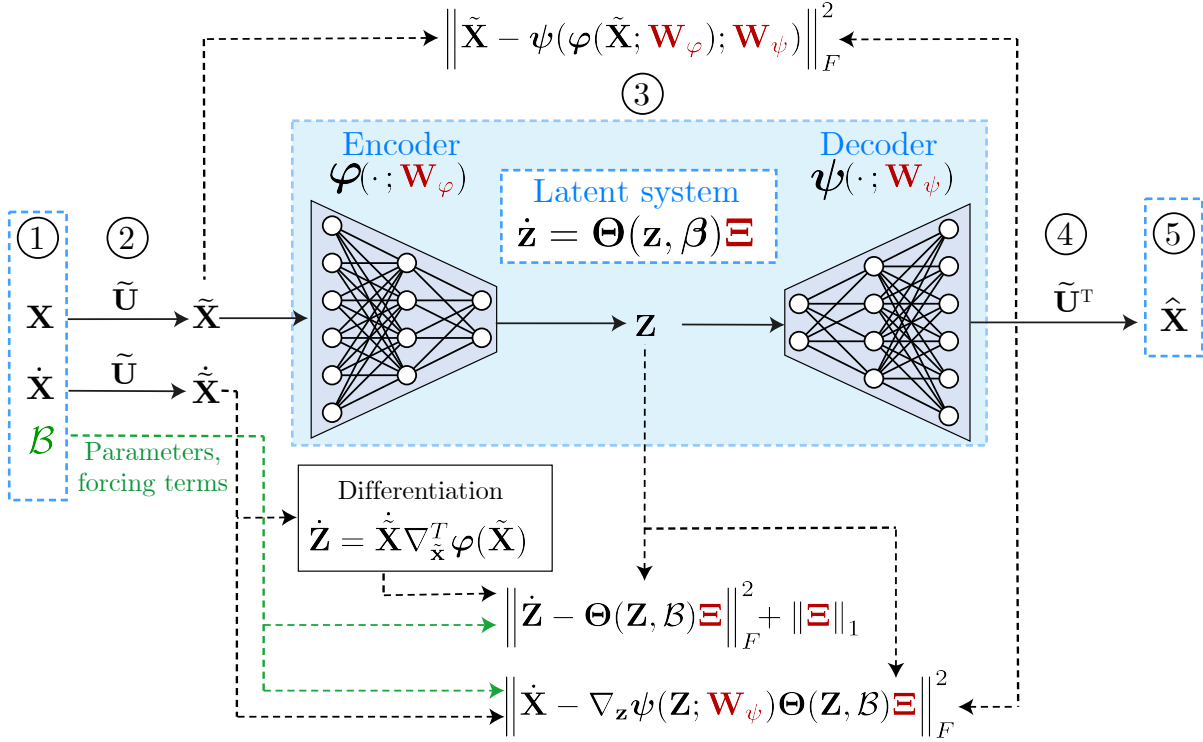
Figure 1: Neural network used during the training stage. The different stages of the training procedure are schematized as follows. $\mathbf{X}$ and $\dot{\mathbf{X}}$ are provided as input data (1) and are fed to the POD linear reduction procedure (2). The POD coordinates $\tilde{\mathbf{X}}$ are processed by the AE+SINDy (3), where the AE extracts the latent coordinates $\mathbf{Z}$ and these, together with $\dot{\mathbf{Z}}$ and $\mathcal{B}$, are used to compute the SINDy approximation of the latent dynamics. A linear POD reconstruction (4) provides the final output (5). The trainable parameters are highlighted in red.

than λ_2 and λ_3, these latter weighting regularization and consistency terms, respectively. More precise values need to be adjusted to the specific dataset at-hand, being the choice of these coefficients problem-dependent.

A further hyperparameter to be tuned is the latent dimension n. This can be set equal to the dimension of the low-dimensional manifold, if this information is available. Otherwise, the behavior of the *autoencoder loss* (4) as function of the number of latent variables may be used as an investigation tool to find the minimal latent dimension n. This value is the one for which further reducing the number of latent variables results in a drastic increase of the error. Although no examples are shown in the present work, AE+SINDy can be also employed by setting a number of variables larger than the dimension of the manifold. This may require placing more emphasis on *sparse regression loss* in (8), since the system identification task becomes more complex as both the number of latent equations and the number of features in the SINDy library increase.

The *offline training* procedure consists in the minimization of (8) through backpropagation with ADAM algorithm [38]. Once the joint training is finished, we freeze the AE weights $\mathbf{W}_\varphi, \mathbf{W}_\psi$ and we fine-tune just the SINDy coefficients $\mathbf{\Xi}$ for a better estimation of latent model coefficients [3].

2.5. Online testing

Once the AE+SINDy model is trained *offline*, we can query it *online* to efficiently compute the entire time evolution of the full order system for different initial conditions and new parameter and forcing instances. The procedure, illustrated in Fig. 2a , consists of the following steps:

1. **encoding**. Once the encoder φ is trained, it is employed as mapping from the physical, full order space to the latent one. Hence, the physical initial condition $\mathbf{x_0} \in \mathbb{R}^N$ is encoded in order to obtain the latent counterpart $\mathbf{z_0} = \varphi(\mathbf{x_0}\tilde{\mathbf{U}}) \in \mathbb{R}^n$;
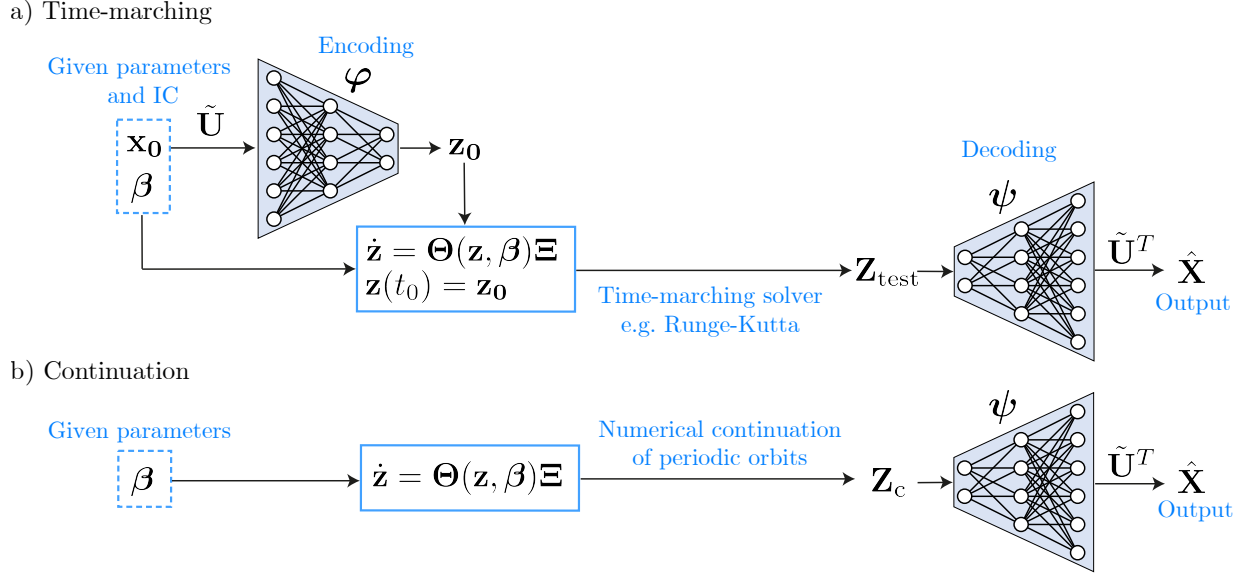
a) Time-marching

b) Continuation

Figure 2: Schematic representation of the setup used during the testing stage with the two approaches considered: time-marching solution and continuation of periodic orbits. In the former case, for a given value of the parameters **beta**, the ROM solution is computed through a time marching scheme, like, e.g., Runge-Kutta methods, starting from a given initial condition. The computed ROM solution spanning the time range $t \in [t_0 : t_{\text{end}}]$ is passed to the decoder. In the latter approach, the numerical continuation computes the periodic orbits of the ROM within a set of parameters. Latent orbits are collected in $\mathbf{Z}_{\text{C}}$ and then decoded.

2. **time-marching ROM solver**. SINDy feature library $\boldsymbol{\Theta}$ is populated with parameter and forcing terms $\boldsymbol{\beta}$, such that – together with the feature coefficients $\boldsymbol{\Xi}$ identified during the training – it defines the right-hand-side of the latent ODE dynamical system $\dot{\mathbf{z}} = \boldsymbol{\Theta}(\mathbf{z}, \boldsymbol{\beta})\boldsymbol{\Xi}$. This equations are numerically integrated, e.g. by Runge-Kutta time-marching scheme, starting from $\mathbf{z}(t_0) = \mathbf{z_0}$ up to an arbitrary time t_{end} – which may be much longer than the final time employed in the training phase. As result, we obtain the temporal evolution of the latent variables, collected in the matrix $\mathbf{Z}_{\text{test}} \in \mathbb{R}^{N_t^{\text{test}} \times n}$, where N_t^{test} is the total number of testing time instants;

3. **decoding**. The integrated latent variables are passed through the decoder and multiplied by the POD modes in order to reconstruct the entire time evolution of the original variables in the physical space, so $\hat{\mathbf{X}} = \boldsymbol{\psi}(\mathbf{Z}_{\text{test}})\tilde{\mathbf{U}}^T \in \mathbb{R}^{N_t^{\text{test}} \times N}$.

These steps are extremely efficient from a computational standpoint: indeed, steps (1) and (3) only require the evaluation of pre-trained neural network functions (encoder and decoder, respectively); step (2) consists instead of the numerical integration of a low-dimensional system, a task which is much cheaper and faster than the resolution of high-dimensional numerical schemes for full order problems.

2.6. Continuation of periodic orbits

As commented in Sect. 1, one of the major benefits behind the identification of an explicit, low-dimensional, parameter-dependent latent dynamical system is the possibility to use continuation algorithms to track the evolution of periodic responses as a function of the system parameters and to perform a bifurcation analysis of the solutions in order to have a complete portrait of the dynamics. As schematically represented in Fig. 2b, once the continuation parameter is defined periodic orbits are computed in the latent space, collected in a matrix $\mathbf{Z}_{\text{C}}$ and the corresponding full field solutions are eventually reconstructed with the decoder.

Periodic responses can be directly computed with several numerical strategies like harmonic balance [40], collocation [15], and shooting methods [54]. Focusing e.g. on the simplest version of collocation methods, the latent variables $\mathbf{z}$ are expressed over one period, of duration τ, as the weighted sum of predefined time-shape

functions N_i as

$$\mathbf{z}(t) = \sum_{i=1}^{n_c} N_i(t)\mathbf{z}_{c,i} \quad t \in (0, \tau) \tag{9}$$

with the additional constraint $\mathbf{z}(0) = \mathbf{z}(\tau)$. $\mathbf{z}_{c,i}$ denotes the value of $\mathbf{z}$ at the $i-th$ collocation point and n_c is the total number of collocation points. The shape functions might be defined globally over the period or locally, resorting to a partition of the period into finite elements in time. Inserting Eq.(9) into Eq.(5), this latter transforms in a nonlinear system of algebraic equations of the form:

$$G(\mathbf{z}_c, \beta) = 0 \tag{10}$$

where $\mathbf{z}_c$ represents the collection of the $\{\mathbf{z}_{c,i}\}_{i=1}^{n_c}$ values in Eq.(9) and β is the selected continuation parameter, i.e. the parameter which is varied in a predefined range over which the solutions for $\mathbf{z}_c$ are sought. For instance, in the numerical applications discussed next, the continuation parameters will be the actuation frequency for the clamped-clamped beam and the Reynolds number for the simulation of the flow past a cylinder. This equation is complemented by suitable phase conditions, see e.g. [41], to guarantee uniqueness.

Let us now suppose that $\mathbf{z}_c^{(k)}$ is a known solution of the system in correspondence of the parameter $\beta^{(k)}$. The simplest choice consists in taking β as continuation parameter, fixing $\beta^{(k+1)} = \beta^{(k)} + \Delta\beta$ and solving (10) for $\mathbf{z}_c^{(k+1)}$ through an iterative Newton–Raphson procedure. However, this procedure fails e.g. in the presence of an unstable branch as in Fig. 3. Indeed, by imposing an increment $\Delta\beta > 0$ at the peak, the solution would jump to stable solutions on the lower branch, completely missing the unstable dashed branch. For this reason, it is customary to introduce an arc-length control in which $\Delta\beta$ is part of the unknowns and the abscissa s along the solution curve is taken as the continuation parameter. A new constraint is added, a typical choice being:

$$F(\Delta\mathbf{z}_c, \Delta\beta) = (\Delta\mathbf{z}_c)^T \Delta\mathbf{z}_c + (\Delta\beta)^2 - (\Delta s)^2 = 0$$

where Δs is the user defined *distance* between successive solutions. An alternative is the Keller's pseudo arc-length method [41] in which the increment $(\Delta\mathbf{z}_c, \Delta\beta)$ such that its projection along a specific direction (typically the tangent to the $\mathbf{z}_c, \beta$ manifold) has length Δs. These methods are implemented in many ready-to-use packages like `Auto07p` [15], that implements collocation methods in `FORTRAN` to perform numerical continuation and bifurcation analysis; `Manlab`, a `Matlab` tool that uses HB methods and Asymptotic Numerical Method [30, 31]; `PyCont`[42] package on `Python`; `Nlvib`, that also exploits HB methods [40], and many others among which we mention `COCO` [12] and `BifurcationKit` [58]. The numerical examples addressed in this work will be solved using the Matlab `MATCONT` package [14], that exploits collocation methods to perform the continuation of periodic orbits.
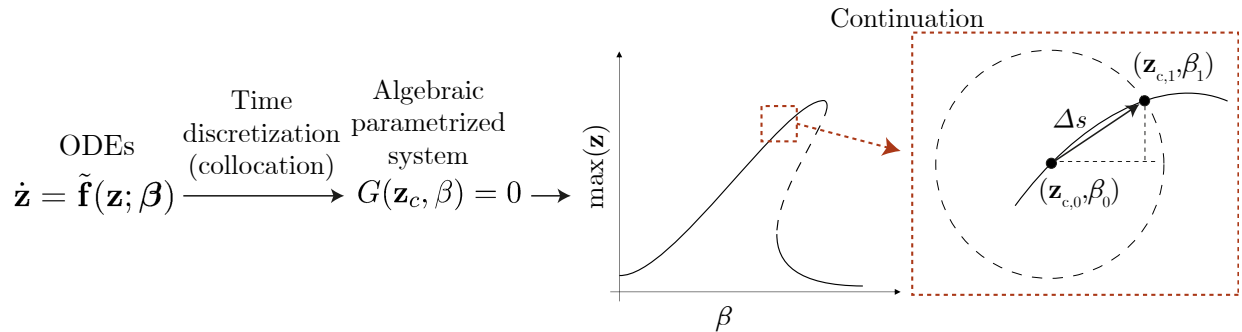

Figure 3: Sketch of the continuation of periodic orbits. Starting from a system of first order ODEs the system response is discretized in time using collocation methods. The continuation approach introduces a distance constraint between successive solutions corresponding to varying β. This allows to follow both stable as well as unstable branches.

3. Applications

3.1. Hardening behaviour of a clamped clamped beam

To highlight the capabilities of the AE+SINDy approach, we first address a structural mechanics problem in which a straight beam MEMS resonator is excited at resonance. Despite its simplicity, such a problem has practical applications like in micro resonators [62], where geometric nonlinearities provide meaningful contributions to the system dynamic response.

3.1.1. Problem description

We consider the doubly clamped beam depicted in Fig. 4 having length $L = 1000\,\mu\text{m}$ with a rectangular cross-section of dimensions $10\,\mu\text{m} \times 24\,\mu\text{m}$, made of isotropic polysilicon [10], with density $\rho = 2330\,\text{Kg/m}^3$, Young modulus $E = 167\,\text{GPa}$ and Poisson coefficient $\nu = 0.22$. The first bending eigenfrequency is $\omega_0 = 0.5475\,\text{rad}/\mu\text{s}$

In the undeformed configuration, the device occupies the domain Ω_0 described by material coordinates $\boldsymbol{x}$. The boundary $\partial\Omega_0$ is partitioned in $\partial\Omega_D$ and $\partial\Omega_N$ where homogeneous Dirichlet and Neumann boundary conditions are enforced, respectively. The external excitation is here provided by fictitious time-periodic body forces $\mathbf{B}(\boldsymbol{x}, t)$. The governing system of PDEs is classical in the context of solid mechanics problems in large transformations [47] and is formulated in terms of the displacement field $\mathbf{u}$ as follows:

$$\rho_0\ddot{\mathbf{u}}(\boldsymbol{x}, t) - \nabla \cdot \mathbf{P}(\boldsymbol{x}, t) - \rho_0\mathbf{B}(\boldsymbol{x}, t; \boldsymbol{\mu}) = \mathbf{0}, \quad (\boldsymbol{x}, t) \in \Omega_0 \times (0, T), \tag{11a}$$

$$\mathbf{P}(\boldsymbol{x}, t) \cdot \mathbf{N}(\boldsymbol{x}) = \mathbf{0}, \quad (\boldsymbol{x}, t) \in \partial\Omega_N \times (0, T), \tag{11b}$$

$$\mathbf{u}(\boldsymbol{x}, t) = \mathbf{0}, \quad (\boldsymbol{x}, t) \in \partial\Omega_D \times (0, T). \tag{11c}$$

Eq. (11a) expresses the conservation of momentum where ρ_0 is the initial density and $\mathbf{P}$ is the first Piola-Kirchhoff stress. The device is made of cubic single crystal silicon or polysilicon, thus admitting only small strains, a condition which is well described by the Saint Venant-Kirchhoff constitutive model $\mathbf{S}(\boldsymbol{x}, t) = \mathcal{A}(\boldsymbol{x}) : \mathbf{E}(\boldsymbol{x}, t)$ where $\mathbf{S}$ is the second Piola-Kirchhoff stress tensor, $\mathcal{A}$ is the fourth-order elasticity tensor and $\mathbf{E}(\boldsymbol{x}, t) = \frac{1}{2}\left(\nabla\mathbf{u}(\boldsymbol{x}, t) + \nabla^T\mathbf{u}(\boldsymbol{x}, t) + \nabla^T\mathbf{u}(\boldsymbol{x}, t) \cdot \nabla\mathbf{u}(\boldsymbol{x}, t)\right)$ is the Green-Lagrange strain tensor. Eq. (11b) and (11c) define the Neumann and Dirichlet boundary conditions respectively. Within the context at hand, it is worthy to highlight that Eq. (11) exactly accounts for geometric (elastic and inertia) nonlinearities, e.g., large rotations or nonlinear mode coupling. The spatial discretization of Eq. (11), e.g., by means of the finite element method, with the additional inclusion of a Rayleigh model damping term, yields a FOM under the form of a system of coupled first-order nonlinear differential equations that reads as:

$$\mathbf{M}\dot{\mathbf{v}}_h(t) + \mathbf{C}\mathbf{v}_h(t) + \mathbf{K}\mathbf{u}_h(t) + \mathbf{G}(\mathbf{u}_h, \mathbf{u}_h) + \mathbf{H}(\mathbf{u}_h, \mathbf{u}_h, \mathbf{u}_h) - \mathbf{F}(t; \boldsymbol{\beta}) = \mathbf{0}, \qquad t \in (0, T) \tag{12a}$$

$$\dot{\mathbf{u}}_h(t) - \mathbf{v}_h(t) = \mathbf{0}, \qquad t \in (0, T) \tag{12b}$$

$$\mathbf{u}_h(0) = \mathbf{0}, \quad \mathbf{v}_h(0) = \mathbf{0}, \tag{12c}$$
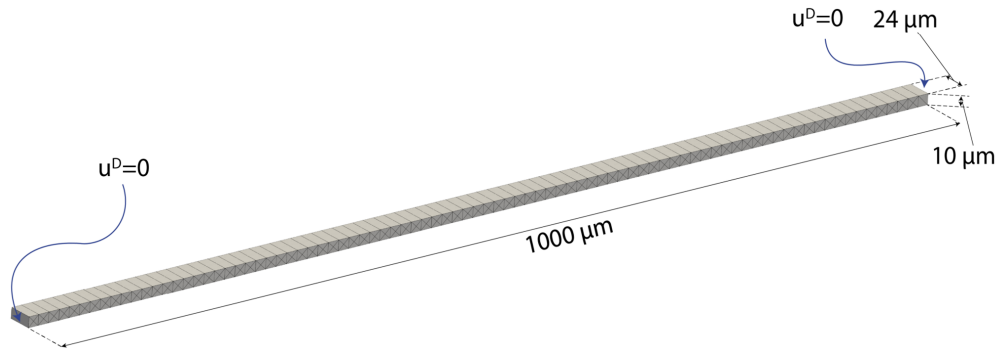

Figure 4: Schematic representation of the beam with the mesh used in the FOM simulations. Dirichlet boundaries are highlighted.

where the vectors $\mathbf{u}_h(t), \mathbf{v}_h(t) \in \mathbb{R}^{N_h}$ collect the N_h unknown displacements and velocity nodal values respectively, $\mathbf{M} \in \mathbb{R}^{N_h \times N_h}$ is the mass matrix, $\mathbf{C} = (\omega_0/Q)\mathbf{M}$ is the Rayleigh mass-proportional damping matrix with quality factor $Q = 50$. The internal force vector has been exactly decomposed in linear, quadratic, and cubic power terms of the displacement: $\mathbf{K} \in \mathbb{R}^{N_h \times N_h}$ is the stiffness matrix related to the linearized system, while $\mathbf{G} \in \mathbb{R}^{N_h}$ and $\mathbf{H} \in \mathbb{R}^{N_h}$ are vectors related to second- and third-order terms, respectively, and indicated with the notation $\mathbf{G}(\mathbf{u}_h, \mathbf{u}_h) = \sum_{k,l} \mathbf{G}_{kl} u_{hk} u_{hl}$ and $\mathbf{H}(\mathbf{u}_h, \mathbf{u}_h, \mathbf{u}_h) = \sum_{k,l,m} \mathbf{H}_{klm} u_{hk} u_{hl} u_{hm}$, where $\mathbf{G}_{kl}$ stands for the vector of coefficients G_{kl}^i, for $i = 1, \dots, N_h$, and analogously $\mathbf{H}_{klm}$ is a vector of coefficients H_{klm}^i. We refer to [53] for a detailed description of these terms. $\mathbf{F}(t; \boldsymbol{\beta}) \in \mathbb{R}^{N_h}$ is the nodal force vector which depends on the vector of parameters $\boldsymbol{\mu}$. In this application the device is excited, for simplicity, by a body load such that $\mathbf{F} = \mathbf{M}\boldsymbol{\phi}_1 F \cos(\omega t)$ is proportional to the first eigenmode $\boldsymbol{\phi}_1$, with $\mathbf{M}$ mass matrix and F load multiplier. In the application at hand the input parameter vector hence becomes $\boldsymbol{\beta} = [\omega, F] \in \mathcal{P}$, with $\mathcal{P}$ a closed and bounded set of dimension 2.

We point out that, independently of the mesh size, the solution of the FOM (12) is usually a challenging task since MEMS devices feature high quality factors while only the steady-state response is of practical interest [51]. Indeed, the performance of such devices is typically characterized by means of the so called Frequency Response Functions (FRFs) (see e.g. Figure 5) which, for different values of the load multipliers, express a steady-state output quantity of interest, like the maximum deflection, for all the values of angular frequency ω of the forcing within a prescribed range. As a consequence, the generation of reliable and efficient ROMs for this type of applications has recently stimulated intensive research with both standard and deep learning approaches; see, e.g., [52, 24, 19, 21].

3.1.2. Dataset

We consider 2 values for the amplitude, $F \in \{0.125, 0.250\} \, \mu N$, and 28 values for the frequency ω in the range $[0.526, 0.564] \, \text{rad}/\mu s$ with a finer sampling around the natural frequency ω_0. Hence the total number of instances is $N_\beta = |\mathcal{P}| = 56$. The parameter set $\mathcal{P}$ is randomly partitioned in training and testing subsets, $\mathcal{P}_{\text{train}}$ and $\mathcal{P}_{\text{test}}$, such that $N_\beta^{\text{train}} = |\mathcal{P}_{\text{train}}| = 54$ and $N_\beta^{\text{test}} = |\mathcal{P}_{\text{test}}| = 2$.

For each choice of training parameters, $(\omega, F) \in \mathcal{P}_{\text{train}}$, we numerically approximate the FOM solutions for displacement and velocity up to time $T = 390 \, \mu s$ and we collect them in the matrices $\mathbf{X}_{\boldsymbol{\beta}}, \dot{\mathbf{X}}_{\boldsymbol{\beta}} \in \mathbb{R}^{N_t \times N}$, where $N_t = 5000$ and $N = 7821$ are the number of time-steps and spatial degrees of freedom of the FOM, respectively, obtained by considering a mesh with 2607 nodes and linear finite elements. The matrices $\{\mathbf{X}_{\boldsymbol{\beta}}\}_{\boldsymbol{\beta} \in \mathcal{P}_{\text{train}}}$ are then stacked together, obtaining a single matrix $\mathbf{X}_{\text{train}} \in \mathbb{R}^{N_t N_\beta^{\text{train}} \times N}$, which contains the displacement snapshots for all the training parameters. Analogously to $\mathbf{X}_{\text{train}}$, we fill a matrix $\dot{\mathbf{X}}_{\text{train}}$ with the corresponding velocity snapshots.

Preprocessing. We preliminary reduce the dimensionality of the system by POD, moving from a dimension of $N = 7821$ to $N_{\text{POD}} = 64$. In order to do so, first a reduced basis is constructed by applying POD on the training displacement snapshot matrix $\mathbf{X}_{\text{train}}$. Next, all the FOM data are projected onto the first $N_{\text{POD}} = 64$ bases, yielding POD coordinates denoted as $\tilde{\mathbf{X}}_{\text{train}}, \dot{\tilde{\mathbf{X}}}_{\text{train}} \in \mathbb{R}^{N_t N_\beta^{\text{train}} \times N_{\text{POD}}}$.

Moreover, the neural network optimization algorithm is more efficient if the training data is in a limited magnitude range, typically in the order of unity, as the standard nonlinear activation functions perform at best in this range. Hence, we chose to rescale each of the 64 input features with respect to its maximum in absolute value, such that they all fall in the range $[-1, 1]$.

3.1.3. AE+SINDy architecture

The encoder block of the AE is composed of 3 hidden dense layers consisting of 64, 32 and 16 neural units. The decoder has a symmetrical structure. The dimensionality of the latent space corresponds to the size of the encoder output (as well as that of decoder input) and it is set to 1. Indeed, as discussed in [19, 21, 24], the dynamics underlying the system lies on a two dimensional manifold in the phase space [52], but the dependence on the velocity is minimal for the master bending mode that we inspect herein. Therefore, we aim to identify a latent dynamical system of the form

$$\ddot{z} = f(z, \dot{z}, \boldsymbol{\beta}), \tag{13}$$

where z is the latent variable, $\dot{z}$ and $\ddot{z}$ are its first- and second-order time derivatives, respectively, and $\boldsymbol{\beta}$ accounts for the parameter and forcing terms. All variables are time-dependent, although the dependency is not explicit for conciseness.

In order to estimate f by SINDy, we rewrite the second-order ODE (13) as a system of two first-order ODEs,

$$\dot{z} = y,$$
$$\dot{y} = f(z, y, \boldsymbol{\beta}), \tag{14}$$

which can be easily expressed in a suitable form for the SINDy framework, namely

$$\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}; \boldsymbol{\beta}) \approx \boldsymbol{\Theta}(\mathbf{z}; \boldsymbol{\beta})\, \boldsymbol{\Xi}. \tag{15}$$

Here, $\boldsymbol{\Theta}$ represents a library of polynomials up to the third degree with respect to $\mathbf{z} = [z, y]^T$ and of the first degree with respect to $\boldsymbol{\beta}$, while $\boldsymbol{\Xi}$ is the sparse unknown matrix of the multiplicative coefficients of the candidate features in $\boldsymbol{\Theta}$.

Regarding the choice of $\boldsymbol{\beta}$, since the beam structure is forced harmonically with amplitude and frequency parameterized by F and ω, these are incorporated in the SINDy model through the library terms $F\cos(\omega t), F\sin(\omega t)$. Therefore, $\boldsymbol{\Theta}$ has the following form

$$\boldsymbol{\Theta}(\mathbf{z}; \boldsymbol{\beta}) = \left[\, z(t)\,,\, y(t)\,,\, z^2(t)\,,\, z(t)\,y(t)\,,\, \dots\,,\, y^3(t)\,,\, F\cos(\omega t)\,,\, F\sin(\omega t)\,\right]$$

and contains $r = 11$ candidate features. Furthermore, we remark that not all the entries of $\boldsymbol{\Xi} \in \mathbb{R}^{r \times n}$, with $n = 2$, need to be estimated. In fact, by construction, the first equation in (14) is known, thus the first column of $\boldsymbol{\Xi}$ – which indeed represents the coefficients of this equation – is all set to 0 except for the second entry, relative to y, which is set to 1. Moreover, we can incorporate prior-knowledge about the dynamics by setting the values of (or imposing constraints on) the coefficients of $\boldsymbol{\Xi}$. In structural problems governed by (14), linear terms reduce to $-\omega_0^2 z$ that just depends on the natural frequency ω_0, which is known as the spectral properties of the linearized system are the classical primal output of the FOM. Thus, $\boldsymbol{\Xi}$ becomes

$$\boldsymbol{\Xi} = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 & 0 \\ -\omega_0^2 & \Xi_{2,2} & \Xi_{3,2} & \dots & \Xi_{r-1,2} & \Xi_{r,2} \end{bmatrix}^T \in \mathbb{R}^{r \times n}.$$

Following the procedure of Sect. 2.4, the overall architecture (see Fig. 1) is trained on $\tilde{\mathbf{X}}_{\text{train}}$ and $\dot{\tilde{\mathbf{X}}}_{\text{train}}$ to simultaneously reconstruct the POD coordinates with the AE as well as to learn the dynamics of the latent system by SINDy. The identified dynamical latent system finally reads:

$$\dot{z} = y,$$
$$\dot{y} = -0.3\,z - 0.011\,y + 0.003\,z^2 - 0.012\,y^2 - 0.113\,z^3 + 0.036\,z^2 y$$
$$+ 0.719\,z y^2 - 0.051\,y^3 - 0.009\,F\cos(\omega t). \tag{16}$$

3.1.4. Results

Once the offline training is concluded and the AE parameters $\mathbf{W}_\varphi$ and $\mathbf{W}_\psi$, as well as the SINDy coefficients $\boldsymbol{\Xi}$, are learned, we can employ the model for the online fast evaluation of new solutions or get insight into the dynamics by the continuation algorithms.

Time-marching solutions. The trained model is used to efficiently compute solutions for new instances of the parameters, i.e. for the new forcing frequencies ω and amplitude values F. In particular, we assess the performance on the test set $\mathcal{P}_{\text{test}}$. Following the procedure presented in Sect. 2.5, initial conditions and forcing parameters are passed to the encoder and to the SINDy model respectively; then the identified latent system is integrated up to $t_{\text{end}} = 1170\,\mu s$, i.e. three times longer than the training final time ($T = 390\,\mu s$). In particular, at t_{end} the system has reached the steady-state, while, during the training, only data covering a transient part of the dynamics have been fed to the network. Finally the physical solution is recovered via the decoder. In Fig. 5 we compare the AE+SINDy predictions with the displacement approximated with the FOM. The model manages to approximate the FOM solution with great accuracy in the whole
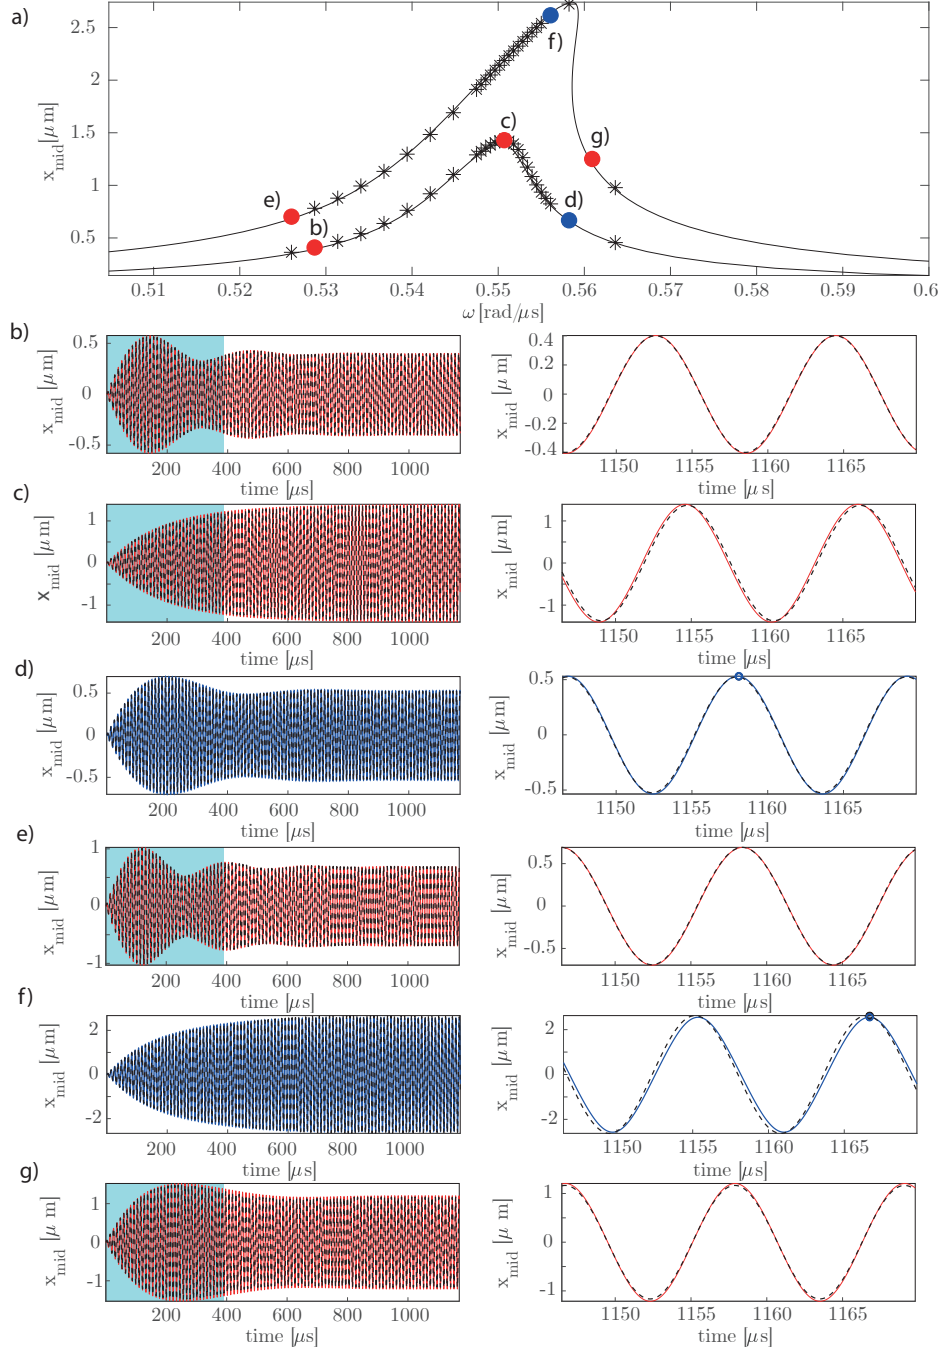
11

Figure 5: Comparison between the time evolution of the FOM numerical solution $\mathbf{X}$ and the predicted displacement $\hat{\mathbf{X}}$ on a specific mesh node positioned at mid-span of the beam, for different frequency-loading pairs. Fig. a) FRF of the system. The continuous line denotes the FOM solution, while the markers represent the solutions obtained with the AE+SINDy after a long time integration of the ROM. The coloured markers correspond to solutions reported in subfigures b)-g). In particular, for each marker the whole time history is plotted in the first column and an enlarged view of the last few periods is presented in the second column. The black dashed lines refer to the FOM solution, while the red and blue ones to the AE+SINDy solution on training and testing load levels, respectively. The light blue shaded regions define the training time interval. The markers on Figs. d) and f) highlight the instants spatially reconstructed in Fig.6.
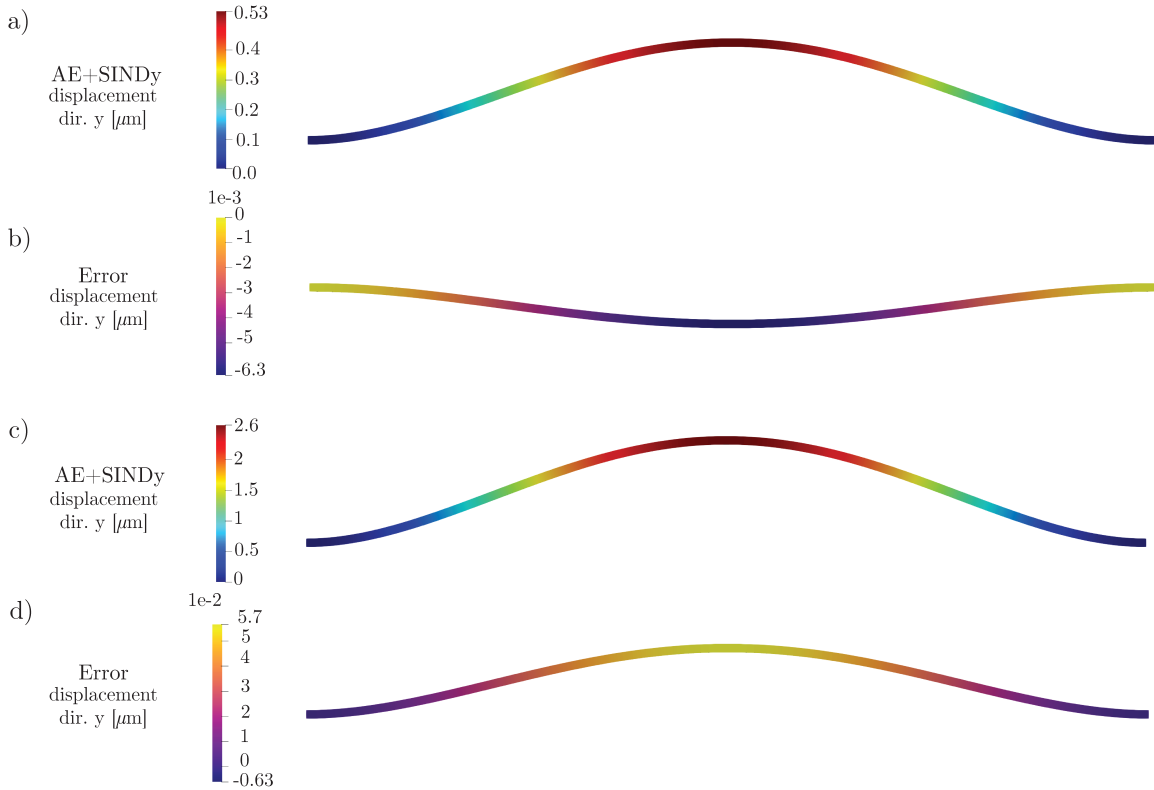
Figure 6: AE+SINDy spatial reconstruction of the displacement field on the y direction and the error with respect to the FOM. Fig. a-b) refer to the testing instance and time instant highlighted in Fig. 5 d), while Fig. c-d) refer to the one in Fig. 5 f).

time window. In addition, long-time predictions for few training parameters are depicted in Fig. 5 for a more comprehensive representation of the capacity of the method to extrapolate in time and to estimate the steady-state behavior in different regions of the parameter space.

These results highlight that *(i)* the identified model is able to explain dynamical phenomena, such as the steady-state behavior, while it has been trained only with snapshots covering a limited portion of the transient; *(ii)* the AE+SINDy framework not only manages to forecast in long-time horizons but it *simultaneously* generalizes with respect to new parameters; *(iii)* the identified latent one-dimensional dynamics physically represent and accurately approximate the whole high-dimensional system.

Continuation of periodic orbits. The low-dimensional latent dynamical system can be analysed with several different approaches and in particular with continuation algorithms that allow tracking directly the steady-state value of a quantity of interest as the parameters of the model are varied. In the problem at hand we are interested in plotting FRFs, that are particularly meaningful because they summarize most of the dynamical information of the mechanical system in a compact form.

Following the continuation strategy presented in Sect. 2.6, steady-state solutions are computed in the latent space, then mapped, by passing them to the decoder, into the physical space, where finally FRFs are drawn. In Fig. 7 we illustrate the comparison between the predicted curves with the exact curves obtained from the FOM [51]. Therefore, even though no steady-state information and unstable solutions have been provided during the training of the model, we still manage to track the evolution of the steady-state response changes with respect to the parameters, even in the unstable regions. Indeed, since the system is nonlinear we encounter bifurcations that affect the number of the solutions and their stability.

In Fig. 7, we plot (in light blue) the FRFs corresponding to the training values of F, i.e. $F \in \{0.125, 0.250\} \, \mu$N. We note that the predicted curves follow with great accuracy the exact ones and, in particular, there is an excellent match even outside the span of training frequencies, that is $[0.526, 0.564] \, \text{rad}/\mu$s, which denotes a
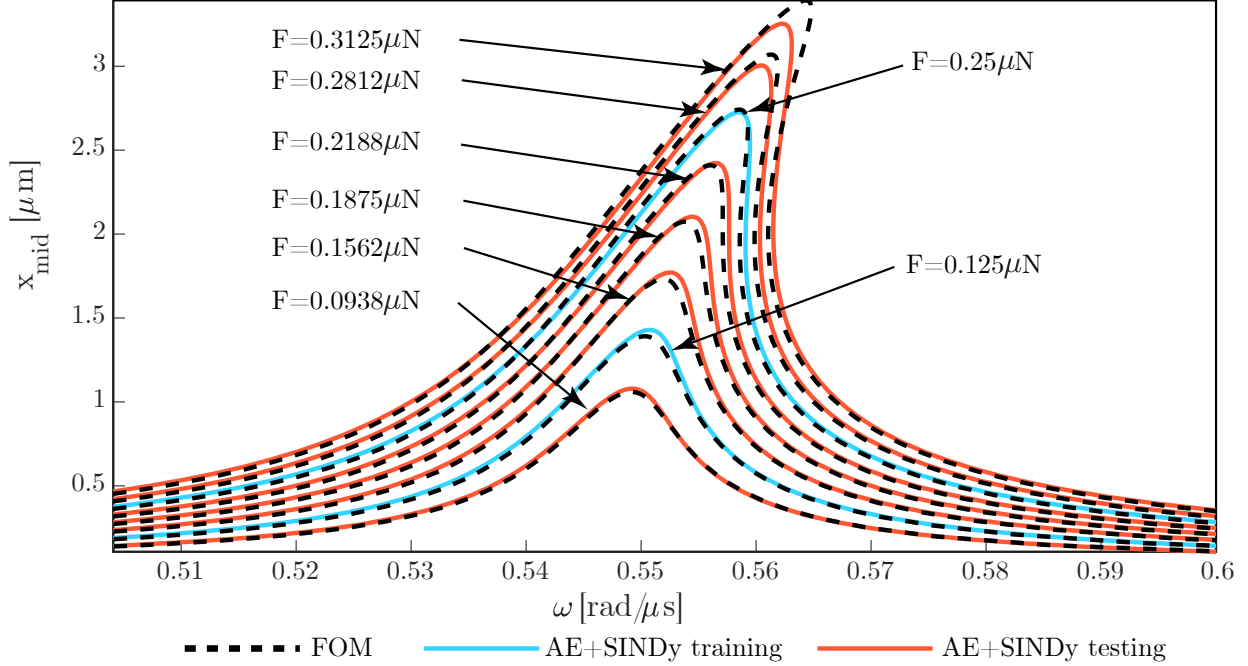
Figure 7: Comparison between FRFs obtained with the FOM [51] and the AE+SINDy. FOM solutions are plotted with dashed lines while AE+SINDy solutions with continuous ones. The red curves refers to testing load levels while the light blue one to training load levels. FRFs are represented as continuous lines for illustrative clarity, since continuation algorithm provides data-points which are very close to each other.

remarkable capability of the AE+SINDy to extrapolate with respect to parameter ω.

Moreover, we plot (in red) FRFs for forcing amplitudes F that have not been seen during the training phase. Also in this case, in general, there is a very good match with the exact curves, suggesting an excellent ability to generalize with respect to both the forcing amplitude F and frequency ω. A progressive degradation of the accuracy can be noted at the peak of FRFs for large values of forcing amplitudes outside the training range, e.g., for the curve corresponding to $F = 0.3125 \, \mu$N. This limitation in extrapolation is mainly due to low data sampling for the frequencies corresponding to peaks and to the fact that the AE performs a data-driven nonlinear coordinate transformation which may be less accurate when the nonlinearity content of the response increases well beyond what experienced during training. It is however worth stressing that far from the peaks the prediction is very accurate, even in the unstable branch of the FRFs. In support of these remarks we further notice that, on the contrary, for the curve at $F = 0.0938 \, \mu$N, i.e., below training range, we have a very accurate approximation.

3.2. Fluid flow around a cylinder

For this second application, our goal is to efficiently approximate the fluid velocity and pressure in the case of a fluid flow problem, at different regimes as the Reynolds number varies, as well as to empirically identify, through continuation algorithms, the bifurcation value of the transition from laminar to unsteady behaviour.

3.2.1. Problem description

We consider the well-known benchmark dealing with a two-dimensional fluid flow around a cylinder. The problem is described by the following unsteady Navier-Stokes equations for a viscous, incompressible Newtonian flow

$$\rho \frac{\partial \mathbf{v}}{\partial t} - \rho \mathbf{v} \cdot \nabla \mathbf{v} - \nabla \cdot \boldsymbol{\sigma}(\mathbf{v}, p) = \mathbf{0}, \qquad (\boldsymbol{x}, t) \in \Omega \times (0, T) ,$$
$$\nabla \cdot \mathbf{v} = 0, \qquad (\boldsymbol{x}, t) \in \Omega \times (0, T) . \tag{17}$$

Here, $\mathbf{v}(\boldsymbol{x}, t)$ and $p(\boldsymbol{x}, t)$ represents respectively the velocity and pressure fields of the flow, $\rho = 1.0 \text{ kg/m}^3$ is the fluid density, $\boldsymbol{\sigma}(\mathbf{v}, p) = -p\mathbf{I} + 2\nu\boldsymbol{\epsilon}(\mathbf{v})$ is the stress tensor, $\boldsymbol{\epsilon}(\mathbf{v})$ is the strain tensor, while ν is the kinematic

14

viscosity. The domain $\Omega = (0, 2.2) \times (0, 0.41) \backslash B_r(0.2, 0.2)$, with $r = 0.05$, consists in a bidimensional channel, where B_r represents a cylindrical obstacle (see Fig. 8). We prescribe the following boundary and initial conditions:

$$
\begin{aligned}
\mathbf{v} &= \mathbf{0}, & (\boldsymbol{x}, t) &\in \Gamma_{\mathrm{D}_1} \times (0, T), \\
\mathbf{v} &= \mathbf{h}, & (\boldsymbol{x}, t) &\in \Gamma_{\mathrm{D}_2} \times (0, T), \\
\boldsymbol{\sigma}(\mathbf{v}, p)\mathbf{n} &= \mathbf{0}, & (\boldsymbol{x}, t) &\in \Gamma_{\mathrm{N}} \times (0, T), \\
\mathbf{v}(\boldsymbol{x}, 0) &= \mathbf{0}, & \boldsymbol{x} &\in \Omega,
\end{aligned}
\tag{18}
$$

which represent a no-slip condition on Γ_{D_1}, a parabolic inflow of the form

$$
\mathbf{h}(\boldsymbol{x}, t) = \left(\frac{4U(t)x_2(0.41 - x_2)}{0.41^2}, 0 \right), \qquad \text{where } U(t) = \begin{cases} 0.75(1 - \cos(\pi t)), & t < 1 \\ 1.5, & t \geq 1 \end{cases}
$$

on the inlet Γ_{D_2}, open boundary conditions on the outlet Γ_{N} and homogeneous initial conditions.

In this framework, $\nu = 1/Re$ [20, 9], where Re is the Reynolds number, a quantity which allows to compare different flow regimes and characteristics. In the present study, we focus on the region $40 \leq Re \leq 60$. For the case at hand, when $Re < 49$ the flow is entirely laminar; instead, for larger Reynolds numbers, a pair of vortices are generated in the wake of the cylinder, periodically alternating between the top and bottom side, and the flow becomes unsteady [61, 57].

For a set of values of the number of Reynolds in the range $\mathcal{P} = [40, 60]$, we approximate the velocity components and pressure solutions up to $T = 30s$, for which the flow becomes fully developed. For the spatial and time discretization we consider finite elements and a semi-implicit backward differentiation formulas (BDF) respectively, employing MATLAB redbKIT library [49].

The geometric mesh employed is shown in Fig. 8, while the temporal discretization step is $\Delta t = 5ms$. The total number of degrees of freedom is $N = 73131$, obtained by considering quadratic finite elements for the velocity field and linear finite element for the pressure field, on a mesh made by 16478 triangular elements and 8239 vertices. Thus, numerical solutions of this type – although extremely accurate – are characterized by a very high dimensionality and they are computationally expensive and time consuming.

Our goal is to exploit the AE+SINDy structure to compute the entire space-time solution for a new instance of the Reynolds number, by exploiting a ROM consisting of a low-dimensional dynamical system. In addition, we are interested in understanding whether the identified system can actually represent the change of fluid behaviour as the Reynolds number varies, in particular, if it manages to detect the bifurcation at the transition from laminar to unsteady flow (between $Re = 49$ and 50, for this application [61, 57]).

3.2.2. Dataset

A uniform grid of $N_\beta = 21$ Reynolds number values over the parameter interval $\mathcal{P} = [40, 60]$ is considered. For each parameter value, we compute numerical solutions for the pressure and the two velocity components of the fluid. For the sake of computational efficiency, we restrict our analysis to the time window $[t_0, T] = [15s, 30s]$, by trimming off the first portion of the transient. $N_\beta^{\mathrm{train}} = 19$ snapshots will be employed as training set and the remaining $N_\beta^{\mathrm{test}} = 2$ for testing. In order to verify the accuracy of the proposed method in reconstructing different fluid behaviors, we consider as test instances the full order snapshots for $Re = \{43, 55\}$, which correspond to solutions in either the laminar or the unsteady regime, respectively.
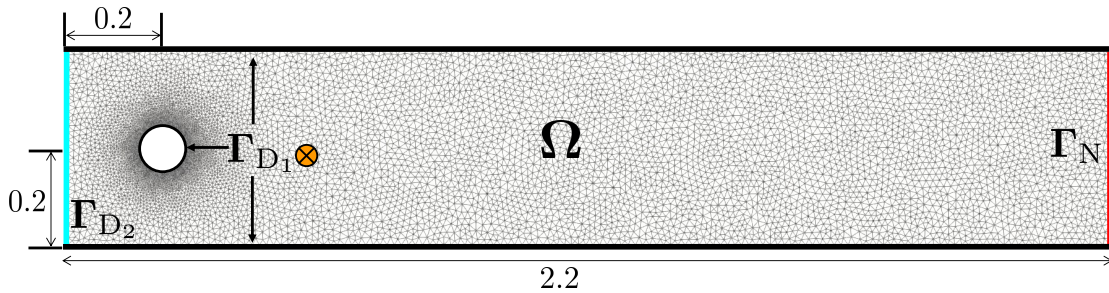

Figure 8: Geometry for the 2-D channel flow around a cylinder. All lengths are measured in meters. The node highlighted in orange is used as reference node for visualizations of results in Figs. 9,11.

Training and testing snapshots are stacked respectively in matrices $\mathbf{X}_{\text{train}} \in \mathbb{R}^{N_t N_\beta^{\text{train}} \times N}$ and $\mathbf{X}_{\text{test}} \in \mathbb{R}^{N_t N_\beta^{\text{test}} \times N}$, where $N_t = 3000$ is the number of time-steps.

As in the previous example, a preliminary dimensionality reduction is performed by POD. In particular, POD is applied separately on the training snapshots of both pressure and velocities, once they have been properly shifted to obtain zero mean snapshots, since they feature different units of measurement. Then, for both velocity and pressure, we retain the projections of the snapshots onto the first 32 bases (for a final dimension of $N_{\text{POD}} = 64$). These POD coordinates are scaled by the square root of the corresponding singular value (to reemphasize low-energy POD modes) and divided by the overall maximum in absolute value. In this way, POD coordinates of velocity and pressure have the same order of magnitude (i.e. unitary), which is also ideal for the neural network training. Finally, the two sets of POD coordinates are collected into a single matrix $\tilde{\mathbf{X}}_{\text{train}} \in \mathbb{R}^{N_t N_\beta^{\text{train}} \times N_{\text{POD}}}$. Here, time derivatives $\dot{\tilde{\mathbf{X}}}_{\text{train}}$ are computed numerically. Analogously, testing snapshots are projected onto the POD training modes, rescaled in the same way and collected in the matrix $\tilde{\mathbf{X}}_{\text{test}} \in \mathbb{R}^{N_t N_\beta^{\text{test}} \times N_{\text{POD}}}$.

3.2.3. AE+SINDy architecture

The AE architecture is such that the encoder block consists of 4 hidden layers of widths $64, 32, 16$ and 8, respectively, while the decoder has a symmetrical structure. Given a low Reynolds number, the dynamical system evolves onto a two-dimensional parabolic manifold [45, 50]. To further account for the parametric dependency of the manifold on the Reynolds number, the number of latent variables, which correspond to the dimension of the bottleneck of the AE, is set to 3. Regarding the SINDy algorithm, we employ a set of polynomials features $\boldsymbol{\Theta}$, up to the third degree, with respect to both the latent variables $\mathbf{z} = [z_1, z_2, z_3]^T$ and the parameter $\frac{1}{Re}$ to approximate the latent dynamics $\mathbf{f}$, such that

$$\dot{\mathbf{z}} = \mathbf{f}\left(\mathbf{z}, \frac{1}{Re}\right) \approx \boldsymbol{\Theta}\left(\mathbf{z}, \frac{1}{Re}\right)\boldsymbol{\Xi}, \tag{19}$$

where $\boldsymbol{\Xi}$ is the sparse matrix of the unknown multiplicative coefficients of the features of library $\boldsymbol{\Theta}$. The parametric dependence is with respect to the reciprocal of Reynolds number – which corresponds to the viscosity ν – for consistency with the full order equation (17). Differently from the previous example in Sect. 3.1 for which the parametric dependency is directly related to the forcing term, thus known a priori, this application provides an example for which no information about how the parameter enters the latent equations is available. For this reason, we consider a extensive polynomial library that includes also parameter interactions with latent variables.

The identified coefficients of the latent system (19), i.e. the entries of $\boldsymbol{\Xi}$, are reported in Table 1.

Table 1: Entries of $\boldsymbol{\Xi}$ obtained by training AE+SINDy. Each column indicates to which candidate feature the coefficient is associated to, while rows indicate the left-hand-side of the equation of the system (19) in which the coefficient enters. Note that the following scaling: $\beta = 10^3/Re = \nu \cdot 10^3$ has been considered.

	1	z_1	z_2	z_3	β	z_1^2	$z_1 z_2$	$z_1 z_3$	$z_1\beta$	z_2^2	$z_2 z_3$	$z_2\beta$
$\dot{z}_1$	2.23	0	3.44	61.8	-0.25	0	-4.17	0	0	0	-8.15	-0.43
$\dot{z}_2$	10.81	-89.0	112.5	-67.5	0	327.2	-623.2	82.17	-2.68	244.7	-54.0	-1.49
$\dot{z}_3$	0	-35.4	-16.1	0	0.04	229.2	-61.2	81.9	2.34	-7.41	1.46	1.77

	z_3^2	$z_3\beta$	β^2	z_1^3	$z_1^2 z_2$	$z_1^2 z_3$	$z_1^2\beta$	$z_1 z_2^2$	$z_1 z_2 z_3$	$z_1 z_2\beta$	$z_1 z_3^2$	$z_1 z_3\beta$
$\dot{z}_1$	-7.55	-29.5	-8.15	0	-33.6	111.3	-1175.1	0.14	-142.5	1449.0	-3.38	-687.0
$\dot{z}_2$	-3.79	-23.0	3.96	0	-734.1	1309.0	-187.5	12.5	-781.0	0	-2.23	0
$\dot{z}_3$	31.0	11.8	0	0	0	154.1	217.0	-12.4	-153.8	-10.9	0	0

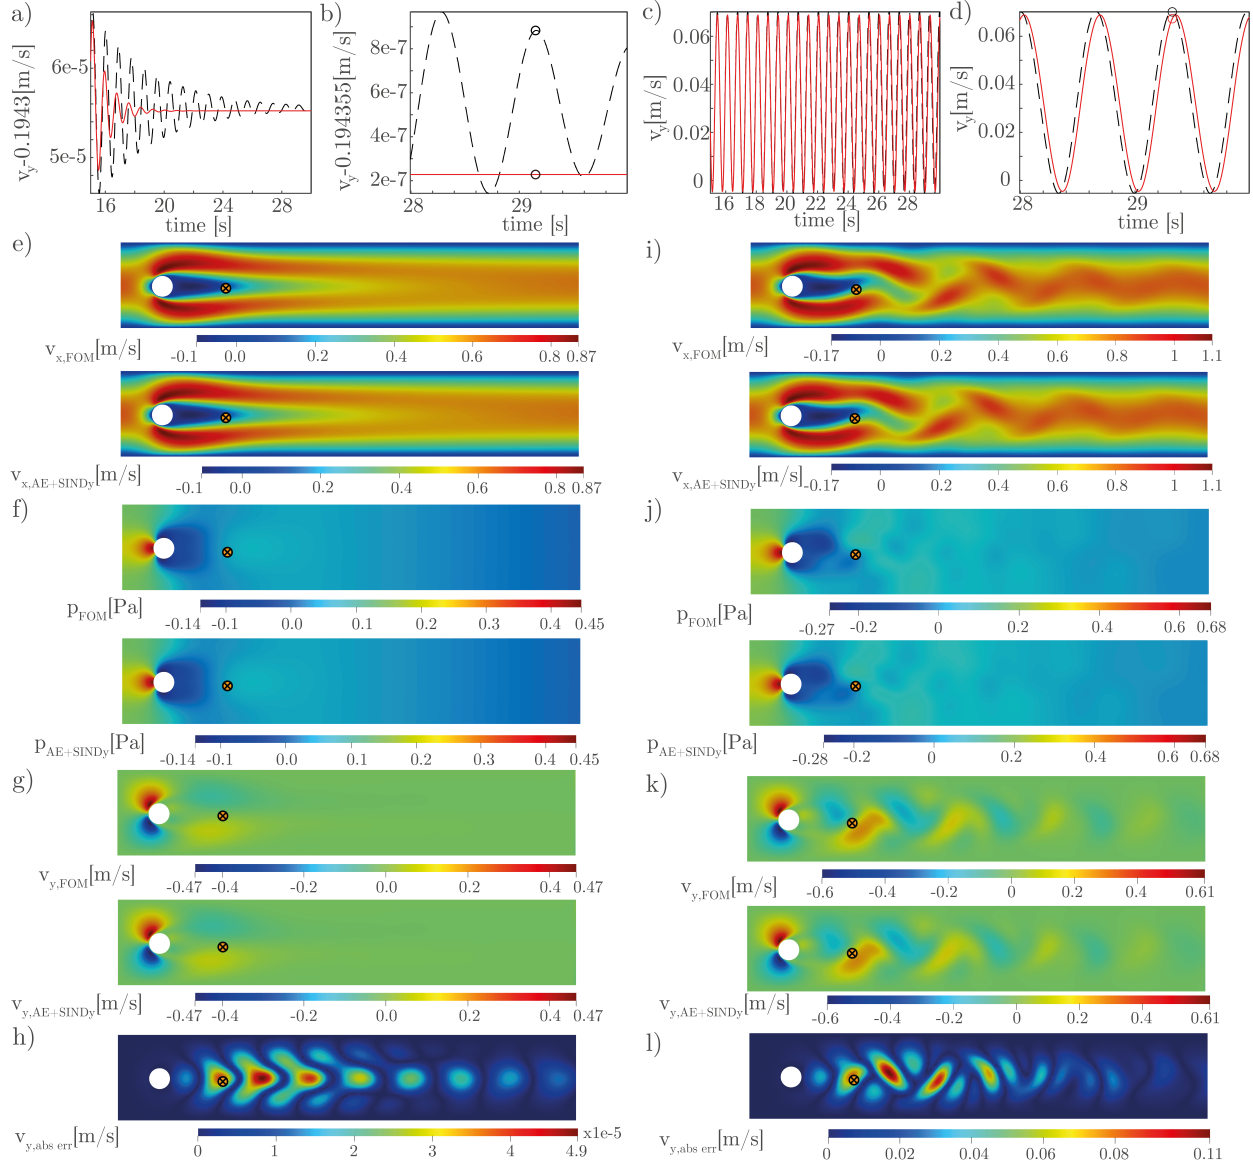|	$z_1\beta^2$	z_2^3	$z_2^2 z_3$	$z_2^2\beta$	$z_2 z_3^2$	$z_2 z_3\beta$	$z_2\beta^2$	z_3^3	$z_3^2\beta$	$z_3\beta^2$	β^3
$\dot{z}_1$	0	4.58	-565.0	0	408.2	-18.5	0	-262.1	7.91	0	0
$\dot{z}_2$	0	151.5	113.0	-2.51	-85.9	6.78	0	47.3	1.26	0	0
$\dot{z}_3$	0	94.4	-121.0	4.03	139.7	-2.4	0	-58.5	0.77	0	0

Figure 9: Time marching solutions: comparison between FOM and AE+SINDy. Fig. a) time histories related to $Re = 43$ (AE+SINDy in red, FOM in black), enlarged view proposed in Fig. b); Fig. c) time histories related to $Re = 55$, enlarged view proposed in Fig. d); Fig. e) FOM and AE+SINDy velocity field along the x direction with $Re = 43$; Fig. f) FOM and AE+SINDY pressure field with $Re = 43$; Fig. g) FOM and AE+SINDY velocity field along the y direction with $Re = 43$; Fig. h) absolute difference velocity field along the y between the FOM and AE+SINDy with $Re = 43$; Fig. i) FOM and AE+SINDY velocity field along the x direction with $Re = 55$; Fig. j) FOM and AE+SINDy pressure field with $Re = 55$; Fig. k) FOM and AE+SINDY velocity field along the y direction with $Re = 55$; Fig. l) absolute difference between the FOM and AE+SINDy with $Re = 55$ on the y direction.

17

3.2.4. Results

Time-marching solutions. The AE+SINDy model is trained on $\tilde{\mathbf{X}}_{\text{train}}, \dot{\tilde{\mathbf{X}}}_{\text{train}}$ data, with the method introduced in Sect. 2.4, then it is evaluated on the testing values $Re = \{43, 55\}$ following the strategy in Sect. 2.5. Approximated solutions are predicted by AE+SINDy up to final time $T = 30s$ and compared to the numerical values $\mathbf{X}_{\text{test}}$ in Fig. 9. The proposed method provides solutions of comparable accuracy to the numerical one, though, at extremely advantageous computational cost and time. Furthermore, we emphasize how indeed the model manages to accurately reconstruct the different behaviors of the fluid, i.e. laminar and unsteady, as the Reynolds number varies.

However, this is not sufficient to ensure that the identified system is indeed representative of the whole dynamics of the phenomenon. To this end, we pass the identified system to the continuation algorithms to estimate the latent dynamic behavior of the system as the Reynolds number changes.

Continuation. We now intend to apply the continuation algorithm to the identified latent system in order to gain insight into the dynamics of the system under observation. Such approach allows us to directly compute the steady-state response, i.e. when the flow is fully developed, avoiding the costs of computing the transient and the error accumulation of numerical integration. In addition, the continuation technique is an extremely useful tool in this context for identifying and characterizing the transition of the dynamics from laminar to unsteady. Indeed, while in the unsteady regime the steady-state solution is characterized by periodic oscillations, in the laminar regime it is stationary. Therefore, we expect the continuation technique, which assumes a-priori the existence of a periodic response, to break down during a downward sweep over the Reynolds number when reaching the bifurcation point (see Figure 11). We aim to verify whether we can effectively recover the expected bifurcation value, which is between $Re = 49$ and 50 [61, 57], and to check if the latent solutions estimated by continuation are physically representative of the dynamics of the phenomenon when projected back to the physical space. We stress that appropriately locating the bifurcation value would not be straightforward with time-marching schemes, since they would require instead the calculation of the entire time evolution of solutions for a dense sampling of parameter instances.

Starting from $Re = 60$, we employ the continuation algorithm to compute the periodic steady-state latent solutions over the parametric domain under consideration, namely for $Re \in \mathcal{P} = [40, 60]$. We obtain that continuation procedure correctly halts at the Reynolds number which matches the expected bifurcation value for the transition from unsteady to laminar regime (see Fig. 10). To convey continuation insights to the original physical system we follow the procedure presented in Sect. 2.6: physical steady-state solutions are reconstructed by mapping the so-computed periodic latent variables via decoder and then by projection to the physical space through POD modes. The comparison between the predicted physical orbits with the
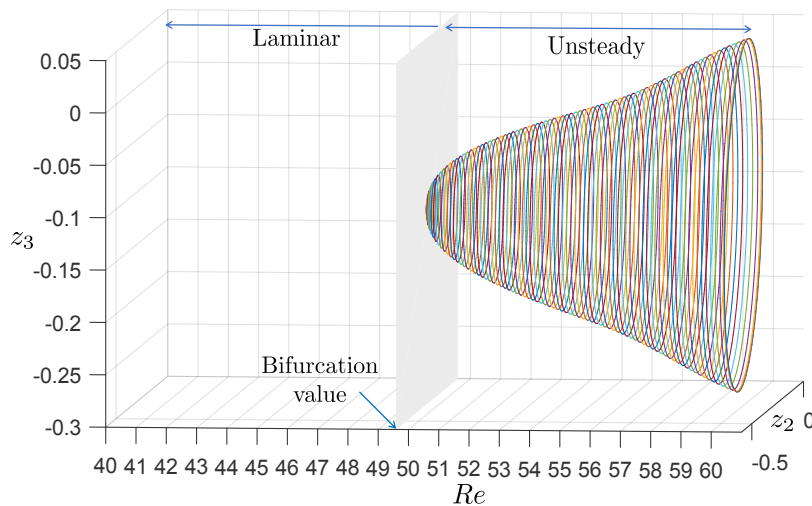

Figure 10: Behaviour of the latent coordinates (z_2, z_3) with respect to the Re parameter. After a critical Re value the fixed point solution (steady state flow) become unstable and a limit cycle arises (Hopf bifurcation). The limit cycle is embedded by the SINDy model integrated with continuation approach.
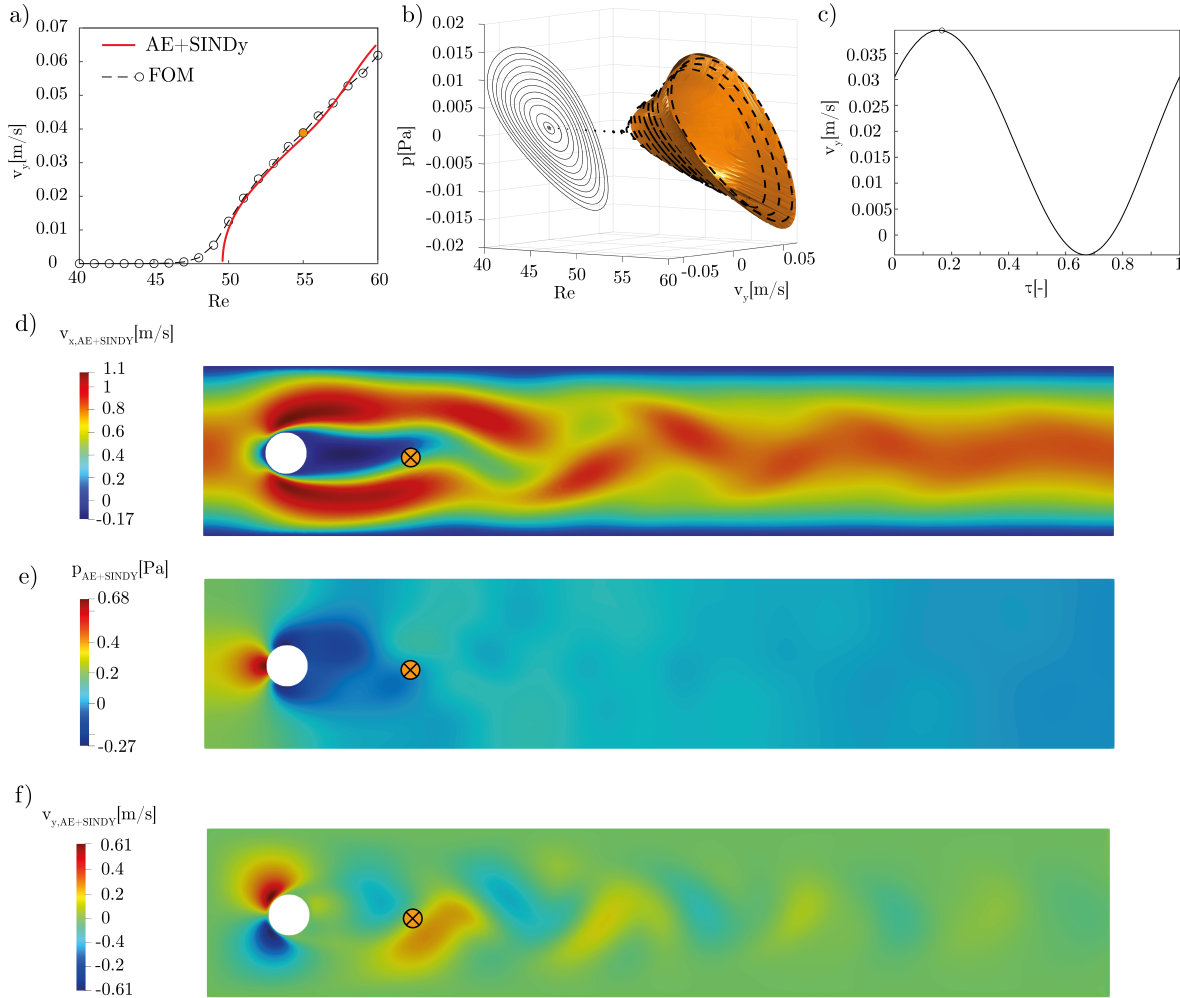
18

Figure 11: Continuation solutions with AE+SINDy. Fig. a) bifurcation diagram obtained from continuation compared with the FOM solution close to the steady state, the orange marker highlights the point having $Re = 55$ used as testing parameter value Fig. b) Surface described by the orbits envelope in the space (Re, v_y, P) referring to the node (highlighted previously) compared with the FOM solutions. Since the FOM is based on time-marching methods, the orbits here represented with dashed lines correspond to the one closer to the steady state. To better highlight the limit cycles, the orbits are also projected in the plane (v_y, P). Fig. c) solution corresponding to $Re = 55$ achieved with AE+SINDy, the marker highlights the time instant kept as reference in the spatial reconstructions Fig. d)-e)-f) AE+SINDy velocity fields along the two directions and the pressure field with $Re = 55$. The marker highlights the mesh node used as reference.

ones drawn from the original data is illustrated in Fig. 11.

Moreover, this technique allows to draw bifurcation diagrams which describes the characteristics of the flow as the Reynolds number changes. For instance, we report in Fig. 11a, the maximum value over time of the velocity of the flow in the cross-channel direction for a spatial point, as a function of the parameter Re. This velocity component is a direct indicator of the presence of vortexes in the flow that characterize the unsteady regime, while it is absent for the laminar flow. We observe that the bifurcation diagram identified by continuation is coherent with the (training and testing) snapshot data of the physical system (see Fig. 11a). For completeness, moreover, we show in Fig. 11d-e-f the physical reconstruction of the steady-state solution obtained by continuation at $Re = 55$. Comparing it with the numerical solution (see Fig. 9i-j-k), we observe that the dynamic behavior of the system is accurately detected.

4. Conclusions

In this work we proposed a framework for building low-dimensional dynamical models which overcomes the major bottlenecks of full order models to approximate solutions of parametrized, nonlinear, time-dependent systems of PDEs. Our method exploits both POD and AEs to construct a ROM, while simultaneously identifying the reduced dynamics via parametric SINDy. The possibility to incorporate parametric dependences in the latent ODE system, besides extending existing approaches [8, 3], opens the door to the application of continuation techniques. On one hand, these techniques allow to derive a portrait of the dynamics of the observed system; on the other hand, they offer the chance to efficiently estimate periodic steady-state solutions, without computing the whole transient dynamics.

We applied the proposed method to study two problems, the first in structural mechanics dealing with the hardening behaviour of a clamped clamped beam, and the second in fluid dynamics dealing with the motion of a fluid flow past a cylinder. In the first application, starting from simulation data covering a limited portion of the transient dynamics, we could efficiently derive a one-dimensional dynamical ROM which predicts the steady-state responses for new parameters outside the training range, thus extrapolating in time for long-term horizons while simultaneously generalizing with respect to system parameters. For the second example, the identified latent model allows to accurately estimate full order solutions, both for laminar and unsteady behaviours, as well as to identify the correct value of the Reynolds number at which the change of the flow regime occurs.

The considered examples highlight that if AE+SINDy successfully converges to a solution on training instances, it also allows numerical integration in long time horizon for testing instances as well as continuation algorithms to work, hence implying robustness and stability with respect to temporal and parameter extrapolation. This suggests that the AE+SINDy algorithm is less prone to overfitting phenomena with respect to standard data-driven method, thanks to the incorporation of physical and dynamical information via SINDy.

In addition to ensuring high accuracy at low computational cost, this method is highly non-intrusive since it just requires a very limited number of snapshots solutions and a choice of basis functions for the latent dynamics, instead of accessing the expensive FOM operators. Therefore, even though only simulation data are considered in the present work, the proposed method can be directly applied to observational and experimental data sets as well. This would allow to construct ROMs and discover the underlying dynamics also in the case of systems for which no information about their generative model is available.

Further improvements of the proposed method include the use of more powerful neural networks than standard feedforward autoencoders and the implementation of more sophisticated techniques for sparsity imposition in order to promote parsimonia and interpretability of the reduced dynamical model.

References

[1] D. Amsallem, M. J. Zahr, and C. Farhat. Nonlinear model order reduction based on local reduced-order bases. *International Journal for Numerical Methods in Engineering*, 92(10):891–916, 2012.

[2] N. Ananthkrishnan and K. Sudhakar. Characterization of periodic motions in aircraft lateral dynamics. *Journal of guidance, control, and dynamics*, 19(3):680–685, 1996.

[3] J. Bakarji, K. Champion, J. N. Kutz, and S. L. Brunton. Discovering governing equations from partial measurements with deep delay autoencoders. *arXiv preprint arXiv:2201.05136*, 2022.

[4] P. Benner, M. Ohlberger, A. Cohen, and K. (Eds.) Willcox. *Model reduction and approximation: theory and algorithms*. SIAM, 2017.

[5] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.

[6] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Sparse identification of nonlinear dynamics with control (sindyc). *IFAC-PapersOnLine*, 49(18):710–715, 2016.

[7] J. L. Callaham, S. L. Brunton, and J.-C. Loiseau. On the role of nonlinear correlations in reduced-order modelling. *Journal of Fluid Mechanics*, 938, 2022.

[8] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019.

[9] P. Conti, M. Guo, A. Manzoni, and J. S. Hesthaven. Multi-fidelity surrogate modeling using long short-term memory networks. *Computer methods in applied mechanics and engineering*, 404:115811, 2023.

[10] A. Corigliano, B. De Masi, A. Frangi, C. Comi, A. Villa, and M. Marchi. Mechanical characterization of polysilicon through on-chip tensile tests. *Journal of Microelectromechanical Systems*, 13(2):200–219, 2004.

[11] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[12] H. Dankowicz and F. Schilder. *Recipes for continuation*. SIAM, 2013.

[13] T. Detroux, L. Renson, L. Masset, and G. Kerschen. The harmonic balance method for bifurcation analysis of large-scale nonlinear mechanical systems. *Computer Methods in Applied Mechanics and Engineering*, 296:18–38, 2015.

[14] A. Dhooge, W. Govaerts, Y. A. Kuznetsov, W. Mestrom, A. Riet, and B. Sautois. Matcont and cl matcont: Continuation toolboxes in matlab. *Universiteit Gent, Belgium and Utrecht University, The Netherlands*, 2006.

[15] E. J. Doedel, A. R. Champneys, F. Dercole, T. F. Fairgrieve, Y. A. Kuznetsov, B. Oldeman, R. Paffenroth, B. Sandstede, X. Wang, and C. Zhang. Auto-07p: Continuation and bifurcation software for ordinary differential equations. 2007.

[16] I. R. Epstein and K. Showalter. Nonlinear chemical dynamics: oscillations, patterns, and chaos. *The Journal of Physical Chemistry*, 100(31):13132–13147, 1996.

[17] N. Franco, A. Manzoni, and P. Zunino. A deep learning approach to reduced order modelling of parameter dependent partial differential equations. *Mathematics of Computation*, 92(340):483–524, 2023.

[18] S. Fresca, L. Dede, and A. Manzoni. A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized pdes. *Journal of Scientific Computing*, 87(2):1–36, 2021.

[19] S. Fresca, G. Gobat, P. Fedeli, A. Frangi, and A. Manzoni. Deep learning-based reduced order models for the real-time simulation of the nonlinear dynamics of microstructures. *International Journal for Numerical Methods in Engineering*, 123(20):4749–4777, 2022.

[20] S. Fresca and A. Manzoni. Real-time simulation of parameter-dependent fluid flows through deep learning-based reduced order models. *Fluids*, 6(7):259, 2021.

[21] S. Fresca and A. Manzoni. Pod-dl-rom: enhancing deep learning-based reduced order models for nonlinear parametrized pdes by proper orthogonal decomposition. *Computer Methods in Applied Mechanics and Engineering*, 388:114181, 2022.

[22] W. D. Fries, X. He, and Y. Choi. Lasdi: Parametric latent space dynamics identification. *Computer Methods in Applied Mechanics and Engineering*, 399:115436, 2022.

[23] K. Fukami, T. Murata, K. Zhang, and K. Fukagata. Sparse identification of nonlinear dynamics with low-dimensionalized flow representations. *Journal of Fluid Mechanics*, 926, 2021.

[24] G. Gobat, S. Fresca, A. Manzoni, and A. Frangi. Virtual twins of nonlinear vibrating multiphysics microstructures: physics-based versus deep learning-based approaches. *arXiv preprint arXiv:2205.05928*, 2022.

[25] G. Gobat, A. Opreni, S. Fresca, A. Manzoni, and A. Frangi. Reduced order modeling of nonlinear microstructures through proper orthogonal decomposition. *Mechanical Systems and Signal Processing*, 171:108864, 2022.

[26] F. J. Gonzalez and M. Balajewicz. Deep convolutional recurrent autoencoders for learning low-dimensional feature dynamics of fluid systems. *arXiv preprint arXiv:1808.01346*, 2018.

[27] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.

[28] P. Goyal and P. Benner. Learning low-dimensional quadratic-embeddings of high-fidelity nonlinear dynamics using deep learning. *arXiv preprint arXiv:2111.12995*, 2021.

[29] P. Goyal and P. Benner. Discovery of nonlinear dynamical systems using a runge–kutta inspired dictionary-based sparse regression approach. *Proceedings of the Royal Society A*, 478(2262):20210883, 2022.

[30] L. Guillot, B. Cochelin, and C. Vergez. A taylor series-based continuation method for solutions of dynamical systems. *Nonlinear dynamics*, 98(4):2827–2845, 2019.

[31] L. Guillot, A. Lazarus, O. Thomas, C. Vergez, and B. Cochelin. A purely frequency based floquet-hill formulation for the efficient stability computation of periodic solutions of ordinary differential systems. *Journal of Computational Physics*, 416:109477, 2020.

[32] L. Guillot, P. Vigué, C. Vergez, and B. Cochelin. Continuation of quasi-periodic solutions with two-frequency harmonic balance method. *Journal of Sound and Vibration*, 394:434–450, 2017.

[33] J. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer, 2016.

[34] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

[35] E. Kaiser, J. N. Kutz, and S. L. Brunton. Sparse identification of nonlinear dynamics for model predictive control in the low-data limit. *Proceedings of the Royal Society A*, 474(2219):20180335, 2018.

[36] M. Kalia, S. L. Brunton, H. G. Meijer, C. Brune, and J. N. Kutz. Learning normal form autoencoders for data-driven discovery of universal, parameter-dependent governing equations. *arXiv preprint arXiv:2106.05102*, 2021.

[37] Y. Kim, Y. Choi, D. Widemann, and T. Zohdi. A fast and accurate physics-informed neural network reduced order model with shallow masked autoencoder. *J. Comput. Phys.*, 451:110841, 2022.

[38] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[39] J. Kneifl, D. Grunert, and J. Fehr. A nonintrusive nonlinear model reduction method for structural dynamical problems based on machine learning. *International Journal for Numerical Methods in Engineering*, 122(17):4774–4786, 2021.

[40] M. Krack and J. Gross. *Harmonic balance for nonlinear vibration problems*, volume 1. Springer, 2019.

[41] B. Krauskopf, H. M. Osinga, and J. Galán-Vioque. *Numerical continuation methods for dynamical systems*, volume 2. Springer, 2007.

[42] D. LaMar. Pycont. *https://pydstool.github.io/PyDSTool/PyCont*, 2021.

[43] K. Lee and K. T. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *J. Comput. Phys.*, 404:108973, 2020.

[44] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.

[45] J.-C. Loiseau, S. Brunton, and B. Noack. From the pod-galerkin method to sparse manifold models, 2021.

[46] B. Lusch, J. N. Kutz, and S. L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1):1–10, 2018.

[47] L. E. Malvern. *Introduction to the Mechanics of a Continuous Medium*. Number Monograph. 1969.

[48] R. Maulik, B. Lusch, and P. Balaprakash. Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders. *Phys. Fluids*, 33(3):037106, 2021.

[49] F. Negri. redbkit version 2.2. *http://redbkit.github.io/redbKIT*, 2016.

[50] B. R. Noack, K. Afanasiev, M. MORZYŃSKI, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *Journal of Fluid Mechanics*, 497:335–363, 2003.

[51] A. Opreni, N. Boni, R. Carminati, and A. Frangi. Analysis of the nonlinear response of piezo-micromirrors with the harmonic balance method. *Actuators*, 10(2):21, 2021.

[52] A. Opreni, A. Vizzaccaro, A. Frangi, and C. Touzé. Model order reduction based on direct normal form: application to large finite element mems structures featuring internal resonance. *Nonlinear Dynamics*, 105(2):1237–1272, 2021.

[53] A. Opreni, A. Vizzaccaro, C. Touzé, and A. Frangi. High-order direct parametrisation of invariant manifolds for model order reduction of finite element structures: application to generic forcing terms and parametrically excited systems. *Nonlinear Dynamics*, pages 1–47, 2022.

[54] M. R. Osborne. On shooting methods for boundary value problems. *Journal of mathematical analysis and applications*, 27(2):417–433, 1969.

[55] S. Pagani, A. Manzoni, and A. Quarteroni. Numerical approximation of parametrized problems in cardiac electrophysiology by a local reduced basis method. *Computer Methods in Applied Mechanics and Engineering*, 340:530–558, 2018.

[56] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced basis methods for partial differential equations: an introduction*, volume 92. Springer, 2015.

[57] B. Rajani, A. Kandasamy, and S. Majumdar. Numerical simulation of laminar flow past a circular cylinder. *Applied Mathematical Modelling*, 33(3):1228–1247, 2009.

[58] R. Veltz. *BifurcationKit. jl*. PhD thesis, Inria Sophia-Antipolis, 2020.

[59] A. Vizzaccaro, A. Opreni, L. Salles, A. Frangi, and C. Touzé. High order direct parametrisation of invariant manifolds for model order reduction of finite element structures: application to large amplitude vibrations and uncovering of a folding point. *Nonlinear Dynamics*, 110(1):525–571, 2022.

[60] A. Vizzaccaro, Y. Shen, L. Salles, J. Blahoš, and C. Touzé. Direct computation of nonlinear mapping via normal form for reduced-order models of finite element nonlinear structures. *Computer Methods in Applied Mechanics and Engineering*, 384:113957, 2021.

[61] M. M. Zdravkovich. *Flow around circular cylinders: Volume 2: Applications*, volume 2. Oxford university press, 1997.

[62] V. Zega, G. Gattere, S. Koppaka, A. Alter, G. D. Vukasin, A. Frangi, and T. W. Kenny. Numerical modelling of non-linearities in mems resonators. *Journal of Microelectromechanical Systems*, 29(6):1443–1454, 2020.