



MOX-Report No. 54/2025

Reduced order modeling with shallow recurrent decoder networks

Tomasetto, M.; Williams, J.P.; Braghin, F.; Manzoni, A.; Kutz, J.N.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Reduced Order Modeling with Shallow Recurrent Decoder Networks

Matteo Tomasetto^a, Jan P. Williams^c, Francesco Braghin^a, Andrea Manzoni^b, J. Nathan Kutz^{d,1}

^a*Department of Mechanical Engineering, Politecnico di Milano, Via privata Giuseppe La Masa 1, 20156, Milano, Italy*

^b*MOX – Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133, Milano, Italy*

^c*Department of Mechanical Engineering, University of Washington, 3900 E Stevens Way NE, 98195, Seattle, WA*

^d*Department of Applied Mathematics, University of Washington, 4182 W Stevens Wy NE, 98105, Seattle, WA*

^e*Department of Electrical and Computer Engineering, University of Washington, 185 E Stevens Way NE AE100R, 98195, Seattle, WA*

Abstract

Reduced Order Modeling is of paramount importance for efficiently inferring high-dimensional spatio-temporal fields in parametric contexts, enabling computationally tractable parametric analyses, uncertainty quantification and control. However, conventional dimensionality reduction techniques are typically limited to known and constant parameters, inefficient for nonlinear and chaotic dynamics, and uninformed to the actual system behavior. In this work, we propose *sensor-driven* SHallow REcurrent Decoder networks for Reduced Order Modeling (SHRED-ROM). Specifically, we consider the composition of a long short-term memory network, which encodes the temporal dynamics of limited sensor data in multiple scenarios, and a shallow decoder, which reconstructs the corresponding high-dimensional states. SHRED-ROM is a robust *decoding-only* strategy that circumvents the numerically unstable approximation of an inverse which is required by encoding-decoding schemes. To enhance computational efficiency and memory usage, the full-order state snapshots are reduced by, e.g., proper orthogonal decomposition, allowing for compressive training of the networks with minimal hyperparameter tuning. Through applications on chaotic and nonlinear fluid dynamics, we show that SHRED-ROM *(i)* accurately reconstructs the state dynamics for new parameter values starting from limited fixed or mobile sensors, independently on sensor placement, *(ii)* can cope with both physical, geometrical and time-dependent parametric dependencies, while being agnostic to their actual values, *(iii)* can accurately estimate unknown parameters, and *(iv)* can deal with different data sources, such as high-fidelity simulations, coupled fields and videos.

1. INTRODUCTION

Reduced order models (ROMs) are widely used computational tools for accelerating engineering design and characterization [4, 3, 49, 27]. Specifically, scientific computing is now an integral part of every field of application, with high-fidelity numerical solvers and methods [35] playing a critically enabling role in the modeling of high-dimensional, complex dynamical systems. However, this might be extremely demanding – or even prohibitive – in case repeated predictions are required for multiple scenarios, in very rapid times, or within sequential design or control loops. ROMs designate any approach aimed at replacing the high-fidelity problem by one featuring a much lower numerical complexity. In emerging modern applications, such as turbulence closure modeling, weather forecasting, powergrid networks, climate modeling and neuroscience, the construction of tractable dynamic models directly from data or high-fidelity simulations enables both engineering design and scientific discovery. Scientific machine learning [6] is an emerging paradigm for constructing data-driven ROMs. As with traditional ROMs, machine learned ROMs aim to learn both low-dimensional embeddings and evolution dynamics which accurately reconstruct (in a least-square or statistical sense) the high-fidelity and high-dimensional state of the original (possibly parametric) system. In what follows, we advocate a new ROM architecture based upon the method of *separation of variables* for solving *partial differential equations* (PDEs). Specifically, we train a recurrent neural network to capture the temporal behavior of limited sensor data, while mapping its latent space to the high-dimensional state via a shallow decoder. The resulting *Shallow REcurrent Decoder-based Reduced Order Model* (SHRED-ROM) turns out to be an ultra-hyperreduced order modeling framework which provides a fully data-driven and

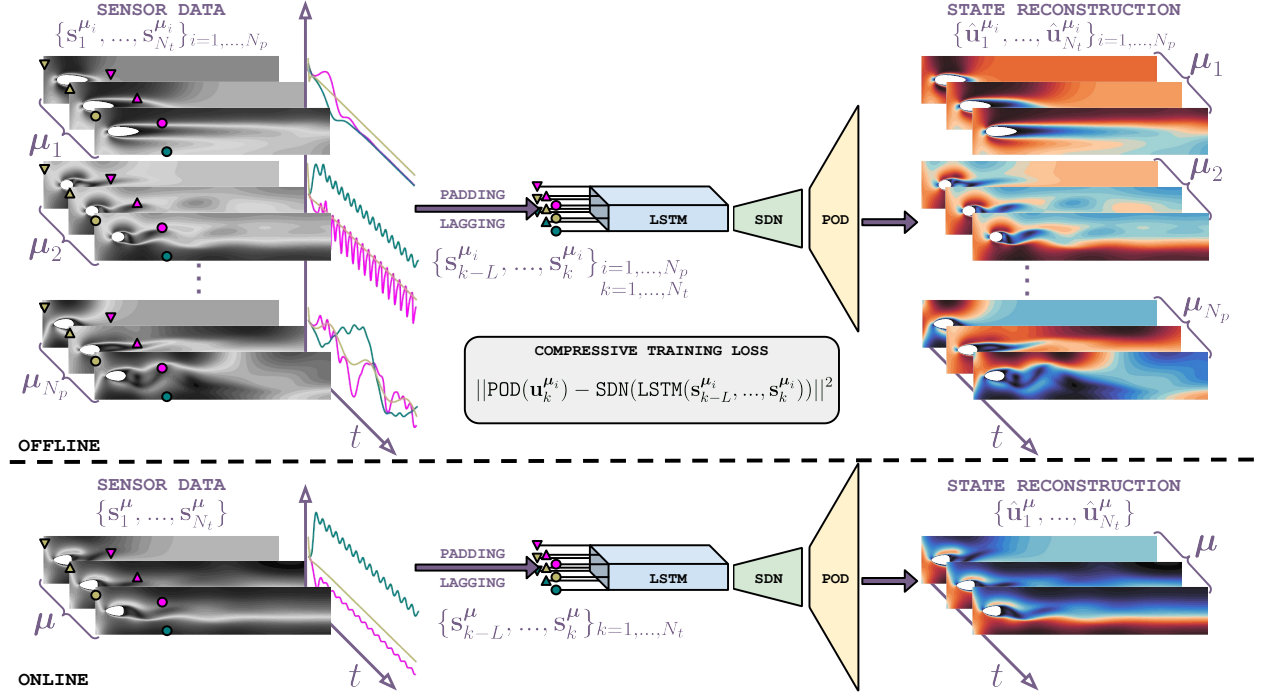


Figure 1: Graphical summary of the *SHallow REcurrent Decoder-based Reduced Order Model* (SHRED-ROM). Sparse sensor values $\mathbf{s}_k^{\mu_i}$ over time windows of length L in multiple scenarios are encoded through a *long short-term memory* (LSTM), while a *shallow decoder network* (SDN) projects the resulting latent representation in the high-dimensional state space. Full-order state snapshots $\mathbf{u}_k^{\mu_i}$ are reduced by *proper orthogonal decomposition* (POD), allowing for compressive training at the POD level. After training, in the online phase, it is possible to reconstruct high-dimensional state trajectories $\hat{\mathbf{u}}_k^{\mu}$ for new parameters μ .

robust ROM architecture for data or high-fidelity simulations. A graphical summary of the SHRED-ROM architecture for parametric state reconstruction from limited sensor data is presented in Figure 1.

PDEs model a diversity of spatio-temporal systems, including those found in the classical physics fields of fluid dynamics, electrodynamics, heat conduction, and quantum mechanics. Indeed, PDEs provide a canonical description of how a system evolves in space and time due the presence of partial derivatives which model the relationships between rates of change of time and space. Governing equations of physics-based systems simply provide a constraint, or restriction, on how these evolving quantities are related. We consider PDEs of the form [13]

$$\dot{u} = N(u, \mathbf{x}, t, \mu) \quad (1)$$

where $u(\mathbf{x}, t)$ is the variable of interest, or the state-space, which we are trying to model. Alternatively, it may be that $N(\cdot)$ is unknown, so a purely data-driven strategy must be considered. The solution $u(\mathbf{x}, t)$ of the PDE depends upon the spatial variable \mathbf{x} , which can be in 1D, 2D or 3D, and time t . Importantly, solutions can often depend on a vector of parameters μ , thus requiring a solution that can model parametric dependencies, i.e. $u(\mathbf{x}, t, \mu)$. Solutions of (1) are typically achieved through numerical simulation, unless $N(\cdot)$ is linear and constant coefficient so that analytic solutions are available. Asymptotic and perturbation methods can also offer analytically tractable solution methods [36]. In many modern applications, discretization of the evolution for $u(\mathbf{x}, t, \mu)$ results in a high-dimensional system for which computations are prohibitively expensive. The goal of building ROMs is to approximate the full simulations of (1) through inexpensive tractable computations.

Traditional ROMs generate a computational proxy of (1) by projecting the governing PDE into a new coordinate system [4, 3, 49, 27]. Coordinate transformations are one of the fundamental paradigms for producing solutions to PDEs [33]. Specifically, ROMs produce coordinate transformations from the simulation data itself. Thus, if the governing PDE (1) is discretized so that $u(\mathbf{x}, t) \rightarrow \mathbf{u}_k = \mathbf{u}(t_k) \in \mathbb{R}^{N_h}$ for

$k = 1, \dots, N_t$, then snapshots of the solution can be collected into the data matrix

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_{N_t} \\ | & | & \cdots & | \end{bmatrix}$$

where $\mathbf{X} \in \mathbb{C}^{N_h \times N_t}$. In case of parametric dependencies, the data matrix contains the discretized states $u(\mathbf{x}, t, \boldsymbol{\mu}) \rightarrow \mathbf{u}_k^\mu = \mathbf{u}(t_k, \boldsymbol{\mu}) \in \mathbb{R}^{N_h}$ for N_p different parametric instances, that is

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | & | \\ \mathbf{u}_1^{\mu_1} & \mathbf{u}_2^{\mu_1} & \cdots & \mathbf{u}_{N_t-1}^{\mu_{N_p}} & \mathbf{u}_{N_t}^{\mu_{N_p}} \\ | & | & \cdots & | & | \end{bmatrix}$$

where $\mathbf{X} \in \mathbb{C}^{N_h \times N_t N_p}$. An optimal coordinate system for ROMs is extracted from the data matrix \mathbf{X} with a *singular value decomposition* (SVD) [35]:

$$\mathbf{X} = \boldsymbol{\Psi} \boldsymbol{\Sigma} \mathbf{V}^*$$

where $\boldsymbol{\Psi} \in \mathbb{C}^{N_h \times r}$, $\boldsymbol{\Sigma} \in \mathbb{R}^{r \times r}$ and $\mathbf{V} \in \mathbb{C}^{N_t N_p \times r}$, with $N_p = 1$ for nonparametric problems, and r is the number of modes extracted to represent the data. The SVD, which is more commonly known in the ROMs community as the *proper orthogonal decomposition* (POD) [30, 4], computes the dominant correlated activity of the data as a set of orthogonal modes. It is guaranteed to provide the best ℓ_2 -norm reconstruction of the data matrix \mathbf{X} for a given number of modes r . Importantly, the r modes $\boldsymbol{\Psi}$ extracted from the data matrix are used to produce a separation of variables solution to the PDE [25]:

$$\mathbf{u} = \boldsymbol{\Psi}(\mathbf{x}) \mathbf{a}(t, \boldsymbol{\mu}) \quad (2)$$

where the vector $\mathbf{a} = \mathbf{a}(t, \boldsymbol{\mu})$ is determined by using this solution form and performing a Galerkin projection of (1) [4]. Thus, a projection-based ROM simply represents the governing PDE evolution (1) in the r -rank subspace spanned by $\boldsymbol{\Psi}$.

The projection-based ROM construction often produces a low-rank model for the dynamics of $\mathbf{a}(t, \boldsymbol{\mu})$ that can be unstable [8]. Machine learning techniques offer a diversity of alternative methods for computing the dynamics in the low-rank subspace. The simplest architecture aims to train a deep neural network that maps the solution forward in time, possibly being robust on different admissible scenario parameters:

$$\mathbf{a}_{k+1}^\mu = \mathbf{f}_\theta(\mathbf{a}_k^\mu) \quad (3)$$

where $\mathbf{a}_k^\mu = \mathbf{a}(t_k, \boldsymbol{\mu})$ and \mathbf{f}_θ represents a deep neural network whose weights and biases θ are determined by minimizing the prediction error on the available reduced snapshots $\mathbf{a}_k^{\mu_i}$ for $i = 1, \dots, N_p$ and $k = 1, \dots, N_t$. A diversity of neural networks can be used to advance solutions, or learn the flow map from time t to $t + \Delta t$ [48, 40]. Indeed, deep learning algorithms provide a flexible framework for constructing a mapping between successive time steps in multiple scenarios. Recently, Parish and Carlberg [47] developed a suite of neural network-based methods for learning time-stepping models for (3), along with an extensive comparison between different neural network architectures and traditional techniques for time-series modeling.

Machine learning algorithms offer options beyond POD-Galerkin projection and deep learning-based modeling of the time-stepping in the variable $\mathbf{a}(t, \boldsymbol{\mu})$. Thus, instead of inserting (2) back into (1) or learning a flow map \mathbf{f}_θ for (3), we can instead think about directly building a model for the dynamics of $\mathbf{a}(t, \boldsymbol{\mu})$, that is

$$\dot{\mathbf{a}} = \mathbf{f}(\mathbf{a}, t, \boldsymbol{\mu})$$

where $\mathbf{f}(\cdot)$ now prescribes the dynamics. Two highly successful options have emerged for producing a model for this dynamical system: (i) the *dynamic mode decomposition* (DMD) [37] and (ii) the *sparse identification of nonlinear dynamics* (SINDy) [7]. The DMD model for $\mathbf{f}(\cdot)$ is assumed to be linear so that

$$\dot{\mathbf{a}} \approx \mathbf{L} \mathbf{a}$$

where \mathbf{L} is a linear operator found by regression. Solutions are then trivial as all that one requires is to find the eigenvalues and eigenvectors of \mathbf{L} in order to provide a general solution by linear superposition. Note

that, in parametric settings, it may be convenient to consider interpolation strategies among a set of linear operators $\mathbf{L}(\boldsymbol{\mu}_i)$ for $i = 1, \dots, N_p$, as proposed by Andreuzzi et al. [2] and Huhn et al. [31]. The SINDy method makes, instead, a different assumption: the dynamics $\mathbf{f}(\cdot)$ can be represented by only a few terms. In this case, the regression has the form

$$\dot{\mathbf{a}} \approx \boldsymbol{\Theta}(\mathbf{a}, \boldsymbol{\mu}) \boldsymbol{\xi}$$

where the columns of the matrix $\boldsymbol{\Theta}(\mathbf{a}, \boldsymbol{\mu})$ are (possibly $\boldsymbol{\mu}$ -dependent) candidate terms from which to construct a dynamical system and $\boldsymbol{\xi}$ contains the loading (or coefficient or weight) of each library term. SINDy assumes that $\boldsymbol{\xi}$ is a sparse vector so that most of the library terms do not contribute to the dynamics. The regression is a simple solve of an overdetermined linear system that is regularized by sparsity, or the sparsity-promoting ℓ_0 or ℓ_1 norms.

In addition to the diversity of methods for building a model for the time dynamics of $\mathbf{a}(t, \boldsymbol{\mu})$, there also exists the possibility of using coordinates other than those defined by $\boldsymbol{\Psi}$. Moving beyond a linear subspace can provide improved coordinate systems for building latent dynamic models. Importantly, there exists the possibility of learning a coordinate system where, for instance, a linear DMD model or a parsimonious SINDy model can be more efficiently imposed. Thus we wish to learn a coordinate transformation

$$\mathbf{z} = \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{u})$$

where \mathbf{z} is the new variable representing the state space \mathbf{u} and $\mathbf{f}_{\boldsymbol{\theta}}$ is a neural network – such as, e.g., the encoder part of an autoencoder – that defines the coordinate transformation. This allows us to find a coordinate system beyond the standard linear, low-rank subspace $\boldsymbol{\Psi}$, which can be advantageous for ROM construction. The goal is then to fit a dynamical model in the new coordinate system

$$\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}, t, \boldsymbol{\mu})$$

through, for instance, a DMD or SINDy approximation [9]. A more recent extension to parametric problems has been proposed by Conti et al. [11], also featuring uncertainty quantification by considering variational autoencoders [12]. Similarly, Regazzoni et al. [50, 51] automatically discover a suitable low-dimensional coordinate system and the corresponding latent dynamics that best describe the input-output data pairs. A more flexible and theoretically rigorous formulation of the problem of latent dynamics learning has been proposed by Farenga et al. [16], whereas alternative approaches to evolve the latent space dynamics using recurrent neural networks or an auto-regressive attention mechanism have been proposed by Koumoutsakos and coauthors [54, 23]. Importantly, in contrast to autoencoders, the decoding-only strategy of SHRED-ROM does not require the computation of inverse pairs, i.e. an encoder and the corresponding decoder. Since the early days in scientific computing, it has been well known that the computation of the inverse of a matrix is highly unstable and non-robust [18, 14, 28]. By decoding only, SHRED-ROM circumvents this problem and learns a single embedding without the corresponding inversion.

Instead of learning a dynamical model at POD or latent level, it is also possible to focus on the map from time-parameters to the POD coefficients $(t, \boldsymbol{\mu}) \rightarrow \mathbf{a}(t, \boldsymbol{\mu})$ or to other low-dimensional coordinate systems $(t, \boldsymbol{\mu}) \rightarrow \mathbf{z}(t, \boldsymbol{\mu})$. This is the case, for instance, of the non-intrusive reduced order modeling frameworks proposed by Hesthaven and Ubbiali [26] and Fresca et al. [22, 20], where long-short term memory networks have also been employed [21].

In what follows, we introduce a new reduced order modeling framework based upon the separation of variable method that combines recurrence and decoding in order to reconstruct high-dimensional states starting from limited sensor measurements, while being agnostic to sensor placement and parameter values. Differently from the aforementioned ROM strategies, we encode the sensor dynamics through a recurrent neural network, mapping its latent representation to the state space with a shallow decoder. The proposed decoding-only strategy turns out to be very efficient and accurate in a wide range of applications, such as the reconstruction of chaotic and nonlinear fluid dynamics in multiple scenarios.

2. SHALLOW RECURRENT DECODER-BASED REDUCED ORDER MODELING

The *shallow recurrent decoder network* (SHRED) proposed by Williams et al. [55] is a promising sensing strategy, with impressive reconstruction and forecasting results in the low-data limit. See [38, 15, 52, 24]

for a complete presentation with possible extensions. In this work, we extend the SHRED strategy to a unified reduced order modeling framework capable of *(i)* reconstructing high-dimensional dynamics from sparse sensor measurements, regardless of sensor placement, *(ii)* dealing with both physical, geometrical and time-dependent parametric dependencies, while being agnostic to their actual values, *(iii)* estimating unknown parameters, and *(iv)* coping with both fixed or mobile sensors, as well as different data sources such as high-fidelity simulations, coupled fields or videos. Importantly, computational efficiency and memory usage are enhanced by reducing the dimensionality of full-order snapshots, allowing for compressive training of the networks.

2.1. The rationale behind SHRED-ROM: separation of variables

SHRED-ROM consists of a combination of recurrence and decoding, as better detailed in Section 2.2. While the former operation captures the temporal behavior of limited sensor data, the latter performs a spatial upscaling to recover the corresponding high-dimensional state. A similar split in spatial and temporal components is taken into account by the well-known method of separation of variables for solving linear PDEs [17], which assumes that the solution has the form $u(\mathbf{x}, t, \boldsymbol{\mu}) = T(t, \boldsymbol{\mu})X(\mathbf{x}, \boldsymbol{\mu})$. For the sake of simplicity, we first review the case of nonparametric problems, postponing at the end of the section the case of parametric PDEs.

Linear PDEs. Let us consider the 1D constant coefficient linear PDE

$$\dot{u} = \mathcal{L}u \quad (4)$$

where $\mathcal{L} = \mathcal{L}(\partial_x, \partial_x^2, \dots)$ is a linear differential operator modeling the dynamics of $u(x, t)$. Note that (4) has to be coupled with suitable initial and boundary conditions in order to guarantee its well-posedness. When looking for solutions in the separated form $u(x, t) = T(t)X(x)$, (4) reduces to two differential equations for, respectively, time $T(t)$ and space $X(x)$, namely

$$\dot{T} = \lambda T; \quad \mathcal{L}X = \lambda X, \quad (5)$$

where $\lambda \in \mathbb{C}$. Finally, thanks to the linearity assumption, the explicit formula for $u(x, t)$ is given by superposition of the solutions of (5), that is

$$u(x, t) = \sum_{n=1}^N a_n \exp(\lambda_n t) \phi_n(x) \quad (6)$$

where $\phi_n(x)$ and λ_n are, respectively, the eigenfunctions and eigenvalues of the linear operator \mathcal{L} , i.e. $\mathcal{L}\phi_n(x) = \lambda_n \phi_n(x)$. Note that the sum in (6) is truncated up to N terms, as standard practice for numerical evaluation. To uniquely determine the coefficients a_n , the initial condition $u(x, 0) = u_0(x)$ is typically imposed in (6), while exploiting the orthogonality of the eigenfunctions $\phi_n(x)$, yielding

$$u_0(x) = \sum_{n=1}^N a_n \phi_n(x) \implies a_n = \langle u_0(x), \phi_n(x) \rangle$$

where $\langle \cdot, \cdot \rangle$ stands for the inner product.

In general, when dealing with real-world high-dimensional problems, it is more likely to measure the dynamical system in few locations over time. Therefore, the coefficients a_n can be determined requiring that the solution (6) matches the available sensor measurements. For instance, if N temporal observations of a single sensor located in $x = x_s$ are available, the system of equations

$$u(x_s, t_k) = \sum_{n=1}^N a_n \exp(\lambda_n t_k) \phi_n(x_s) \quad \text{for } k = 1, \dots, N$$

uniquely determine the solution $u(x, t)$. Similarly, the high-dimensional state may be prescribed by employing multiple sensors over time. For example, if two sensors are available in the considered domain, then $N/2$

temporal observations steps are needed to compute the coefficients a_n . In general, we can exploit N_s sensors on N/N_s trajectory points.

Besides the initial condition and the temporal history of stationary sensors, mobile sensors can be exploited, as shown in Section 3.1 and Section 3.4 with sensors passively transported by the underlying fluid. For example, if we take into account a single mobile sensor with time-dependent position $x_s = x_s(t)$, the constraints become

$$u(x_s(t_k), t_k) = \sum_{n=1}^N a_n \exp(\lambda_n t_k) \phi_n(x_s(t_k))$$

for $k = 1, \dots, N$. The temporal trajectory information at a single spatial location, as well as a mobile sensor, is therefore equivalent to knowing the full spatial field at a given time instant. The above arguments guarantee that, in the linear case, SHRED-ROM reconstructs exactly the high-dimensional spatio-temporal state. Note that the same argument may be easily extended to systems in higher spatial dimensions.

Nonlinear PDEs. Let us consider the 1D nonlinear PDE

$$\dot{u} = N(u) \tag{7}$$

equipped with suitable initial and boundary conditions, where $N = N(\partial_x, \partial_x^2, \dots)$ is a nonlinear differential operator. The dynamical system in (7) is typically solved through numerical techniques such as, e.g., finite differences, finite element methods, finite volume methods and spectral methods [35]. In the latter case, the solution is approximated by a spectral basis expansion

$$u(x, t) = \sum_{n=1}^N a_n(t) \phi_n(x). \tag{8}$$

Standard choices of spectral basis functions $\{\phi_n(x)\}_{n=1}^N$ include, e.g., Fourier modes or Chebyshev polynomials. Inserting (8) back into (7) yields a system of N coupled ODEs for $a_n(t)$:

$$\dot{a}_n = f_n(a_1, a_2, \dots, a_N) \quad \text{for } n = 1, \dots, N. \tag{9}$$

Similarly to the method of separation of variables in the linear limit, the solution of this N -dimensional ODE depends on N constants of integration to be determined by imposing initial condition and orthogonality in (8), that is

$$a_n(0) = \langle u_0(x), \phi_n(x) \rangle.$$

If a fixed sensor monitors the PDE solution in $x = x_s$ over N time steps, we can still uniquely determine the dynamics of the time-dependent coefficients a_n in (9). Specifically, we impose the constraints

$$u(x_s, t_k) = \sum_{n=1}^N a_n(t_k) \phi_n(x_s) \quad \text{for } k = 1, \dots, N$$

to determine the N constants of integration that arise while solving (9), thus obtaining the state approximation through (8). In the nonlinear case, establishing rigorous theoretical bounds for SHRED-ROM poses challenges, consistently with the difficulties encountered to rigorously bound both analytical and numerical solutions in computational PDE settings. However, time-series of sparse sensor measurements allows us to recover standard numerical techniques, such as the spectral method. The same argument can be then extended to multiple or mobile sensors, as detailed for the linear case.

Coupled PDEs. Let us now consider coupled, constant coefficient 1D linear PDEs, such as the first-order coupled system

$$\dot{u} = \mathcal{L}_1 u + \mathcal{L}_2 v \tag{10a}$$

$$\dot{v} = \mathcal{L}_3 u + \mathcal{L}_4 v \tag{10b}$$

in the unknowns $u(x, t)$ and $v(x, t)$, where $\mathcal{L}_i = \mathcal{L}_i(\partial_x, \partial_x^2, \dots)$ for $i = 1, 2, 3, 4$ are linear differential operators. If we differentiate (10a) with respect to time and substitute v and \dot{v} with, respectively, (10a) and (10b), we rewrite the coupled system as a second-order PDE

$$\ddot{u} = \mathcal{L}_1 \dot{u} + \mathcal{L}_2 \mathcal{L}_3 u + \mathcal{L}_2 \mathcal{L}_4 (\mathcal{L}_2^{-1}(\dot{u} - \mathcal{L}_1 u))$$

dependent only on the spatio-temporal field $u(x, t)$. This suggests that the solution fields $u(x, t)$ and $v(x, t)$ can be reconstructed by only knowing $u(x, t)$. For instance, in the application detailed in Section 3.1, we reconstruct the high-dimensional flow velocity exploiting sensors monitoring a coupled quantity, that is the height deviation of the pressure surface from its mean height. Second-order PDEs require both an initial condition $u(x, 0)$ and an initial velocity $\dot{u}(x, 0)$ to uniquely determine the solution. As with previous arguments, we can show that the high-dimensional coupled fields can be prescribed by, e.g., a single sensor measurement over $2N$ temporal trajectory points, as well as by multiple or mobile sensors.

Parametric PDEs. Let us consider the parametric, constant coefficient 1D linear PDE

$$\dot{u} = \mathcal{L}^\mu u \quad (11)$$

with suitable initial and boundary conditions, where $\mathcal{L}^\mu = \mathcal{L}(\partial_x, \partial_x^2, \dots, \mu)$ is a parameters-dependent linear differential operator. As with the linear nonparametric case, looking for solutions of (11) in the separated form $u(x, t, \mu) = T(t, \mu)X(x, \mu)$ leads to the differential equations

$$\dot{T} = \lambda T; \quad \mathcal{L}^\mu X = \lambda X,$$

with $\lambda \in \mathbb{C}$. The parametric high-dimensional state $u(x, t, \mu)$ is then given by

$$u(x, t, \mu) = \sum_{n=1}^N a_n \exp(\lambda_n^\mu t) \phi_n^\mu(x)$$

where $\phi_n^\mu(x)$ and λ_n^μ are, respectively, the eigenfunctions and eigenvalues of the parametric linear operator \mathcal{L}^μ , i.e. $\mathcal{L}^\mu \phi_n^\mu(x) = \lambda_n^\mu \phi_n^\mu(x)$. As with previous arguments, the coefficients a_n and thus the solution $u(x, t, \mu)$ can be uniquely determined by employing a fixed sensor in position $x = x_s$ monitoring the state on N trajectory points in the scenario identified by the set of parameters μ , that is

$$u(x_s, t_k, \mu) = \sum_{n=1}^N a_n \exp(\lambda_n^\mu t_k) \phi_n^\mu(x_s) \quad \text{for } k = 1, \dots, N.$$

The same arguments can be generalized to multiple or mobile sensors, to nonlinear and coupled parametric PDEs, as detailed in the previous cases, as well as to systems in higher spatial dimensions or *ordinary differential equations* (ODEs), as highlighted in the following example.

Example 1. Let us consider the 3×3 parametric linear system of ODEs

$$\dot{\mathbf{u}} = \begin{bmatrix} 1 & 2 & 0 \\ 0 & -1 & \mu \\ 0 & 0 & 2 \end{bmatrix} \mathbf{u} \quad (12)$$

where $\mathbf{u}(t, \mu) = [u_1(t, \mu), u_2(t, \mu), u_3(t, \mu)]^\top$. The explicit solution of (12) is given by

$$\begin{bmatrix} u_1(t, \mu) \\ u_2(t, \mu) \\ u_3(t, \mu) \end{bmatrix} = c_1 e^{2t} \begin{bmatrix} 2\mu \\ \mu \\ 3 \end{bmatrix} + c_2 e^{-t} \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} + c_3 e^t \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

where the coefficients c_1, c_2, c_3 has to be determined. To this aim, we may exploit the initial condition $\mathbf{u}(0, \mu) = [1, 2, 3]^\top$, that is we solve the system

$$\begin{bmatrix} u_1(0, \mu) \\ u_2(0, \mu) \\ u_3(0, \mu) \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 2\mu & -1 & 1 \\ \mu & 1 & 0 \\ 3 & 0 & 0 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$$

yielding $c_1 = 1$, $c_2 = 2 - \mu$ and $c_3 = 3 - 3\mu$. Equivalently, we can determine the coefficients by employing one sensor monitoring u_1 over time. Specifically, if we measure $[u_1(0, 1), u_1(1, 1), u_1(2, 1)]^\top = [1, e^2 - e^{-1}, 2e^4 - e^{-2}]$ in the scenario identified by $\mu = 1$, we obtain the same results by solving the system

$$\begin{bmatrix} u_1(0, 1) \\ u_1(1, 1) \\ u_1(2, 1) \end{bmatrix} = \begin{bmatrix} 1 \\ 2e^2 - e^{-1} \\ 2e^4 - e^{-2} \end{bmatrix} = \begin{bmatrix} 2 & -1 & 1 \\ 2e^2 & -e^{-1} & e \\ 2e^4 & -e^{-2} & e^2 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$$

The same strategy can be then repeated for different values of the parameter μ . Note that, if we consider measurements of u_2 and/or u_3 only, the coefficient c_3 does not appear in the corresponding system of constraints, and it is not possible to uniquely determine the exact solution. This limitation is due to the non-observability of the system with respect to u_2 and/or u_3 . Thus, in this setting, full state reconstruction from sensor data is guaranteed whenever we have observability with respect to the measured quantity.

2.2. SHRED-ROM architecture and compressive training

SHRED-ROM aims to reconstruct a high-dimensional spatio-temporal field starting from N_s sensors in multiple scenarios. Rather than learning the *one-shot* reconstruction map $\mathbf{s}_k^\mu = \{u(\mathbf{x}_s, t_k, \boldsymbol{\mu})\}_{s=1}^{N_s} \rightarrow \mathbf{u}_k^\mu$ for $k = 1, \dots, N_t$ – as considered by, e.g., [45, 44, 43, 56] – we take advantage of the temporal history of sensor values. Specifically, a recurrent neural network \mathbf{f}_T encodes the sensor measurements over a time window of length $L \leq N_t$ into a latent representation of dimension N_l , that is

$$\mathbf{f}_T : \underbrace{\mathbb{R}^{N_s} \times \dots \times \mathbb{R}^{N_s}}_{L \text{ times}} \rightarrow \mathbb{R}^{N_l}, \quad \mathbf{h}_k^\mu = \mathbf{f}_T(\mathbf{s}_{k-L}^\mu, \dots, \mathbf{s}_k^\mu),$$

where the hyperparameter L identifies the number of lags considered by SHRED-ROM, and may be selected according to the problem-specific evolution rate or periodicity. The high-dimensional state is then approximated through a decoder \mathbf{f}_X , which performs a nonlinear projection of the latent representation onto the state space

$$\mathbf{f}_X : \mathbb{R}^{N_l} \rightarrow \mathbb{R}^{N_h}, \quad \mathbf{u}_k^\mu \approx \hat{\mathbf{u}}_k^\mu = \mathbf{f}_X(\mathbf{h}_k^\mu) = \mathbf{f}_X(\mathbf{f}_T(\mathbf{s}_{k-L}^\mu, \dots, \mathbf{s}_k^\mu)).$$

In this work, we consider a *long short-term memory* (LSTM) network [29] to model the temporal dependency of sensor data, and a *shallow decoder network* (SDN) as latent-to-state map. In general, SHRED-ROM provides a modular framework where different architectures – such as, e.g., convolutional neural networks [39], gated recurrent units [10], echo state networks [42], transformers [53] and variational autoencoders [34] – may be exploited.

In the offline phase, once for all, the neural networks appearing in SHRED-ROM have to be properly trained. To this aim, suppose to collect snapshots of the high-dimensional state – either in the form of experimental measurements or synthetic data generated by high-fidelity simulations – along with the corresponding sensor data

$$\mathbf{u}_k^{\mu_i} = \mathbf{u}(t_k, \boldsymbol{\mu}_i) \in \mathbb{R}^{N_h}, \quad \mathbf{s}_k^{\mu_i} = \{u(\mathbf{x}_s, t_k, \boldsymbol{\mu}_i)\}_{s=1}^{N_s} \in \mathbb{R}^{N_s}$$

for different time instants t_1, \dots, t_{N_t} and scenario parameters $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{N_p}$ in the parameter space \mathcal{P} . Note that, as discussed in the previous section, mobile sensors can be easily taken into account, resulting in the measurements $\mathbf{s}_k^{\mu_i} = \{u(\mathbf{x}_s(t_k), t_k, \boldsymbol{\mu}_i)\}_{s=0}^{N_s} \in \mathbb{R}^{N_s}$. Note also that, in order to strengthen the generalization capabilities of SHRED-ROM, it is crucial to adequately explore the variability in time and in the parameter space. Starting from the sequences of sensor measurements in multiple scenarios, we extract time-series of length L

$$\{\mathbf{s}_{k-L}^{\mu_i}, \dots, \mathbf{s}_k^{\mu_i}\} \quad \text{for } i = 1, \dots, N_p \text{ and } k = 1, \dots, N_t \quad (13)$$

where pre-padding ($\mathbf{s}_k^{\mu_i} = \mathbf{0}$ for $k \leq L$ and $i = 1, \dots, N_p$) is applied to obtain state reconstructions on the whole time interval $[t_1, t_{N_t}]$, avoiding any burn-in period. After a training-validation-test splitting of the available input-output pairs, the recurrent neural network \mathbf{f}_X and the shallow decoder \mathbf{f}_T are trained by minimizing the reconstruction error

$$J(\mathbf{u}_k^{\mu_i}, \hat{\mathbf{u}}_k^{\mu_i}) = \sum_{\substack{i \in I_{\text{train}} \\ k \in K_{\text{train}}}} \|\mathbf{u}_k^{\mu_i} - \hat{\mathbf{u}}_k^{\mu_i}\|^2 = \sum_{\substack{i \in I_{\text{train}} \\ k \in K_{\text{train}}}} \|\mathbf{u}_k^{\mu_i} - \mathbf{f}_X(\mathbf{f}_T(\mathbf{s}_{k-L}^{\mu_i}, \dots, \mathbf{s}_k^{\mu_i}))\|^2, \quad (14)$$

where I_{train} and K_{train} identify the times and scenarios in the training set, while $\|\cdot\|$ stands for the Euclidean norm. Once trained, in the online phase, SHRED-ROM provides a real-time reconstruction of the state trajectory starting from the corresponding sensor history in new time instants and/or new scenarios μ unseen during training.

Whenever the state dimension N_h is remarkably high, compressive training strategies may be employed to enhance computational efficiency and memory usage. Specifically, it is possible to reduce the state dimensionality through a data- or physics-driven basis expansion of the snapshots $\mathbf{u}_k^{\mu_i}$. Doing so, SHRED-ROM estimates only the $r \ll N_h$ basis expansion coefficients, rather than the entire high-dimensional state. For example, we can project the state snapshots onto a lower-dimensional subspace by POD, that is $\mathbf{u}_k^{\mu_i} = \Psi \mathbf{a}_k^{\mu_i}$, resulting in the SHRED reconstruction

$$\hat{\mathbf{u}}_k^{\mu_i} = \Psi \hat{\mathbf{a}}_k^{\mu_i} = \Psi \mathbf{f}_X(\mathbf{f}_T(\mathbf{s}_{k-L}^{\mu_i}, \dots, \mathbf{s}_k^{\mu_i})).$$

An example of an alternative basis is provided in Section 3.1, where spherical harmonics are taken into account to properly describe and compress synthetic state snapshots simulated on a sphere.

3. NUMERICAL RESULTS

This section presents five test cases where SHRED-ROM is employed to reconstruct high-dimensional spatio-temporal fields starting from limited sensor measurements. In Section 3.1, we employ SHRED-ROM to reconstruct synthetic spherical snapshots in a nonparametric setting, while exploiting both fixed and mobile sensors, as well as different compressive training strategies based on POD and spherical harmonics. To demonstrate the wide applicability of SHRED-ROM, we cope with the reconstruction of videos given limited pixel values in Section 3.2. The remaining three test cases are devoted to the reconstruction of high-dimensional states of chaotic, nonlinear and parametric fluid dynamics in multiple scenarios. Note that, when dealing with nonparametric problems, 80% of the time instants are randomly selected as training sequences in K_{train} , while the remaining ones are equally split in validation and test. Instead, in parametric contexts, we consider a parameter-wise splitting – i.e. 80% of the trajectories are regarded as training set, while the remaining trajectories related to different parameter values constitute validation and test sets.

Differently to many deep learning-based models, SHRED-ROM requires minimal hyperparameter tuning. Indeed, in all the test cases presented, the following lightweight architecture is exploited: the LSTM \mathbf{f}_T shows 2 hidden layers with 64 neurons each, while the SDN \mathbf{f}_X is made of 2 hidden layers having 350 and 400 neurons, respectively. Moreover, we select ReLU as activation function, and we prevent overfitting through dropout with rate equal to 0.1. Regarding the neural networks training, we exploit Adam optimizer for 200 epochs, half with learning rate equal to 0.001 and half with learning rate equal to 0.0001, considering a batch size equal to 64. The relatively light and simple SHRED-ROM architecture, along with compressive training strategies, allows us to efficiently solve a wide range of challenging reconstruction problems with laptop-level computing.

The generalization capabilities of SHRED-ROM on new times and new scenarios in the test set are quantitatively assessed with the following mean relative error

$$\varepsilon(\mathbf{u}_k^{\mu_i}, \hat{\mathbf{u}}_k^{\mu_i}) = \frac{1}{N_{\text{test}}} \sum_{\substack{i \in I_{\text{test}} \\ k \in K_{\text{test}}}} \frac{\|\mathbf{u}_k^{\mu_i} - \hat{\mathbf{u}}_k^{\mu_i}\|}{\|\mathbf{u}_k^{\mu_i}\|} = \frac{1}{N_{\text{test}}} \sum_{\substack{i \in I_{\text{test}} \\ k \in K_{\text{test}}}} \frac{\|\mathbf{u}_k^{\mu_i} - \mathbf{f}_X(\mathbf{f}_T(\mathbf{s}_{k-L}^{\mu_i}, \dots, \mathbf{s}_k^{\mu_i}))\|}{\|\mathbf{u}_k^{\mu_i}\|},$$

where I_{test} and K_{test} identify the time and parameters in the test set, while N_{test} is the test set cardinality.

3.1. Shallow Water Equations

The first test case we consider deal with the reconstruction of synthetic spherical data in a nonparametric setting. The snapshots correspond to the solution of the Shallow Water Equations (SWE) on a sphere with earth-like topography and daily/annual periodic external forces, available in the dataset *The Well* [46]. The SWE solution includes three high-dimensional fields ($N_h = 131072$), corresponding to the height deviation of pressure surface from its mean height, and to the velocity components with respect to the polar and azimuthal angles, over a year ($N_t = 1008$ time steps).

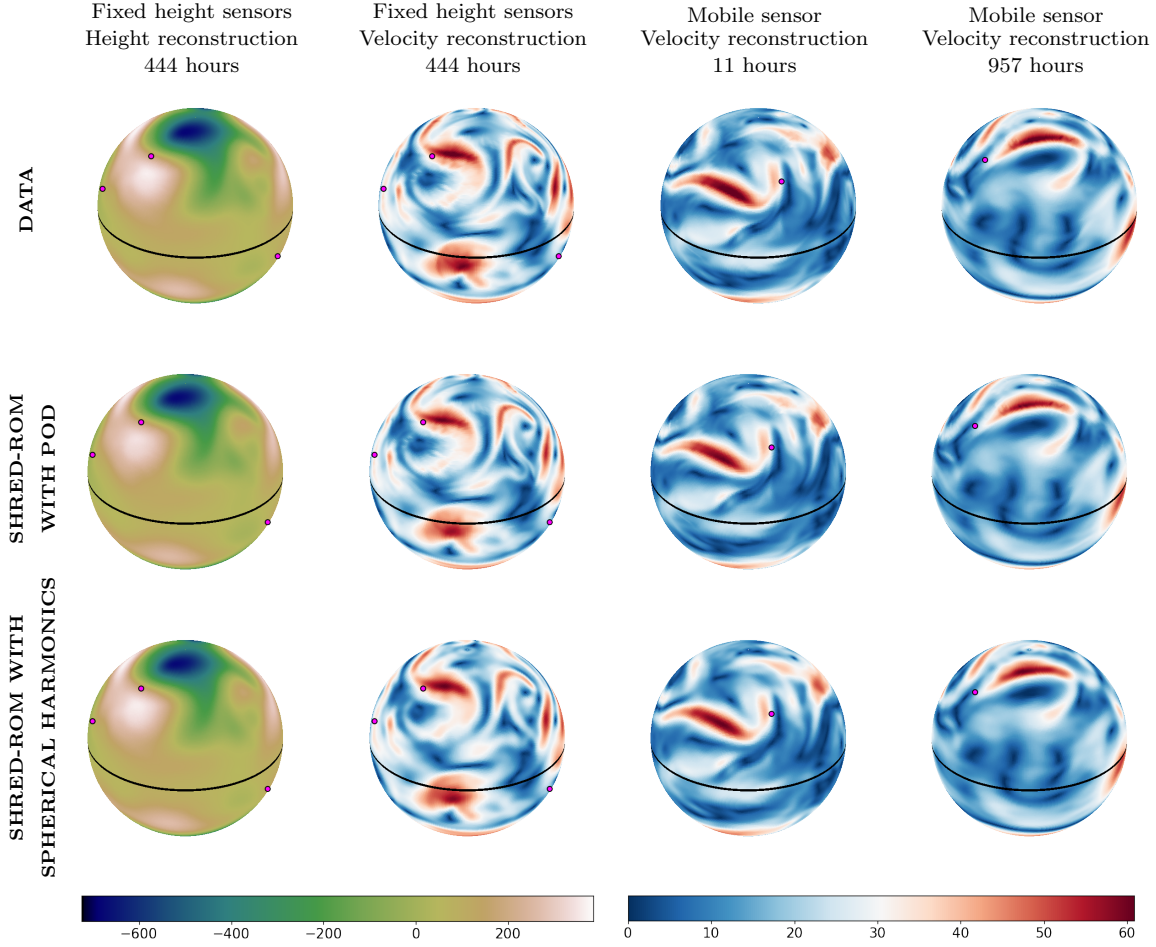


Figure 2: *Shallow Water Equations*. Ground truth (first row), SHRED-ROM reconstructions with POD-based compressive training (second row) and SHRED-ROM reconstructions with spherical harmonics-based compressive training (third row). The following test cases are considered: height reconstruction from 3 fixed height sensors (first column); velocity reconstruction from 3 fixed height sensors (second column); velocity reconstruction from the coordinates of 1 mobile sensor at $t = 11$ hours (third column) and $t = 957$ hours (fourth column). The sparse sensors exploited by SHRED-ROM are depicted with magenta dots.

We first consider randomized SVD to reduce the data dimensionality and allow for compressive training. With only $r = 50$ POD modes (compression ratio equal to 99.96%), it is possible to compress the height snapshots with a relative reconstruction error on test data equal to 1.34%. Instead, $r = 75$ are retrieved for the velocity components reductions (compression ratio equal to 99.94%), with relative test errors equal 1.71% and 0.91%. Note that the number of POD modes to retrieve for accurate compressions is problem-dependent, and can be determined by looking at the singular values decay.

SHRED-ROM is then applied to reconstruct the high-dimensional spatio-temporal fields starting from 3 fixed sensors at as many random locations. To show the wide applicability of SHRED-ROM with coupled fields, we suppose that the sensors can monitor the height variable only, without having access to the velocity values. After data preprocessing with a lag $L = 100$ and neural networks compressive training, SHRED-ROM can accurately reconstruct the trajectories of the three high-dimensional fields starting from the limited sensors randomly chosen. Specifically, the reconstruction errors on the time instants unseen during training are equal to 3.14%, 9.59% and 5.90% for, respectively, the height variable and the polar and azimuthal components of the velocity.

The system configuration may also be captured by mobile sensors, such as drifting buoys. In particular, we consider 1 sensor passively transported by the velocity field on the sphere. Starting only from the history of the sensor coordinates, thus without monitoring any quantity, SHRED-ROM accurately predicts the

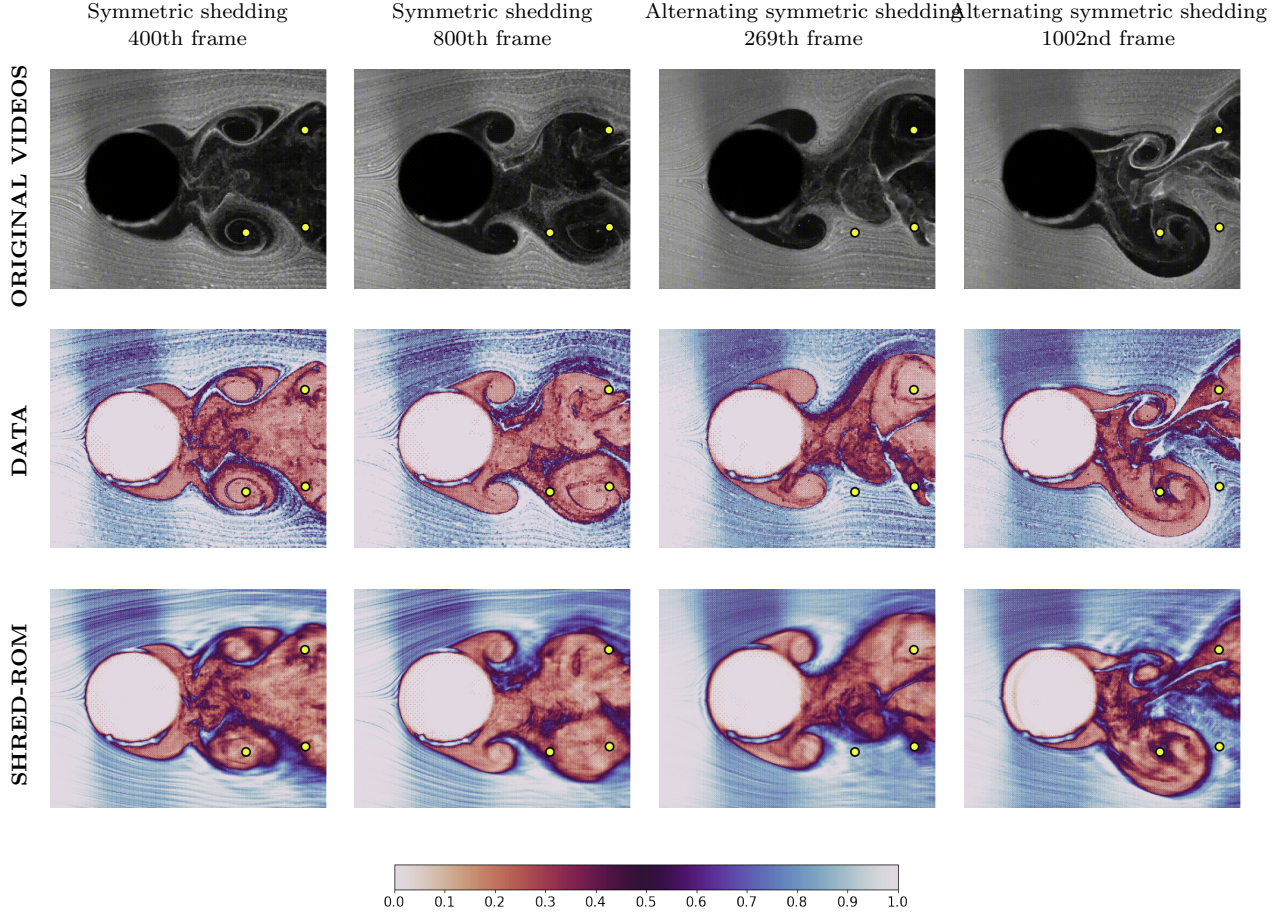


Figure 3: *GoPro Physics*. Original videos (first row), preprocessed data (second row) and SHRED-ROM reconstructions from 3 fixed pixels in the cylinder wake (first column). The following test cases are considered: symmetric shedding reconstruction of the 400th (first column) and the 800th frame (second column); alternating symmetric shedding reconstruction of the 269th (third column) and the 1002nd frame (fourth column). The sparse pixels exploited by SHRED-ROM are depicted with yellow dots.

velocity components, with test relative errors equal to 7.43% and 4.80%, respectively.

Instead of considering data-driven compression techniques for the snapshots, such as POD, we can also exploit a different physics-based basis of functions that properly capture the variability of the SWE solution. In this context, we employ the spherical harmonics up to degree 50, resulting in a dimensionality reduction with $r = 2601$ and with test reconstruction errors equal to 0.25%, 9.45% and 6.37% for the height and velocity components, respectively. By measuring the height values in 3 random locations over time, SHRED-ROM can still accurately predict the spherical harmonics expansion coefficients. The reconstruction errors committed on the test snapshots of height and velocity components are, respectively, 2.15%, 10.71% and 7.13%. If, instead, we take into account a mobile sensor passively transported by the velocity field, SHRED-ROM reconstructs the high-dimensional polar and azimuthal components of the velocity with test errors equal to 11.64% and 7.68%.

Figure 2 shows four height and velocity magnitude snapshots in the test set, along with the corresponding SHRED-ROM reconstructions obtained considering both POD and spherical harmonics as compression techniques. By a visual inspection, it is possible to assess the high level of accuracy of SHRED-ROM despite the limited input information, that are 3 fixed sensors capturing the height variable or the temporal history of the mobile sensor coordinates.

3.2. GoPro Physics

The second test case focuses on reconstructing videos of experimental studies in fluid dynamics, typically referred to as *GoPro physics*, in contrast with the other examples where we exploit synthetic snapshots simulated through high-fidelity solvers. Specifically, we consider two videos recording the vortex shedding in the wake of an oscillating cylinder provided by Boersma et al. [5]. Due to different inflow velocities considered, two different patterns are displayed, that are the so-called *symmetric* and *alternating symmetric shedding*. Starting from sparse pixels values, we aim to build a SHRED-ROM capable of reconstructing the videos frames, while discriminating between the two scenarios.

Each video consists of $N_t = 1077$ frames of dimension 400×504 ($N_h = 201600$). In the preprocessing stage, the original frames are converted to gray scale, the background is removed and the pixel values are normalized. POD is then applied to allow for compressive training. Specifically, retaining the first $r = 100$ POD coefficients for every frame with randomized SVD (that is a compression ratio equal to 99.95%), we achieve a POD reconstruction error equal to 11.96%.

SHRED-ROM is taken into account to reconstruct the frames starting from 3 pixels randomly sampled within the region behind the cylinder. With a lag parameter equal to $L = 150$, SHRED-ROM accurately approximates the test snapshots with a relative error equal to 13.17%. In the direction of parametric settings, it is interesting to note that SHRED-ROM is able to automatically discriminate among the two videos from the pixel values, without requiring any knowledge about the underlying setting. Figure 3 presents 4 test snapshots of the original videos, along with the corresponding preprocessed data and the SHRED-ROM reconstructions.

3.3. Kuramoto-Sivashinsky Equation

In this section, we cope with a nonlinear and chaotic fluid dynamics exhibiting parametric dependency. In particular, we model the dynamics of a 1D state variable $u : [0, L] \times [0, T] \rightarrow \mathbb{R}$ with the Kuramoto-Sivashinsky equation (KS)

$$u_t + u_{xx} + \nu u_{xxx} + uu_x = 0$$

with periodic boundary conditions. To deal with chaotic patterns, we set $L = 22$ and $T = 200$. Regarding the initial condition, we consider

$$u(x, 0, \omega) = \cos\left(\frac{2\pi\omega}{L}x\right) \left(1 + \sin\left(\frac{2\pi\omega}{L}x\right)\right)$$

The viscosity ν and the frequency of the initial datum are regarded as parameters, that is $\boldsymbol{\mu} = [\nu, \omega]^\top$. As discussed in Section 2.2, our goal is to employ SHRED-ROM to rapidly and accurately approximate the state trajectory for a new set of parameters, unseen in the offline phase. To this aim, we generate 500 trajectories corresponding to parameters ν and ω randomly sampled in the intervals $[1, 2]$ and $[1, 5]$, respectively. Note that an adequate exploration of the parameter space is crucial due to the chaotic and instable patterns given by KS, where small changes in the parameter values result in completely different solutions. To solve the KS equation, we consider the Exponential Time Differencing fourth-order Runge-Kutta (ETDRK4) method [32] with a uniform space discretization of $N_h = 100$ spatial points and a time step equal to $\Delta t = 0.01$. Finally, starting from the initial values at $t = 0$, we save the KS solution with a frequency of 100, resulting in trajectories of $N_t = 201$ snapshots. The dimensionality of the generated data is then reduced through randomized SVD: $r = 20$ POD modes are sufficient to capture most of the variability (compression ratio equal to 80%), with a relative error on test data equal to 0.35%.

After generating and compressing the data, it is now possible to train SHRED-ROM. We consider 2 fixed sensors randomly sampled in the domain and a lag parameter equal to $L = 50$. When reconstructing the state trajectories corresponding to new parameters unseen during training, the relative error with respect to the ground truth is equal to 9.13%. Note that only the sensor values are exploited as input, disregarding any information about the parameters values – which, in general, may be unknown or uncertain – as typically considered by other ROM strategies. The performance of SHRED-ROM can be also visually assessed in Figure 4, panel (a), where we compare three test spatio-temporal KS solutions along with the corresponding SHRED-ROM approximations.

The panel (b) of Figure 4 presents, instead, an analysis of the SHRED-ROM test reconstruction errors with respect to the lag parameter L (while considering 2 fixed sensors) and the number of sensors (with lag

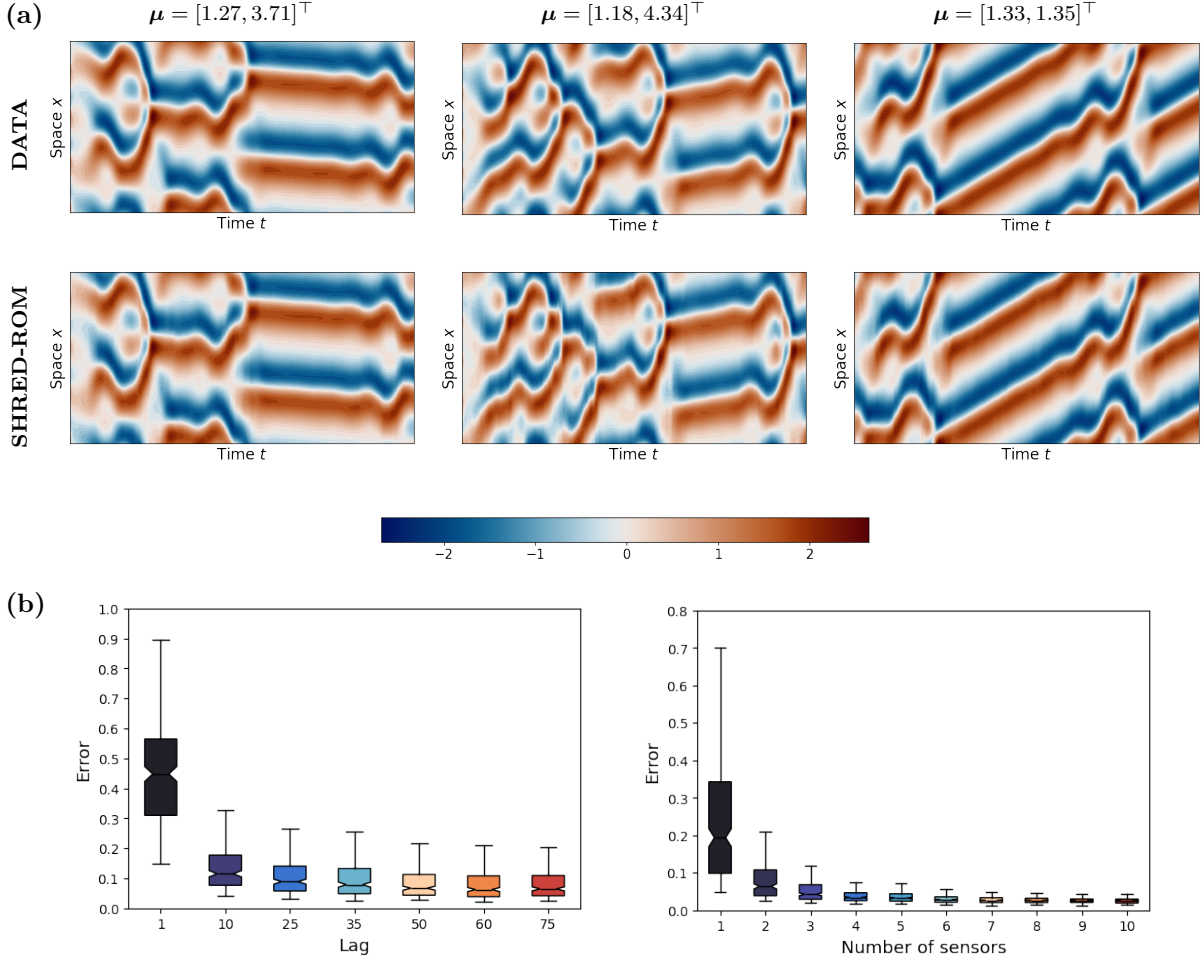


Figure 4: *Kuramoto-Sivashinsky*. (a) Ground truth (first row) and SHRED-ROM reconstructions (second row). The following test cases are considered: state reconstruction from 2 fixed sensors with $\mu = [1.27, 3.71]^\top$ (first column), $\mu = [1.18, 4.34]^\top$ (second column) and $\mu = [1.33, 1.35]^\top$ (third column). (b) Relative reconstruction errors committed by SHRED-ROM on test data for different lag values (first column) and different number of sensors (second column), while considering different sensor placements.

$L = 50$). To show that the proposed architecture is agnostic to sensor placement, every setting takes into account the errors committed by 10 different SHRED-ROMs trained on as many sensor configurations. The fast decay of the error distributions suggest that even a small temporal history, as well as a tiny amount of sensors, are enough to achieve acceptable results. Note that the case with unitary lag ($L = 1$) corresponds to the one-shot reconstruction through a shallow recurrent decoder with compressive training, equivalent to the *POD-based deep state estimation* (PDS) proposed by Nair and Goza [45].

Sensor measurements in real-world problems are typically corrupted by noise. If we corrupt the sensor values with Gaussian random noise with zero mean and standard deviation equal to 0.25 (which corresponds approximately to the 5% of the state data range), the generalization capabilities of SHRED-ROM in test scenarios get worse, with a relative reconstruction error equal to 23.24%. In this case, thanks to the lightweight SHRED-ROM architecture and the efficient compressive training strategy, we can consider an ensemble of SHRED-ROMs in order to provide a robust state estimation, as proposed by Riva et al. [52]. We thus consider 20 different SHRED-ROMs trained on 2 sensor trajectories corrupted by as many Gaussian random noises. The state reconstruction can be now obtained by averaging the predictions of every SHRED-ROM, with a mean relative reconstruction errors on test data equal to 15.51%

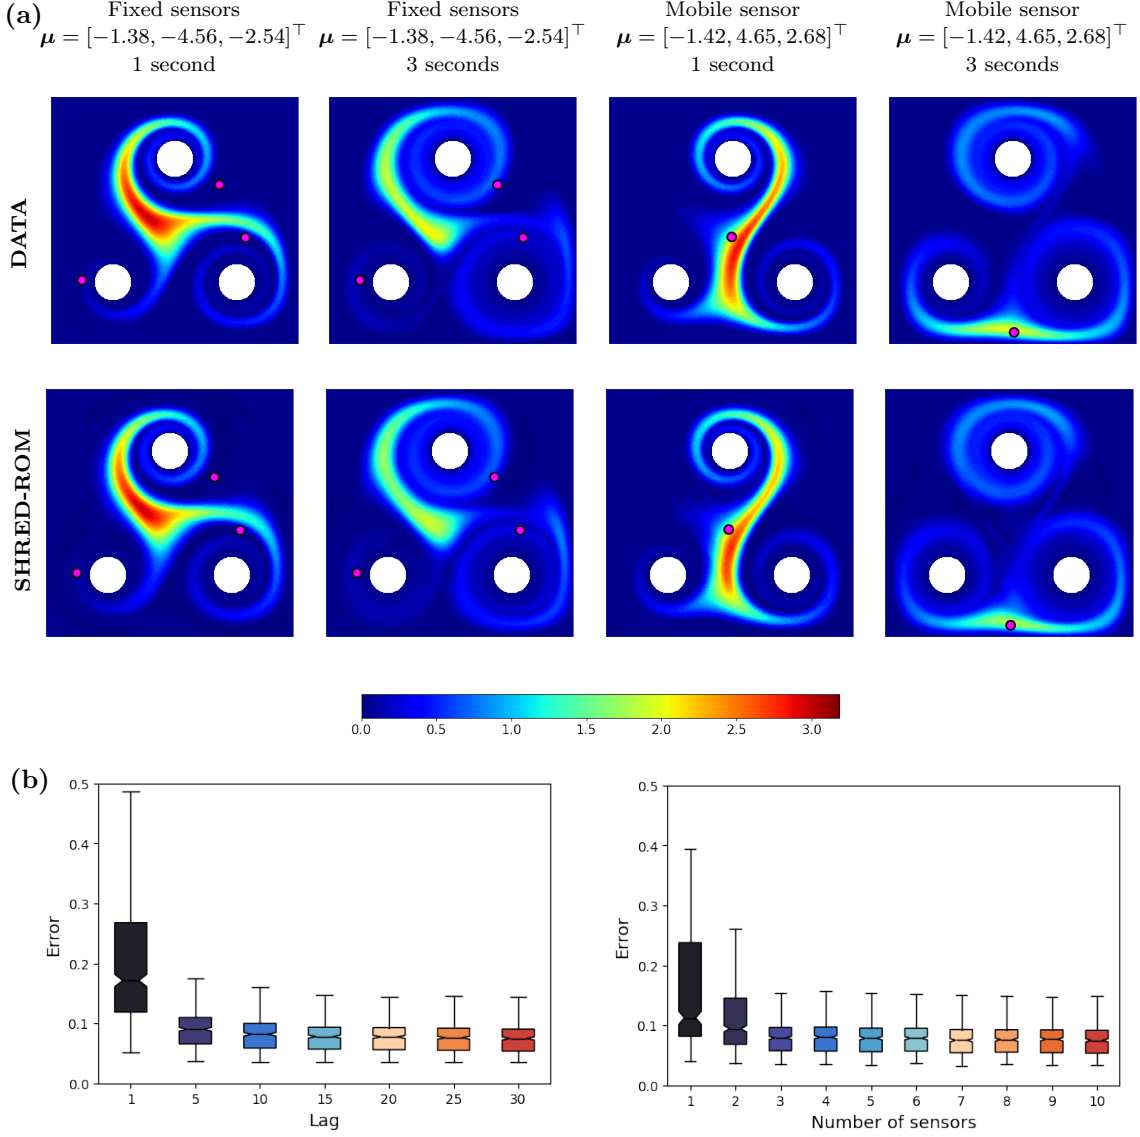


Figure 5: *Fluidic Pinball*. (a) Ground truth (first row) and SHRED-ROM reconstructions (second row). The following test cases are considered: state reconstruction from 3 fixed sensors with $\mu = [-1.38, -4.56, -2.54]^\top$ at $t = 1$ seconds (first column) and $t = 3$ seconds (second column); state reconstruction from the coordinates of 1 mobile sensor with $\mu = [-1.42, 4.65, 2.68]^\top$ at $t = 1$ seconds (third column) and $t = 3$ seconds (fourth column). (b) Relative reconstruction errors committed by SHRED-ROM on test data for different lag values (first column) and different number of sensors (second column), while considering different sensor placements. The sparse sensors exploited by SHRED-ROM are depicted with magenta dots.

3.4. Fluidic Pinball

The fourth test case copes with an involved 2D advection-diffusion problem characterized by an implicit parametric dependence. We consider an incompressible fluid constrained in a square domain $(-1, 1)^2$ with three cylinders centered at $(-0.5, -0.5)$, $(0.5, -0.5)$ and $(0.0, 0.5)$, and with radius 0.15. The cylinders can rotate with constant velocities v_i for $i = 1, 2, 3$ – here regarded as parameters $\mu = [v_1, v_2, v_3]^\top$ – resulting in a fluid motion inside the box. Specifically, the fluid velocity $\mathbf{v} : [-1, 1]^2 \rightarrow \mathbb{R}^2$ and pressure $p : [-1, 1]^2 \rightarrow \mathbb{R}$ are determined by the steady Navier-Stokes equations

$$\begin{cases} -\nu \Delta \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} + \nabla p = 0 \\ \nabla \cdot \mathbf{v} = 0 \end{cases}$$

with viscosity $\nu = 1.0$, no-slip boundary conditions on the external walls and Dirichlet boundary conditions on the three cylinders.

Let us now consider a time-dependent quantity $y : [-1, 1]^2 \times [0, T] \rightarrow \mathbb{R}$ – which may represent, for instance, heat, mass or density of particles – that spreads and moves in the domain according to the advection-diffusion PDE

$$y_t + \nabla \cdot (-\eta \nabla y + \mathbf{v}(\boldsymbol{\mu})y) = 0$$

with homogeneous Neumann boundary conditions and initial condition

$$y(\mathbf{x}, 0) = \frac{10}{\pi} \exp(-10x_1^2 - 10x_2^2).$$

Therefore, starting from a Gaussian function with mean $(0, 0)$ and variance 0.05 , the quantity y is transported and deformed by the parametric fluid flow with velocity $\mathbf{v}(\boldsymbol{\mu})$. To mainly focus on the advection effect over the diffusion one, we set the viscosity $\eta = 0.001$. Moreover, we set the final time $T = 3$.

SHRED-ROM is now exploited to provide the state evolution for new combinations of cylinder velocities, unseen during training. To train the LSTM \mathbf{f}_T and the SDN \mathbf{f}_X , we generate 500 trajectories for different parameter values in the parameter space $\mathcal{P} = [-5, 5]^3$. Note that both clockwise and counter-clockwise rotations are considered according to the signs of the sampled parameters. Both the steady Navier-Stokes equations and the advection-diffusion PDE are solved with finite element solvers in **fenics** [1], taking into account a time step $\Delta t = 0.1$ – resulting in state trajectories of length $N_t = 31$ – and a mesh with $N_h = 7525$ vertices. The dimensionality of the simulated snapshots is then reduced to $r = 200$ thanks to randomized SVD. Despite a compression ratio equal to 97%, the low-dimensional features are enough to accurately reconstruct the high-dimensional data with a test error equal to 1.54%.

We now build two SHRED-ROMs to deal with, respectively, fixed and mobile sensors. In the first case, we measure the state y at 3 fixed locations within the domain. In the second setting, we suppose to know only the spatial coordinates of 1 mobile sensor located in $(0, 0)$ at $t = 0$ and passively transported by the fluid flow. Note that, to show the versatility of the method, we suppose to know the parameter values $\boldsymbol{\mu}$, and we thus exploit this information as input, in addition to the sensor values. After training the corresponding LSTM networks and the SDNs with a time window of length $L = 10$, both SHRED-ROMs are able to reconstruct the state patterns in new scenarios with mean relative errors equal to, respectively, 7.90% and 8.44%. The panel (a) of Figure 5 displays four different test snapshots along with the corresponding reconstructions provided by SHRED-ROM, both taking into account fixed and mobile sensors. The panel (b) of Figure 5 shows, instead, the behavior of the mean relative error on test data with respect to different lag values and different number of fixed sensors, while taking into account 10 different random sensor placements in each setting. Similarly to the findings in Section 3.3, a small time window L and a very limited amount of sensors are enough to achieve accurate results. In particular we note that, in both the parametric applications presented so far, accurate results are obtained by exploiting $d + 1$ fixed sensors, where d is the problem dimension ($d = 1$ in the Kuramoto-Sivashinsky setting, $d = 2$ in the fluidic pinball one). Similarly, an object in a d -dimensional space can be localized by $d + 1$ constraints, such as $d + 1$ sensors measuring the distance to that object.

3.5. Flow Around an Obstacle

The last test case is devoted to the reconstruction of fluid flows where both time-dependent, physical and geometrical parametric dependencies are employed. Let us consider a 2D channel $[0, 10] \times [0, 2]$ with a circular obstacle centered in $(1, 1)$ and with radius equal to 0.2 . Through Radial Basis Function (RBF) interpolation, it is possible to deform the reference setting in order to obtain different obstacle shapes, without changing the number of degrees of freedom N_h . In particular, we lengthen the circle to the left and right by, respectively, γ_l and γ_r , so that the surface of the deformed obstacle passes through the points $(0.8 - \gamma_l, 1.0)$, $(1.0, 1.2)$, $(1.2 + \gamma_r, 1.0)$ and $(1.0, 0.8)$. Let us also consider an incompressible fluid flow around the obstacle, whose dynamics is described by the unsteady Navier-Stokes equations

$$\begin{cases} \mathbf{v}_t - \nu \Delta \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} + \nabla p = 0 \\ \nabla \cdot \mathbf{v} = 0 \end{cases}$$

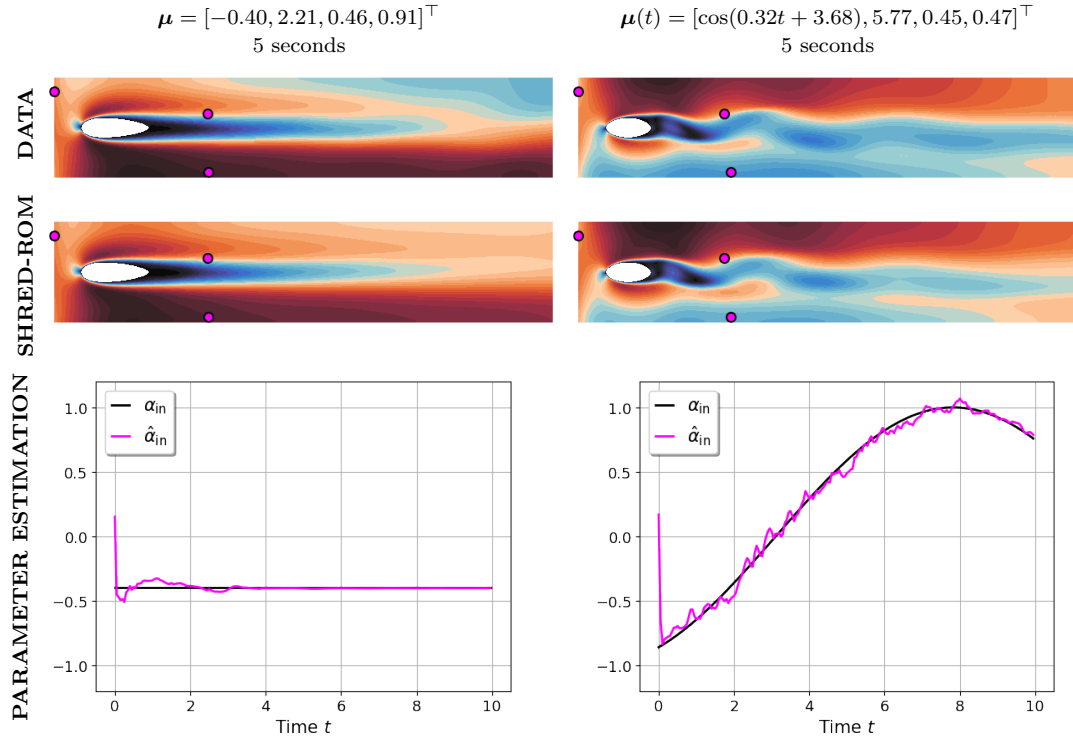


Figure 6: *Fluid Around an Obstacle*. Ground truth (first row), SHRED-ROM reconstructions (second row) and parameter estimations (third row). The following test cases are considered: state reconstruction at $t = 5$ seconds and parameter estimation over $[0, 10]$ seconds from 3 fixed horizontal velocity sensors with $\mu = [-0.40, 2.21, 0.46, 0.91]^\top$ (first column) and $\mu(t) = [\cos(0.32t + 3.68), 5.77, 0.45, 0.47]^\top$ (second column). The sparse sensors exploited by SHRED-ROM are depicted with magenta dots.

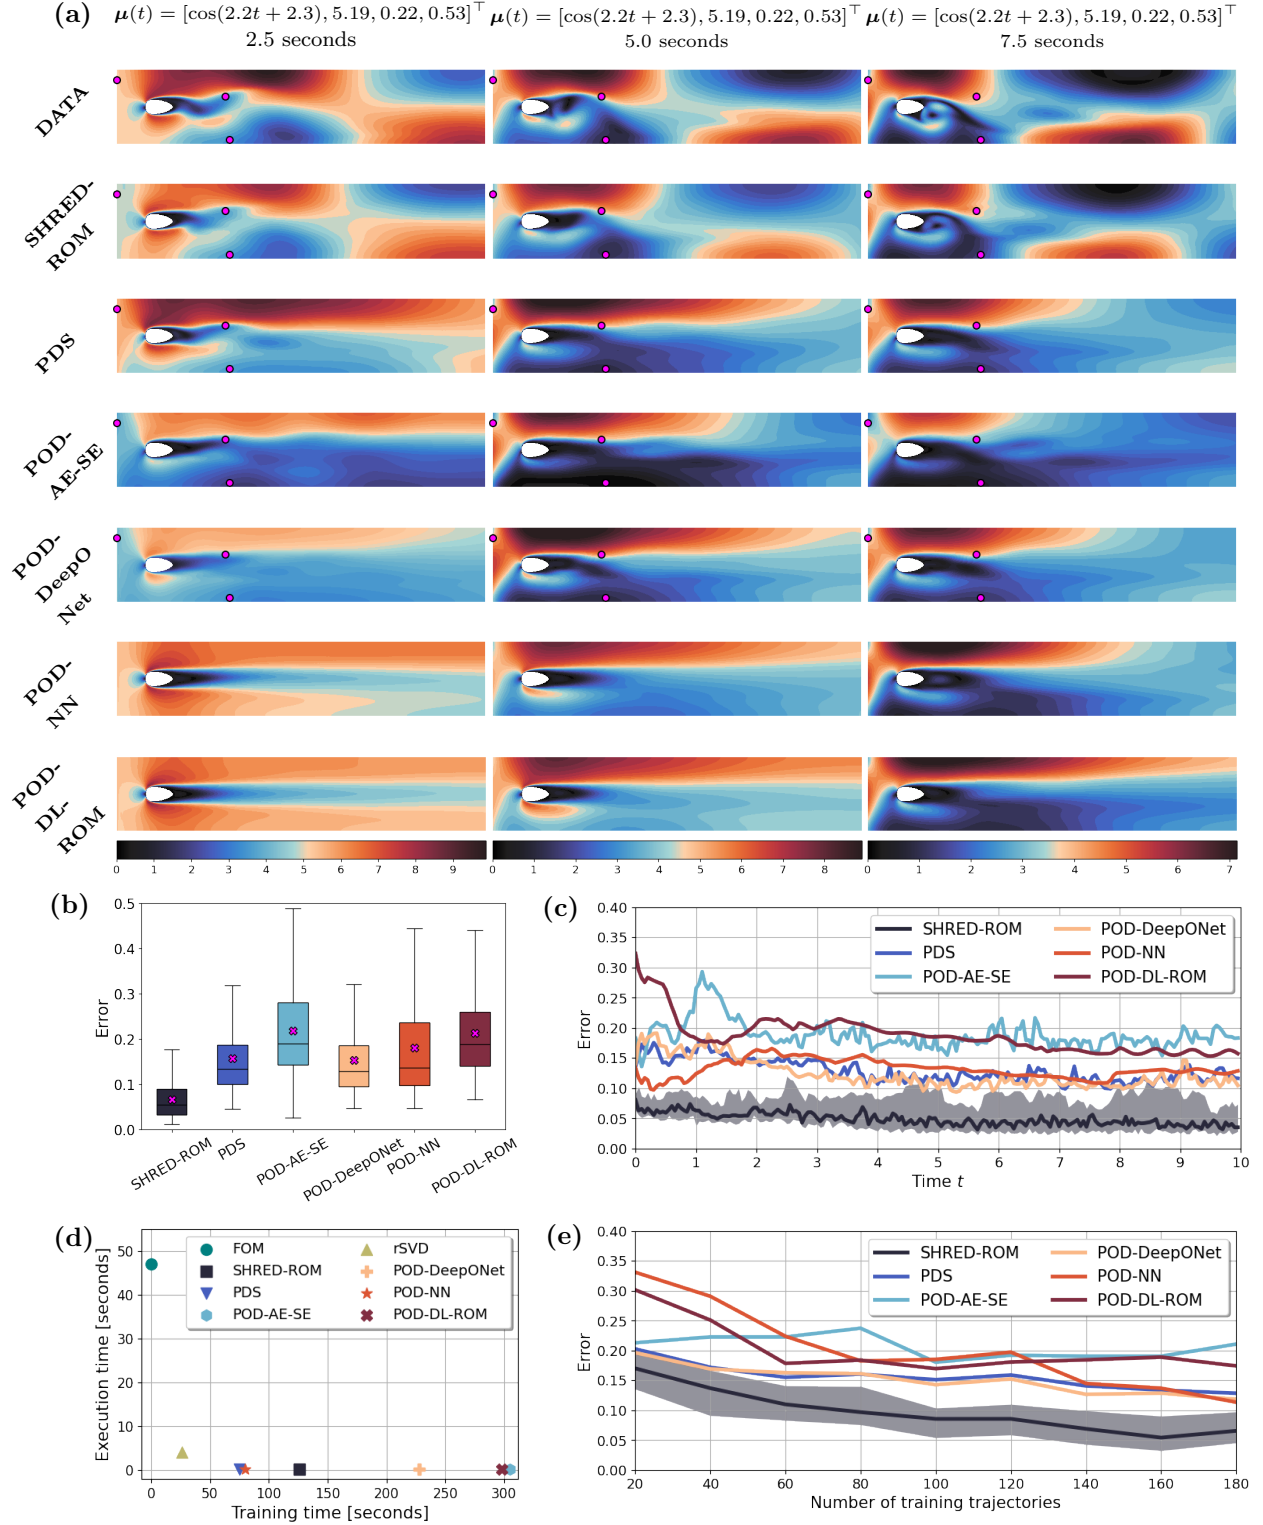


Figure 7: *Fluid Around an Obstacle*. (a) Ground truth and competing methods velocity reconstructions with test parameters $\mu(t) = [\cos(2.2t + 2.3), 5.19, 0.22, 0.53]^\top$ at $t = 2.5$ (first column), $t = 5.0$ (second column) and $t = 7.5$ seconds (third column). (b) Distributions of relative reconstruction errors committed on test data by competing methods. (c) Median relative reconstruction errors committed on test data by competing methods over the time horizon $[0, T]$. For SHRED-ROM, also the corresponding interquartile range is displayed. (d) Training and execution times required by the *full-order model* (FOM), randomized SVD (rSVD) and competing methods. (e) Median relative reconstruction errors committed on test data by competing methods for different training set dimensions. For SHRED-ROM, also the corresponding interquartile range is displayed.

in terms of the velocity $\mathbf{v} : \Omega(\gamma_l, \gamma_r) \times [0, T] \rightarrow \mathbb{R}^2$ and pressure $p : \Omega(\gamma_l, \gamma_r) \times [0, T] \rightarrow \mathbb{R}$, where $\Omega(\gamma_l, \gamma_r)$ stands for the spatial parametric domain. Starting from a zero velocity at time $t = 0$, the fluid enters the domain from the left boundary with angle of attack α_{in} and intensity γ_{in} , that is

$$\mathbf{v}((0, x_2), t) = (\gamma_{\text{in}} \cos(\alpha_{\text{in}}), x_2(2 - x_2)\gamma_{\text{in}} \sin(\alpha_{\text{in}})) \quad (15)$$

where the parabolic profile $x_2(2 - x_2)$ is useful to prevent discontinuities. Moreover, free-slip and no-slip boundary conditions are employed to, respectively, the walls and the obstacle.

Our goal is to reconstruct the high-dimensional velocity field over time for different sets of parameters $\boldsymbol{\mu} = [\alpha_{\text{in}}, \gamma_{\text{in}}, \gamma_l, \gamma_r]^\top$. To this aim, we generate 200 velocity trajectories by solving (15) through the incremental Chorin-Temam projection method in **fenics** [1] with $T = 10$ and $\Delta t = 0.05$ ($N_t = 200$ time steps). The first 100 scenarios consider constant parameters randomly sampled in the space $\mathcal{P} = [-1.0, 1.0] \times [1.0, 10.0] \times [0.2, 0.6] \times [0.2, 1.0]$. The other 100 velocity trajectories take into account as many time-dependent angles of attack $\alpha_{\text{in}}(t) = \cos\left(\frac{2\pi}{\tau}t + \vartheta\right)$ with τ and ϑ randomly sampled in $(2.5, 20.0)$ and $(0.0, 2\pi)$, respectively. The dimensionality of the simulated snapshots ($N_h = 80592$) is then reduced to $r = 150$ through POD. Despite a remarkably high compression ratio (99,81%), the latent representations are enough to recover the high-dimensional velocities up to a test relative error equal to 1.16%.

The spatio-temporal velocity behavior in multiple scenarios can be approximated by SHRED-ROM starting from 3 fixed sensors monitoring the horizontal velocity only. Besides the POD coefficients of the velocity, we consider the (possibly time-dependent) angle of attack as an additional output of the proposed model, allowing for a real-time parameter estimation through SHRED-ROM. After training the neural networks with a lag equal to $L = 50$, SHRED-ROM is capable of reconstructing the high-dimensional velocities in new scenarios with a mean relative error equal to 6.65%, and to estimate the corresponding angle of attack with a mean absolute error equal to 0.037 radians. Figure 6 shows the results obtained in two test cases: by a visual inspection it is possible to assess the high accuracy of SHRED-ROM for both the state reconstructions and the parameter estimations. Note that the estimate of the angle of attack $\hat{\alpha}_{\text{in}}$ is not accurate in the first few time steps. However, a very limited temporal history of sensor values allows the estimation to align with respect to the ground truth α_{in} .

With a different training-validation-test splitting, we can estimate state and angle of attack for future times with respect to the ones considered during training. In particular, we now exploit 160 trajectories in the time range $[0, 5]$ for randomized SVD and SHRED-ROM training, while the remaining data (that are the same 160 trajectories for $5 < t \leq 10$ and 40 full trajectories over $[0, 10]$) are used for evaluation purposes only. SHRED-ROM is now able to reconstruct the high-dimensional velocity both in new scenarios and future time instants with a mean relative error equal to 7.96%. Similar results are obtained for the angle of attack estimation, with a mean absolute error equal to 0.045 radians.

The SHRED-ROM performances are compared to 5 different state-of-the-art frameworks arising from different fields. In particular, we take into account (i) sensor-based state estimation techniques such as *POD-based deep state estimation* (PDS) [45] and *POD-enhanced autoencoder state estimation* (POD-AE-SE) [43], (ii) the POD-DeepONet [41] operator learning strategy, and (iii) non-intrusive reduced order modeling techniques such as POD-NN [26] and *POD-enhanced deep learning-based reduced order model* (POD-DL-ROM) [22]. Note that, differently from SHRED-ROM, the first three frameworks rely on a one-shot state reconstruction from sensor data, while the last two learn the parameters-to-state map. To perform fair comparisons, every setting takes into account POD-based compressive training with $r = 150$ modes, 3 fixed sensors measuring the horizontal velocity and neural networks with comparable numbers of parameters with respect to SHRED-ROM. For instance, the mappings from sensors or parameters to reduced states show similar complexities with respect to the proposed LSTM, while the encoders and decoders exploited by POD-AE-SE and POD-DL-ROM share the same architecture of the proposed SDN, with a hidden dimension equal to 11 as suggested by Franco et al. [19]. Panel (a) of Figure 7 displays the velocity fields predicted by the aforementioned strategies for $\boldsymbol{\mu}(t) = [\cos(2.2t + 2.3), 5.19, 0.22, 0.53]^\top$ at three different time instants. The temporal history of sensor values encoded through the LSTM in SHRED-ROM is crucial to achieve more accurate reconstructions with respect to the competing methods. The superiority of SHRED-ROM is also confirmed by the distributions of the relative test errors shown in panels (b) and (c) of the same figure. Panel (d) compares, instead, the training and execution times of the different frameworks, along with the computational times of the *full-order model* (FOM) and randomized SVD (rSVD). While the evaluation times of all the competing methods are extremely fast (ranging between 0.15 and 0.3 seconds), with a

remarkably high speed up with respect to the FOM, SHRED-ROM is faster to train than POD-DeepONet and autoencoder-based techniques, such as POD-AE-SE and POD-DL-ROM. Note that the training time differences would be even higher whenever non-compressive training strategies were considered. Finally, panel (e) of Figure 7 presents a data requirement analysis across all the competing models. Even with few training trajectories, SHRED-ROM is able to achieve impressive generalization capabilities, superior to the other competitors investigated.

4. CONCLUSIONS

In this work, we advocate a new decoding-only reduced order modeling framework to reconstruct high-dimensional fields from temporal history of sparse sensor measurements in multiple scenarios. Importantly, state snapshots reduction through, e.g., POD or spherical harmonics expansion, allows for lightweight neural networks and efficient compressive training strategies, feasible with laptop-level computing. Moreover, differently from other deep learning-based models, minimal hyperparameter tuning is required, as demonstrated throughout the test cases detailed in Section 3.

With several applications dealing with chaotic and nonlinear fluid dynamics, as well as video data, we show that the temporal history of at most 3 sensors is enough to achieve accurate reconstructions for new scenario parameters, with a remarkably high speed up with respect to numerical simulations. Moreover, we demonstrate that SHRED-ROM serves as a versatile framework capable of dealing with fixed and mobile sensors, synthetic and video data, noisy measurements, coupled fields, time extrapolation in periodic regimes, physical and geometric (possibly time-dependent) parametric dependencies, while being agnostic to sensor placement and parameter values. Indeed, differently from traditional ROMs strategies, the parameter values are not required both at training and evaluation stage, allowing for unknown or uncertain parametric dependencies. As shown in Section 3.5, SHRED-ROM outperforms competing sensing methods focusing on one-shot reconstructions, as well as non-intrusive ROMs, both in terms of accuracy, training time and data requirement.

The proposed SHRED-ROM framework may be extended in future works in multiple directions. For instance, forecasting may be easily taken into account, as proposed by Williams et al. [55], in order to reconstruct the state variables in multiple scenarios even in the absence of sensor values. To this aim, different deep learning-based recurrence strategies may be considered to embed temporal sensor measurements, with accurate forecasts in autoregressive or free-running mode. Moreover, interpretable and sparse dynamical models can be identified at latent level through a SINDy-based regularization, as proposed by Gao et al. [24] in nonparametric settings, promoting smooth predictions and stable forecasts. Finally, a variational SHRED-ROM may be taken into account to quantify uncertainties in the model reconstructions and in the parameter estimates.

DATA AND CODE AVAILABILITY

The data and the code can be found in our repositories
<https://doi.org/10.5281/zenodo.14524524>
<https://github.com/MatteoTomasetto/SHRED-ROM>.

ACKNOWLEDGEMENTS

The work of JNK was supported in part by the US National Science Foundation (NSF) AI Institute for Dynamical Systems (dynamicsai.org), grant 2112085. AM acknowledges the Project “Reduced Order Modeling and Deep Learning for the real-time approximation of PDEs (DREAM)” (Starting Grant No. FIS00003154), funded by the Italian Science Fund (FIS) - Ministero dell’Università e della Ricerca and the project FAIR (Future Artificial Intelligence Research), funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence).

References

- [1] M. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. Rognes, and G. Wells. The fenics project version 1.5. *Archive of Numerical Software*, 3(100):9–23, 2015.
- [2] F. Andreuzzi, N. Demo, and G. Rozza. A dynamic mode decomposition extension for the forecasting of parametric dynamical systems. *SIAM Journal on Applied Dynamical Systems*, 22(3):2432–2458, 2023.
- [3] A. C. Antoulas. *Approximation of large-scale dynamical systems*. SIAM, 2005.
- [4] P. Benner, S. Gugercin, and K. Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM review*, 57(4):483–531, 2015.
- [5] P. Boersma, E. DeWitt, C. Quinteros, F. Thurber, T. Gurian, P. Riahi, and Y. Modarres-Sadeghi. Video: Vortex arms - symmetric vortices in the wake of an oscillating cylinder. In *74th Annual Meeting of the APS Division of Fluid Dynamics - Gallery of Fluid Motion*. American Physical Society, 2021.
- [6] S. L. Brunton, J. N. Kutz, K. Manohar, A. Y. Aravkin, K. Morgansen, J. Klemisch, N. Goebel, J. Buttrick, J. Poskin, A. Blom-Schieber, et al. Data-driven aerospace engineering: Reframing the industry with machine learning. *arXiv preprint arXiv:2008.10740*, 2020.
- [7] 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.
- [8] K. Carlberg, M. Barone, and H. Antil. Galerkin v. least-squares petrov–galerkin projection in nonlinear model reduction. *Journal of Computational Physics*, 330:693–734, 2017.
- [9] 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.
- [10] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [11] P. Conti, G. Gobat, S. Fresca, A. Manzoni, and A. Frangi. Reduced order modeling of parametrized systems through autoencoders and sindy approach: continuation of periodic solutions. *Computer Methods in Applied Mechanics and Engineering*, 411:116072, 2023.
- [12] P. Conti, J. Kneifl, A. Manzoni, A. Frangi, J. Fehr, S. L. Brunton, and J. N. Kutz. Veni, vindy, vici: a variational reduced-order modeling framework with uncertainty quantification. *arXiv preprint arXiv:2405.20905*, 2024.
- [13] R. Courant and D. Hilbert. *Methods of mathematical physics: partial differential equations*. John Wiley & Sons, 2008.
- [14] J. J. D. CROZ and N. J. Higham. Stability of methods for matrix inversion. *IMA Journal of Numerical Analysis*, 12(1):1–19, 1992.
- [15] M. R. Ebers, J. P. Williams, K. M. Steele, and J. Nathan Kutz. Leveraging arbitrary mobile sensor trajectories with shallow recurrent decoder networks for full-state reconstruction. *IEEE Access*, 12:97428–97439, 2024.
- [16] N. Farenga, S. Fresca, S. Brivio, and A. Manzoni. On latent dynamics learning in nonlinear reduced order modeling. *Neural Networks*, page 107146, 2025.
- [17] G. B. Folland. *Introduction to partial differential equations*. Princeton University Press, Princeton, NJ, 1995.
- [18] G. E. Forsythe, M. A. Malcolm, and C. B. Moler. Computer methods for mathematical computation prentice-hall. *Englewood Cliffs, New Jersey*, 1977.

- [19] N. R. Franco, A. Manzoni, and P. Zunino. A Deep Learning approach to Reduced Order Modeling of parameter dependent Partial Differential Equations. *Mathematics of Computation*, 92(340):483–524, 2023.
- [20] 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.
- [21] S. Fresca, F. Fatone, A. Manzoni, et al. Long-time prediction of nonlinear parametrized dynamical systems by deep learning-based reduced order models. *Mathematics in Engineering*, 5(6):1–36, 2023.
- [22] 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.
- [23] H. Gao, S. Kaltenbach, and P. Koumoutsakos. Generative learning for forecasting the dynamics of high-dimensional complex systems. *Nature Communications*, 15(1):8904, 2024.
- [24] M. L. Gao, J. P. Williams, and J. N. Kutz. Sparse identification of nonlinear dynamics and koopman operators with shallow recurrent decoder networks. *arXiv preprint arXiv:2501.13329*, 2025.
- [25] R. Haberman. *Elementary applied partial differential equations*, volume 987. Prentice Hall Englewood Cliffs, NJ, 1983.
- [26] J. Hesthaven and S. Ubbiali. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*, 363:55–78, 2018.
- [27] J. S. Hesthaven, G. Rozza, B. Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.
- [28] N. J. Higham. *Accuracy and stability of numerical algorithms*. SIAM, 2002.
- [29] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [30] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge university press, 2012.
- [31] Q. A. Huhn, M. E. Tano, J. C. Ragusa, and Y. Choi. Parametric dynamic mode decomposition for reduced order modeling. *Journal of Computational Physics*, 475:111852, 2023.
- [32] A.-K. Kassam and L. N. Trefethen. Fourth-order time-stepping for stiff pdes. *SIAM Journal on Scientific Computing*, 26(4):1214–1233, 2005.
- [33] J. P. Keener. *Principles of applied mathematics: transformation and approximation*. CRC Press, 2018.
- [34] D. P. Kingma and M. Welling. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.
- [35] J. N. Kutz. *Data-driven modeling & scientific computation: methods for complex systems & big data*. Oxford University Press, 2013.
- [36] J. N. Kutz. Advanced differential equations: Asymptotics & perturbations. *arXiv preprint arXiv:2012.14591*, 2020.
- [37] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.
- [38] J. N. Kutz, M. Reza, F. Faraji, and A. Knoll. Shallow recurrent decoder for reduced order modeling of plasma dynamics. *arXiv preprint arXiv:2405.11995*, 2024.
- [39] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.

- [40] Y. Liu, J. N. Kutz, and S. L. Brunton. Hierarchical deep learning of multiscale differential equation time-steppers. *arXiv preprint arXiv:2008.09768*, 2020.
- [41] L. Lu, X. Meng, S. Cai, Z. Mao, S. Goswami, Z. Zhang, and G. E. Karniadakis. A comprehensive and fair comparison of two neural operators (with practical extensions) based on fair data. *Computer Methods in Applied Mechanics and Engineering*, 393:114778, 2022.
- [42] M. Lukoševičius. *A Practical Guide to Applying Echo State Networks*, pages 659–686. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- [43] Z. Luo, L. Wang, J. Xu, M. Chen, J. Yuan, and A. C. C. Tan. Flow reconstruction from sparse sensors based on reduced-order autoencoder state estimation. *Physics of Fluids*, 35(7):075127, 2023.
- [44] R. Maulik, K. Fukami, N. Ramachandra, K. Fukagata, and K. Taira. Probabilistic neural networks for fluid flow surrogate modeling and data recovery. *Physical Review Fluids*, 5:104401, 2020.
- [45] N. J. Nair and A. Goza. Leveraging reduced-order models for state estimation using deep learning. *Journal of Fluid Mechanics*, 897:R1, 2020.
- [46] R. Ohana, M. McCabe, L. T. Meyer, R. Morel, F. J. Agocs, M. Beneitez, M. Berger, B. Burkhart, S. B. Dalziel, D. B. Fielding, D. Fortunato, J. A. Goldberg, K. Hirashima, Y.-F. Jiang, R. Kerswell, S. Maddu, J. M. Miller, P. Mukhopadhyay, S. S. Nixon, J. Shen, R. Watteaux, B. R.-S. Blancard, F. Rozet, L. H. Parker, M. Cranmer, and S. Ho. The well: a large-scale collection of diverse physics simulations for machine learning. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.
- [47] E. J. Parish and K. T. Carlberg. Time-series machine-learning error models for approximate solutions to parameterized dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 365:112990, 2020.
- [48] T. Qin, K. Wu, and D. Xiu. Data driven governing equations approximation using deep neural networks. *Journal of Computational Physics*, 395:620–635, 2019.
- [49] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced basis methods for partial differential equations: an introduction*, volume 92. Springer, 2015.
- [50] F. Regazzoni, L. Dede, and A. Quarteroni. Machine learning for fast and reliable solution of time-dependent differential equations. *Journal of Computational physics*, 397:108852, 2019.
- [51] F. Regazzoni, S. Pagani, M. Salvador, L. Dede, and A. Quarteroni. Learning the intrinsic dynamics of spatio-temporal processes through latent dynamics networks. *Nature Communications*, 15(1):1834, 2024.
- [52] S. Riva, C. Intorini, A. Cammi, and J. N. Kutz. Robust state estimation from partial out-core measurements with shallow recurrent decoder for nuclear reactors. *arXiv preprint arXiv:2409.12550*, 2024.
- [53] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2023.
- [54] P. R. Vlachas, G. Arampatzis, C. Uhler, and P. Koumoutsakos. Multiscale simulations of complex systems by learning their effective dynamics. *Nature Machine Intelligence*, 4(4):359–366, 2022.
- [55] J. P. Williams, O. Zahn, and J. N. Kutz. Sensing with shallow recurrent decoder networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 480(2298):20240054, 2024.
- [56] Q. Zhang, D. Krotov, and G. E. Karniadakis. Operator learning for reconstructing flow fields from sparse measurements: an energy transformer approach. *arXiv preprint arXiv:2501.08339*, 2025.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 53/2025** Zecchi, A. A.; Sanavio, C.; Perotto, S.; Succi, S.
Telescopic quantum simulation of the advection-diffusion-reaction dynamics
- 51/2025** Tomasetto, M.; Braghin, F.; Manzoni, A.
Latent feedback control of distributed systems in multiple scenarios through deep learning-based reduced order models
- 50/2025** Bonetti, S.; Botti, M.; Antonietti, P.F.
Conforming and discontinuous discretizations of non-isothermal Darcy–Forchheimer flows
- 49/2025** Zanin, A.; Pagani, S.; Corti, M.; Crepaldi, V.; Di Fede, G.; Antonietti, P.F.; the ADNI
Predicting Alzheimer's Disease Progression from Sparse Multimodal Data by NeuralODE Models
- 48/2025** Temellini, E.; Ballarin, F.; Chacon Rebollo, T.; Perotto, S.
On the inf-sup condition for Hierarchical Model reduction of the Stokes problem
- 47/2025** Gimenez Zapiola, A.; Consolo, A.; Amaldi, E.; Vantini, S.
Penalised Optimal Soft Trees for Functional Data
- 46/2025** Mirabella, S.; David, E.; Antona, A.; Stanghellini, C.; Ferro, N.; Matteucci, M.; Heuvelink, E.; Perotto, S.
On the Impact of Light Spectrum on Lettuce Biophysics: A Dynamic Growth Model for Vertical Farming
- 45/2025** Calì, G.; Ragazzi, F.; Popoli, A.; Cristofolini, A.; Valdetaro, L.; De Falco, C.; Barbante, F.
Hierarchical Multiscale Modeling of Positive Corona Discharges
- 44/2025** Brivio, S.; Fresca, S.; Manzoni, A.
Handling geometrical variability in nonlinear reduced order modeling through Continuous Geometry-Aware DL-ROM
- 43/2025** Tomasetto, M.; Manzoni, A.; Braghin, F.
Real-time optimal control of high-dimensional parametrized systems by deep-learning based reduced order models