

Neural-network learning of SPOD latent dynamics

Andrea Lario^a, Romit Maulik^b, Oliver T. Schmidt^c, Gianluigi Rozza^a,
Gianmarco Mengaldo^{d,*}

^a Scuola Internazionale Superiore di Studi Avanzati (SISSA), Italy

^b Argonne National Laboratory (ANL), USA

^c University of California San Diego (UCSD), USA

^d National University of Singapore (NUS), Singapore

ARTICLE INFO

Article history:

Received 31 January 2022

Received in revised form 8 July 2022

Accepted 11 July 2022

Available online 18 July 2022

Keywords:

Dynamical systems

Reduced order modeling

Neural networks

Deep learning

ABSTRACT

We aim to reconstruct the latent space dynamics of high dimensional, quasi-stationary systems using model order reduction via the spectral proper orthogonal decomposition (SPOD). The proposed method is based on three fundamental steps: in the first, once that the mean flow field has been subtracted from the realizations (also referred to as snapshots), we compress the data from a high-dimensional representation to a lower dimensional one by constructing the SPOD latent space; in the second, we build the time-dependent coefficients by projecting the snapshots containing the fluctuations onto the SPOD basis and we learn their evolution in time with the aid of recurrent neural networks; in the third, we reconstruct the high-dimensional data from the learnt lower-dimensional representation. The proposed method is demonstrated on two different test cases, namely, a compressible jet flow, and a geophysical problem known as the Madden-Julian Oscillation. An extensive comparison between SPOD and the equivalent POD-based counterpart is provided and differences between the two approaches are highlighted. The numerical results suggest that the proposed model is able to provide low rank predictions of complex statistically stationary data and to provide insights into the evolution of phenomena characterized by specific range of frequencies. The comparison between POD and SPOD surrogate strategies highlights the need for further work on the characterization of the interplay of error between data reduction techniques and neural network forecasts.

© 2022 Elsevier Inc. All rights reserved.

1. Introduction

Deep learning strategies are gaining traction as a tool for surrogate modeling (also referred to as emulation or non-intrusive reduced order modeling) of dynamical systems [1–10]. Data associated to dynamical systems has usually a space and a time component. For instance, one can think of the evolution in time of the temperature at various locations worldwide. The various locations where the temperature is provided constitute the space component of the data, while the time snapshots constitute the time component. A common case in computational physics is when we have access to values of a given quantity at S spatial locations (or grid points), for N_t time snapshots. In this case, the dimension of the problem is $S \times N_t$, where the spatial component S might be extremely large. The key in surrogate modeling or emulation is to iden-

* Corresponding author.

E-mail addresses: alario@sissa.it (A. Lario), rmaulik@anl.gov (R. Maulik), oschmidt@eng.ucsd.edu (O.T. Schmidt), gianluigi.rozza@sissa.it (G. Rozza), mpegim@nus.edu.sg (G. Mengaldo).

<https://doi.org/10.1016/j.jcp.2022.111475>

0021-9991/© 2022 Elsevier Inc. All rights reserved.

tify a suitable representation of the data that reduces the high dimensionality of S to a more computationally manageable number, that we denote with S_r . The new reduced space of dimension S_r , often referred to as the *latent space*, is where we seek to learn the time evolution of the system. The learnt reduced space can then be used to reconstruct the original high-dimensional space by inverting the data compression. Indeed, the three key steps that constitute the foundations of surrogate modeling or emulation are: (i) data compression from high- to lower-dimensional representation, (ii) learning the time evolution in lower-dimensional space, (iii) reconstruction of the high-dimensional data from lower-dimensional representation learnt. Given the nature of surrogate modeling, there are two errors that one needs to take into account, the (a) data compression/projection error, and the (b) learning error. The trade-off is usually to find a data compression efficient enough to make the problem computationally feasible while retaining sufficient accuracy. This balance is obviously task dependent, and may vary across disciplines and specific applications. In addition, the interplay of data compression and time-evolution learning must be accounted for carefully, as the learning algorithms adopted might not be optimal for the specific latent space built. Because of this delicate trade-off and complex interplay, it is not surprising that several learning approaches to emulate latent-space dynamics have been recently proposed in the literature. These include approaches where proper orthogonal decomposition (POD) and autoencoders have been used to compress the data, while recurrent and convolutional neural networks, among others, have been used for the time-evolution learning task [11–15]. Many works are currently exploring non-linear techniques for dimensional reductions relying, for example, either on unsupervised learning approaches such as generative adversarial network (GAN) [16] or on diffusion maps [17].

In computational physics, especially in the context of engineering and geophysical flows, one of the most popular methods for reducing the dimensionality of the problem is the POD [18,19]. This technique seeks to find a latent space of dimension S_r that is composed of a basis vector $\phi_j(\mathbf{x})$ that represents the spatial component of the data, and time coefficients $a_j(t)$ that represent the time component

$$\mathbf{q}(\mathbf{x}, t) = \sum_{j=1}^{S_r} a_j(t) \phi_j(\mathbf{x}) \quad t = 1, \dots, N \quad (1)$$

where $t = 1, \dots, N$ are the time snapshots available, and the coefficient constitute the latent space. One typically computes the basis vector and the coefficients on a portion of the snapshots available, e.g., $t = 1, \dots, N_{\text{train}}$, where $N_{\text{train}} < N$. This portion of the data is usually referred to as training data, as it is commonly used to train a data-driven method – e.g., a neural network – to reconstruct the time evolution of the latent space (i.e., the coefficients). Once the data-driven method is trained and the evolution of the coefficients learnt, one reconstructs the original high-dimensional data using the pre-computed basis vector ϕ . The learnt dynamics of the latent space can then be used as a non-intrusive reduced order model on e.g., unseen coefficients.

In this paper, we construct a surrogate model for the spectral proper orthogonal decomposition (SPOD) latent space via recurrent neural networks. In particular, we use the SPOD to reduce the dimensionality of the problem, thereby identifying the SPOD latent space through singular value decomposition and oblique projection of the modes, and use recurrent neural networks to learn the time evolution of the SPOD latent space. The SPOD-based surrogate model is then compared against more popular POD-based surrogate strategies. The key difference between the two dimensionality reduction methods, POD and SPOD, consists of the space where the basis vector is sought for. On the one hand, POD extracts the basis vector (also referred to as modes) as the eigenvectors of the (spatially weighted) anomaly covariance matrix. Therefore, POD provides a space-only picture of the problem. On the other hand, SPOD [20–22] is based on an eigenvalue decomposition of the estimated cross-spectral density (CSD) matrix of the data, that is we can leverage the frequency space to extract coherent patterns of interest at specific wavenumbers. Indeed, at each frequency, SPOD yields a set of optimally ranked, time-harmonic and orthogonal modes, that are coherent both in space and time. Furthermore, these modes are identical to optimal response modes of a stochastically forced non-normal linear dynamical system if the forcing consists of white noise. Hence, they provide a direct bridge for dynamical system modeling in computational physics, where the generative process of the underlying data comes from a set of differential equations. This useful property, could be used, for instance, in the context of geophysical flows, where it can open several opportunities in the analysis and physical understanding of Earth system model results [23–25]. In particular, we can construct different SPOD latent spaces depending on the set of wavenumbers that we are interested in. For instance, in geophysical flows, it is common to encounter phenomena that are characterized by a compact frequency band and whose evolution in time might be of interest. This is the case of several modes of variability, such as the Madden-Julian Oscillation, or El Niño–Southern Oscillation. In these cases, one might be interested in a detailed analysis of the frequency spectrum, and the evolution of the latent space associated to only a portion of the frequencies provided by the SPOD, and/or to perform data prewhitening and de-noising [26–28].

In this paper, we show how SPOD surrogate modeling with a frequency-selective latent space learning can be achieved, and we compare it against more standard POD data reduction methods. We also show how the POD can be reconstructed from the SPOD modes, and discuss, for a neural network architecture, the trade-off between data reduction error and learning error. In Fig. 1, we depict the conceptual workflow of the proposed approach for learning the latent space of SPOD, eventually selecting a frequency band of the spectrum provided by the method. The left-hand side plot of Fig. 1 consists of high-dimensional data on which we apply the SPOD. From this, we obtain a set of space-time coherent SPOD structures at different periods (or frequencies), that constitute that basis to build the latent (or reduced) space that the




Fig. 1. Conceptual workflow of SPOD latent space learning.

neural network learns. We use the model that the neural network learnt to forecast future states, and we finally reconstruct high-dimensional future snapshots from the reduced space forecasts.

The rest of the paper is organized as follows. In section 2, we briefly introduce the SPOD method, and detail the neural-network-based SPOD emulation strategy. In section 3, we compare POD and SPOD emulation, pointing out their key differences. In section 4, we present the results on a compressible jet and on a geophysical problem associated to the Madden-Julian Oscillation. We also present a direct comparison between POD and SPOD emulation. Finally, in section 5, we provide a critical discussion of this work, and draw the key conclusions.

2. SPOD emulation via non-intrusive reduced order modeling

This section follows the workflow presented in Fig. 1. In particular, we introduce the SPOD method (section 2.1) that allows us to reduce the dimensionality of the data and construct the SPOD latent space (section 2.2). These are depicted as step 1 and step 2 in Fig. 1. We then introduce the neural-network learning strategy (section 2.3), depicted as step 3 in Fig. 1, along with its connections to theoretical findings. Finally, we present the reconstruction of the high-dimensional space from the SPOD latent space (section 2.4), depicted as step 4 in Fig. 1.

2.1. The SPOD method

We use the SPOD approach to identify a suitable latent space of high-dimensional data, as a first step of our dynamical system emulation strategy. The high-dimensional data consists of multidimensional fluctuating and wide-sense stationary time-series $\mathbf{q}(\mathbf{x}, t_i)$, from which we subtract the temporal mean $\bar{\mathbf{q}}$. The data $\mathbf{q}(\mathbf{x}, t_i)$ is composed by $t_i = 1, \dots, N_t$ time snapshots, spatial coordinates $\mathbf{x} \in \mathbb{R}^d$ for a total of S spatial points available, and $\mathbf{q} = q_1, \dots, q_{N_{\text{var}}}$ variables.

To derive the SPOD method, we recast each time snapshot of the multidimensional data into a vector of dimension $\mathbf{q}(\mathbf{x}, t_i) = \mathbf{q}_i \in \mathbb{R}^{M \times N_t}$, where M is the number of spatial points S multiplied by the number of variables N_{var} , that is $M = S \times N_{\text{var}}$. We can then assemble the data (also referred to as snapshot) matrix

$$\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{N_t}] \in \mathbb{R}^{M \times N_t}, \quad (2)$$

that constitutes the starting point of the SPOD approach. The task is to identify a deterministic (set of) function(s) $\phi(\mathbf{x}, t)$ that best approximates the statistically stationary (zero-mean) process $\mathbf{q}(\mathbf{x}, t)$ underlying our data $\mathbf{Q}$. In mathematical terms, this translates into maximizing the quantity

$$\lambda = \frac{E\{|\langle \mathbf{q}(\mathbf{x}, t), \phi(\mathbf{x}, t) \rangle_{\mathbf{x}, t}|^2\}}{\langle \phi(\mathbf{x}, t), \phi(\mathbf{x}, t) \rangle_{\mathbf{x}, t}}, \quad (3)$$

where we assume that any realization of $\{\mathbf{q}(\mathbf{x}, t)\}$ span a Hilbert space H with inner product $\langle \cdot, \cdot \rangle_{\mathbf{x}, t}$ and expectation operator $E\{\cdot\}$, here taken to be the ensemble mean. The inner product $\langle \cdot, \cdot \rangle_{\mathbf{x}, t}$ is defined as

$$\langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{x}, t} = \int_{-\infty}^{\infty} \int_{\Omega} \mathbf{u}^*(\mathbf{x}, t) \mathbf{W} \mathbf{v}(\mathbf{x}, t) d\Omega dt, \quad (4)$$

where $d\Omega$ denotes the spatial integration, dt the time integration, and $\mathbf{W}$ is a spatial weighting matrix. Given the assumption that any realization of $\{\mathbf{q}(\mathbf{x}, t)\}$ span H , and thanks to Karhunen–Loève (KL) theorem, we know that there exists a set of mutually orthogonal deterministic functions that forms a complete basis in the Hilbert space H . This can be defined

as $\mathbf{q}(\mathbf{x}, t) = \sum_{j=1}^{\infty} a_j \phi_j(\mathbf{x}, t)$, where ϕ_k are eigenfunctions whose associated eigenvalues λ_k arise from the solution of the Fredholm integral equation, that owing to the stationarity assumption on the underlying process, can be written in the frequency domain

$$\int_S \mathcal{S}(\mathbf{x}, \mathbf{x}', f) \psi(\mathbf{x}', f) d\mathbf{x}' = \lambda(f) \phi(\mathbf{x}, f), \quad (5)$$

where $\mathcal{S}(\mathbf{x}, \mathbf{x}', f)$ is the Fourier transform of the two-point space-time correlation tensor $C(\mathbf{x}, \mathbf{x}', \tau) = E\{\mathbf{q}(\mathbf{x}, t) \mathbf{q}^*(\mathbf{x}', t')\}$, also referred to as the cross-spectral density tensor

$$\mathcal{S}(\mathbf{x}, \mathbf{x}', f) = \int_{-\infty}^{\infty} C(\mathbf{x}, \mathbf{x}', \tau) e^{-i2\pi f \tau} d\tau. \quad (6)$$

To form the discrete analog of (5), we use the Welch's method, and segment the data into L (possibly overlapping) blocks

$$\mathbf{Q}^{(\ell)} = [\mathbf{q}_1^{(\ell)}, \dots, \mathbf{q}_{N_{\text{fft}}}^{(\ell)}] \in \mathbb{R}^{M \times N_{\text{fft}}}, \quad \ell = 1, \dots, L, \quad (7)$$

where each block contains N_{fft} snapshots, overlaps by ℓ_{ovlp} snapshots with the next block, and is regarded as equally representative of the whole data (i.e. it is assumed ergodic). An optimal overlap value is 50% as suggested by Welch [29], since it avoids data leakages when windowing functions are applied on the ensembles; different strategies must be adopted in some cases, for example nonoverlapping blocks minimize the variance of spectral estimations while on the other hand heavily duplicated data (i.e. high overlap value) allow to extract a higher number of blocks from the starting set of data $\mathbf{q}(\mathbf{x}, t_i)$ for a fixed number of snapshots N_{fft} . The latter strategy was adopted for the test case described in 4.3 in order to extract from a relatively small data set a sufficient number of blocks to train a neural network. To obtain the discrete analog of the cross-spectral density tensor, we apply the temporal discrete Fourier transform to our data in equation (7)

$$\hat{\mathbf{Q}}^{(\ell)} = [\hat{\mathbf{q}}_1^{(\ell)}, \hat{\mathbf{q}}_2^{(\ell)}, \dots, \hat{\mathbf{q}}_{N_{\text{fft}}}^{(\ell)}] \in \mathbb{R}^{M \times N_{\text{fft}}}, \quad \ell = 1, \dots, L. \quad (8)$$

We then reorganize the Fourier-transformed data by collecting all realizations at the k -th frequency into a single data matrix

$$\hat{\mathbf{Q}}_k = [\hat{\mathbf{q}}_k^{(1)}, \hat{\mathbf{q}}_k^{(2)}, \dots, \hat{\mathbf{q}}_k^{(L)}] \in \mathbb{R}^{M \times L}, \quad \text{for frequency } k. \quad (9)$$

From this, we can construct the cross-spectral density matrix $\mathbf{S}_k = \mathbf{X}_k \mathbf{X}_k^* \in \mathbb{R}^{M \times M}$, where $\mathbf{X}_k = \frac{1}{\sqrt{L}} \mathbf{W}^{\frac{1}{2}} \hat{\mathbf{Q}}_k = [\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \dots, \mathbf{x}_k^{(L)}] \in \mathbb{R}^{M \times L}$. The SPOD modes and associated energies can be computed as the eigenvectors and eigenvalues of the $\mathbf{S}_k$, that is $\mathbf{S}_k = \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^*$, where $\mathbf{\Lambda}_k = \text{diag}(\lambda_k^{(1)}, \lambda_k^{(2)}, \dots, \lambda_k^{(L)}) \in \mathbb{R}^{L \times L}$ and $\mathbf{U}_k = [\mathbf{u}_k^{(1)}, \mathbf{u}_k^{(2)}, \dots, \mathbf{u}_k^{(L)}] \in \mathbb{R}^{M \times L}$ are the diagonal matrix of eigenvalues, or modal energies, and the eigenvector matrix, respectively. In practice, we solve the more computationally amenable eigenvalue problem $\mathbf{X}_k^* \mathbf{X}_k \mathbf{V}_k = \mathbf{V}_k \mathbf{\Lambda}_k$ for $\mathbf{X}_k^* \mathbf{X}_k \in \mathbb{R}^{L \times L}$ instead of $\mathbf{X}_k \mathbf{X}_k^* \in \mathbb{R}^{M \times M}$, following [30], and obtain the SPOD modes for the k -th frequency from

$$\Phi_k = \mathbf{W}^{-\frac{1}{2}} \mathbf{X}_k \mathbf{V}_k \mathbf{\Lambda}_k^{-\frac{1}{2}} = \frac{1}{\sqrt{L}} \hat{\mathbf{Q}}_k \mathbf{V}_k \mathbf{\Lambda}_k^{-\frac{1}{2}}. \quad (10)$$

We note that if we group the frequencies together, we can rewrite the SPOD modes as follows:

$$\Phi = [\phi_1, \phi_2, \dots, \phi_{N_f}] = [\underbrace{\phi_1^{(1)}, \phi_1^{(2)}, \dots, \phi_1^{(L)}}_{\phi_1}, \underbrace{\phi_2^{(1)}, \phi_2^{(2)}, \dots, \phi_2^{(L)}}_{\phi_2}, \dots, \underbrace{\phi_{N_f}^{(1)}, \phi_{N_f}^{(2)}, \dots, \phi_{N_f}^{(L)}}_{\phi_{N_f}}]. \quad (11)$$

noting that for real data the total number of frequencies is $N_f = \lceil \frac{N_{\text{fft}}}{2} \rceil + 1$. From these modes, it is possible to construct the latent space. This is discussed next. For additional details on the SPOD method, the interested reader can refer to [21,22].

2.2. SPOD latent space

As shown by Nekkanti and Schmidt [26], there are two strategies to reconstruct the SPOD latent space, namely a time-domain and a frequency-domain reconstruction. In this paper, we mainly use the time-domain reconstruction. Yet, we perform a comparison between time- and frequency-domain learning of SPOD latent space in section 4.3.

2.2.1. Time-domain reconstruction

The SPOD latent space is constituted by a matrix of expansion coefficients $\mathbf{A}$. Following [26], this can be constructed using a weighted oblique projection of the data onto the modal basis

$$\mathbf{A} = (\Phi^* \mathbf{W} \Phi)^{-1} \Phi^* \mathbf{W} \mathbf{Q} = \underbrace{[a_1^{(1)}, a_1^{(2)}, \dots, a_1^{(L)}]}_{\mathbf{a}_1}, \underbrace{[a_2^{(1)}, a_2^{(2)}, \dots, a_2^{(L)}]}_{\mathbf{a}_2}, \dots, \underbrace{[a_{N_f}^{(1)}, a_{N_f}^{(2)}, \dots, a_{N_f}^{(L)}]}_{\mathbf{a}_{N_f}}. \quad (12)$$

where $\mathbf{A} \in \mathbb{R}^{(L \times N_f) \times N_t}$ is the matrix containing the expansion coefficients and $\tilde{\Phi} \in \mathbb{R}^{M \times (L \times N_{fr})}$ is a matrix which gathers all the SPOD modes arranged by frequency as in equation (11). The full matrix of expansion coefficients, constructed using all modes and frequencies has dimensions $L \times N_f$ (frequencies). In practice, it is common to use only a portion L_r of the total number of modes, and (eventually) a portion N_{fr} of the total number of frequencies, from a frequency id flb to fub. This recasts the original high-dimensional data into a smaller SPOD latent space of dimension $L_r \times N_{fr}$, defined as follows

$$\begin{aligned} \mathbf{A}_r &= (\Phi_r^* \mathbf{W} \Phi_r)^{-1} \Phi_r^* \mathbf{W} \mathbf{Q} = \\ &= [\underbrace{a_{flb}^{(1)}, a_{flb}^{(2)}, \dots, a_{flb}^{(L_r)}}_{\mathbf{a}_{flb}}, \underbrace{a_{flb+1}^{(1)}, a_{flb+1}^{(2)}, \dots, a_{flb+1}^{(L_r)}}_{\mathbf{a}_{flb+1}}, \dots, \underbrace{a_{fub}^{(1)}, a_{fub}^{(2)}, \dots, a_{fub}^{(L_r)}}_{\mathbf{a}_{fub}}]. \end{aligned} \quad (13)$$

The calculation of the matrix coefficients (12) can be problematic, because $\tilde{\Phi}^* \mathbf{W} \tilde{\Phi}$ is potentially ill conditioned especially when a high number of modes are taken into account. Computing its inverse numerically can therefore be difficult. In this work, we follow Nekkanti and Schmidt [26], and use a pseudo-inverse based on the singular value decomposition of $\tilde{\Phi}^* \mathbf{W} \tilde{\Phi}$. Additionally, we note that one cannot use an orthogonal projection because the modes at different frequencies are not orthogonal when a space inner product is considered.

Once the reduced SPOD latent space $\mathbf{A}_r \in \mathbb{R}^{(L_r \times N_{fr}) \times N_t}$ is available, we can learn its time evolution using data-driven methods.

2.2.2. Frequency-domain reconstruction

An alternative approach for the computation of the expansion coefficients, usually referred as frequency-domain reconstruction, relies on the orthogonal projection of the Fourier-transformed data $\hat{\mathbf{Q}}$ onto the SPOD modes:

$$\hat{\mathbf{A}} = \Phi^* \mathbf{W} \hat{\mathbf{Q}} \quad (14)$$

where $\hat{\mathbf{A}} \in \mathbb{R}^{(L \times N_f) \times L}$ is the matrix containing the coefficients [26]. Hence the snapshots which belong to the l -th block can be reconstructed by computing the corresponding Fourier transformed data starting from the modes and the expansion coefficients:

$$\hat{\mathbf{Q}}^{(l)} = \left[\left(\sum_{i=1}^L \hat{a}_{ikl} \phi_i \right)_{k=1}, \left(\sum_{i=1}^L \hat{a}_{ikl} \phi_i \right)_{k=2}, \dots, \left(\sum_{i=1}^L \hat{a}_{ikl} \phi_i \right)_{k=N_f} \right], \quad (15)$$

and then by applying an Inverse Fast Fourier Transform to the vector just computed. In order to perform properly this operation, the number of frequencies N_f must be equal to the number of the snapshots contained in each block; if only some subsets of frequencies have to be retained during the snapshot reconstruction phase, one has to put equal to zero the expansion coefficients corresponding to the frequencies to be cut off.

In this paper, we propose to learn the latent space evolution via neural networks, and a long short-term memory neural network in particular. This is described next.

2.3. Learning of dynamics via neural networks

We use deep learning strategies based on neural networks to learn the time evolution of the latent space constituted by the expansion coefficients constructed in equation (13). The problem here is in the realm of autoregressive time-series forecasting, where one seek to learn the future values of a variable¹ $\mathbf{f} = [f(t+1), f(t+2), \dots, f(t+n_{TF})]$, provided past values of the same variable $\mathbf{f} = [f(t-1), f(t-2), \dots, f(t-n_{TP})]$, where n_{TF} can be different from n_{TP} . In our case, we seek to learn the time evolution of $\mathbf{A}_r$, that is: we want to forecast the future values $\mathbf{A}_r = [\mathbf{a}_r(t+1), \mathbf{a}_r(t+2), \dots, \mathbf{a}_r(t+n_{TF})]$, provided past values of these coefficients $\mathbf{A}_r = [\mathbf{a}_r(t-1), \mathbf{a}_r(t-2), \dots, \mathbf{a}_r(t-n_{TP})]$. To perform autoregressive time-series forecasting, using deep learning, several neural network architectures have been proposed, including recurrent neural networks, such as long short-term memory networks (LSTMs) [31] and gated recurrent networks [32], neural ordinary differential equations [33], transformer networks [34], and temporal convolutional networks [35]. Since an exhaustive comparison of these architectures is beyond the scope of this article - we shall utilize the LSTM as an exemplar architecture for learning

¹ In the more general case, autoregressive time-series forecasting is performed on a set of variables vs. just one variable.

the dynamics of our observed data on the space spanned by the SPOD modes. Moreover, viable theoretical connections are relatively simple to establish for reduced-order modeling and LSTM networks.

The LSTM network is a special type of the recurrent type of neural networks. The term ‘recurrent’ alludes to the fact that certain neuron outputs from a neural network, at each time step in the prediction, are utilized to influence predictions in future time steps. Embedding this type of inductive bias allows for the neural network to learn long and short-term correlations between instances in time from the training data set. The basic equations of the LSTM in our context for the coefficients $\mathbf{A}_r$ are given by

$$\begin{aligned} \text{input gate: } \mathbf{G}_i &= \boldsymbol{\varphi}_S \circ \mathcal{F}_i^{N_c}(\mathbf{A}_r), \\ \text{forget gate: } \mathbf{G}_f &= \boldsymbol{\varphi}_S \circ \mathcal{F}_f^{N_c}(\mathbf{A}_r), \\ \text{output gate: } \mathbf{G}_o &= \boldsymbol{\varphi}_S \circ \mathcal{F}_o^{N_c}(\mathbf{A}_r), \\ \text{internal state: } \mathbf{s}_t &= \mathbf{G}_f \odot \mathbf{s}_{t-1} + \mathbf{G}_i \odot (\boldsymbol{\varphi}_T \circ \mathcal{F}_a^{N_c}(\mathbf{A})), \\ \text{output: } \mathbf{h}_t &= \mathbf{G}_o \circ \boldsymbol{\varphi}_T(\mathbf{s}_t), \end{aligned} \quad (16)$$

where $\mathbf{A}_r \in \mathbb{R}^{(L_r \times N_{fr}) \times N_t}$. Here, $\boldsymbol{\varphi}_S$ and $\boldsymbol{\varphi}_T$ refer to tangent-sigmoid and tangent-hyperbolic activation functions respectively (although these may be modified with a view to improving the accuracy or robustness of learning), N_c is the number of hidden layer units in the specific neural network and is another modifiable parameter. Finally, $\mathcal{F}^n$ refers to an affine transformation and $(\odot)$ refers to a Hadamard product of two vectors. The LSTM implementation is used to advance $\mathbf{A}$ as a function of time. The interested reader may also find connections between classical implicit time-integration methods and the construction of such time-delay embedded neural architectures [36].

The most obvious connection between dynamical systems and LSTM neural networks is through the framework of time-delay embeddings for simulation of the former. A convenient starting point is the Mori-Zwanzig formalism [37–39] which divides the variables of a dynamical system into resolved and unresolved components. Subsequently, unresolved components are ‘estimated’ by functions that are parameterized by the resolved degrees of freedom leading to reduced system of equations for capturing the dynamics of the system. This leads to a generalized Langevin equation for the dynamical system or in other words an equation of the resolved variables with memory and noise.

The following formulation, derived from the first step of the Mori-Zwanzig treatment, can be used to approximate the dynamics of our observed data on the reduced space

$$\frac{d\mathbf{A}_r}{dt} = e^{\mathcal{W}t} \mathcal{W} \mathbf{A}_r, \quad (17)$$

where $\mathcal{W}$ is the unknown evolutionary operator of the dynamics in this space. Subsequently, we define two self-adjoint projection operators into orthogonal subspaces given by

$$P\mathbf{A}_h = \frac{(\mathbf{A}_r^0, \mathbf{A}_h^0)}{(\mathbf{A}_r^0, \mathbf{A}_r^0)} \mathbf{A}_r, \quad Q = I - P, \quad (18)$$

with $QP = 0$ and $\mathbf{A}^0 = \mathbf{A}(t=0)$. Therefore, P may be assumed to be a projection of our full-order representation in the space spanned by SPOD modes. Subsequently, we can further expand our system as

$$\frac{d\mathbf{A}_r}{dt} = e^{\mathcal{W}t} (Q + P) \mathcal{W} \mathbf{A}_r, \quad (19)$$

which may further be decoupled to a Markov-like projection operator $\mathcal{M}$ given by

$$e^{\mathcal{W}t} P \mathcal{W} \mathbf{A}_r = \frac{(\mathbf{A}_r^0, \mathbf{A}_h^0)}{(\mathbf{A}_r^0, \mathbf{A}_r^0)} e^{\mathcal{W}t} \mathbf{A}_r = \mathcal{M} \mathbf{A}_r \quad (20)$$

and a memory operator $\mathcal{G}$ given by

$$e^{\mathcal{W}t} Q \mathcal{W} \mathbf{A}_r^0 = e^{Q \mathcal{W}t} Q \mathcal{W} \mathbf{A}_r^0 + \int_0^t e^{\mathcal{W}(t-t_1)} P \mathcal{W} e^{Q \mathcal{W}t_1} Q \mathcal{W} \mathbf{A}_r^0 dt_1 = \mathcal{G} \mathbf{A}_r, \quad (21)$$

from Dyson’s formula [40] with t_1 corresponding to the length of memory retention. The second relationship, given by Equation (21), may be assumed to be a combination of memory effects and noise. The final evolution then becomes,

$$\frac{d\mathbf{A}_r}{dt} = \mathcal{G} \mathbf{A}_r + \mathcal{M} \mathbf{A}_r. \quad (22)$$

This expression can be compared to that of the internal state update within an LSTM (Equation (16)), which we repeat below

$$\mathbf{s}_t = \mathbf{G}_f \odot \mathbf{s}_{t-1} + \mathbf{G}_i \odot \left(\varphi_T \circ \mathcal{F}_{\mathbf{A}_r}^{N_c}(\mathbf{A}_r) \right), \quad (23)$$

where a linear combination of a nonlinearly transformed input vector at time t with the gated result of a hidden state at a previous time $t - 1$ is used to calculate the result vector at the current time. The process of carrying a state through time via gating may be assumed to be a representation of the memory integral, whereas the utilization of the current input may be assumed to be the Markovian component of the map.

2.4. SPOD high-dimensional space reconstruction

After learning an LSTM-based surrogate model for the dynamics of the latent space, we can reconstruct predictions from this model, $\tilde{\mathbf{A}}_r$, by reconstructing the predicted coefficients using the SPOD modes Φ_r , i.e.,

$$\tilde{\mathbf{Q}} = \Phi_r \tilde{\mathbf{A}}_r, \quad (24)$$

where one seeks $\tilde{\mathbf{Q}} \approx \mathbf{Q}$ and $\tilde{\mathbf{A}}_r \approx \mathbf{A}_r$. Indeed, from this procedure there are two types of errors that we need to consider: (1) the projection error, and the (2) learning error. The first can be easily computed by comparing the reconstruction of the training data coefficients vs. the true high-dimensional data prior to learning the evolution of the latent space via the neural network, i.e.,

$$\text{err}_{\text{projection}} = \tilde{\Phi}_r \mathbf{A}_r - \mathbf{Q}. \quad (25)$$

The second can be computed using the true coefficient vs. the predictions of the LSTM neural-network, i.e.,

$$\text{err}_{\text{learning}} = \tilde{\Phi}_r \tilde{\mathbf{A}}_r - \tilde{\Phi}_r \mathbf{A}_r. \quad (26)$$

We note that if we use only a portion of the total number of frequencies, N_{fr} , we do not have access to the “true dynamics” $\mathbf{Q}$, as this considers only the time evolution of a frequency subspace. This can be however useful in some cases. If we believe that just a portion of the frequency spectrum is relevant to the problem at hand, and the rest is, for instance, noise, hence not contributing to the dynamics of the underlying generative process, then we are evolving a physically relevant phenomenon cleaning the frequencies associated to noise.

The total error is a combination of the two errors in equations (25) and (26), and can be calculated as follows

$$\text{err}_{\text{total}} = \tilde{\Phi}_r \tilde{\mathbf{A}}_r - \mathbf{Q}. \quad (27)$$

3. On the connections between POD and SPOD emulation

In this section, we introduce the POD method, outline how POD modes relate to SPOD modes, and detail how POD latent space differs from the SPOD one, thereby resulting in a learning task that has a different characteristics and degree of difficulty.

3.1. The POD method

Similarly to SPOD, in the POD method we seek to identify a basis $\phi_j^{(\text{POD})}$ that best approximates the process $\mathbf{q}(\mathbf{x}, t)$ underlying our data $\mathbf{Q}$, defined in equation (2). In contrast to SPOD, for POD we seek to maximize

$$\lambda^{(\text{POD})} = \frac{E\{|\langle \mathbf{q}(\mathbf{x}, t), \phi^{(\text{POD})}(\mathbf{x}, t) \rangle_{\mathbf{x}}|^2\}}{\langle \phi^{(\text{POD})}(\mathbf{x}, t), \phi^{(\text{POD})}(\mathbf{x}, t) \rangle_{\mathbf{x}}}, \quad (28)$$

over $\phi^{(\text{POD})}(\mathbf{x}, t) \in H$, where H is a Hilbert space. The analogy with the SPOD modes described in the previous section is evident, but in this case the inner product $\langle \cdot, \cdot \rangle_{\mathbf{x}}$ is defined only in space and therefore the two-point correlation is written as $C(\mathbf{x}, \mathbf{x}') = E\{|\langle \mathbf{q}(\mathbf{x}, t), \phi^{(\text{POD})}(\mathbf{x}, t) \rangle_{\mathbf{x}}|^2\}$, being the dependence on time dropped. The function $\phi(\mathbf{x}, t)$ which maximizes the quantity $\lambda^{(\text{POD})}$ can be found by solving the eigenvalue problem: $\langle C(\mathbf{x}, \mathbf{x}'), \phi^{(\text{POD})}(\mathbf{x}') \rangle = \lambda^{(\text{POD})} \phi^{(\text{POD})}(\mathbf{x})$. In terms of discrete data $\mathbf{Q}$, the eigenvalue problem reads as $\mathbf{Q}^* \mathbf{Q} \mathbf{V}^{(\text{POD})} = \mathbf{V}^{(\text{POD})} \mathbf{\Lambda}^{(\text{POD})}$, and the POD modes can be constructed as $\Phi^{(\text{POD})} = \mathbf{Q} \mathbf{V}^{(\text{POD})}$. Hence POD modes identify an orthonormal basis which is coherent in space and optimal in sense that it minimizes the error in terms of energy between the original snapshots (2) and their orthogonal projection.

3.2. Relation between POD and SPOD modes

POD modes can be seen as a superimposition of space-time coherent structures, and can be directly related to SPOD modes. In order to show this, we start from the equivalence between the space correlation tensor and the space-time correlation tensor reported in Equation (6). For a null translation in time, the latter reads

$$C(\mathbf{x}, \mathbf{x}', 0) = \int_{-\infty}^{\infty} S(\mathbf{x}, \mathbf{x}', f) df. \quad (29)$$

This relation shows the equivalence between a space correlation and a zero-time-lag space correlation. As described in [21], the cross-spectral density tensor has the following diagonal representation:

$$S(\mathbf{x}, \mathbf{x}', f) = \sum_{k=1}^{\infty} \lambda_k(f) \psi_k(\mathbf{x}, f) \psi_k^*(\mathbf{x}', f). \quad (30)$$

Putting together Equation (29) and (30), and applying the operator $\langle \cdot, \phi(\mathbf{x}) \rangle$ to both sides one eventually obtains:

$$\lambda_j^{(\text{POD})} \phi_j^{(\text{POD})}(\mathbf{x}) = \int_{-\infty}^{\infty} \sum_{k=1}^{\infty} \lambda_k(f) c_{jk}(f) \psi_k(\mathbf{x}, f) df \quad (31)$$

being $c_{jk} = \langle \phi_k(\mathbf{x}, f), \phi_j^{(\text{POD})}(\mathbf{x}) \rangle$. This last equation shows how the POD modes can be seen as a summation of SPOD modes at many frequencies; in particular the POD modes are mixing together space-coherent dynamics which are characterized by different frequencies.

3.3. Comparison between POD and SPOD latent space

Using the POD modes $\Phi^{(\text{POD})}$ defined in section 3.1, it is possible to compute the latent space for POD, that is

$$\mathbf{A}_r^{(\text{POD})} = \Phi^{(\text{POD})} \mathbf{Q} = [a^{(\text{POD},1)}, a^{(\text{POD},2)}, \dots, a^{(\text{POD},L_r)}], \quad (32)$$

where $\mathbf{A}_r^{(\text{POD})} \in \mathbb{R}^{L_r \times N_t}$, with L_r being the number of modes retained to compute the coefficients (also referred to as truncation of the POD basis). If we compare equation (32) with equation (13), we note that, unlike POD, the SPOD latent space is richer, because it contains the frequency content of the data, that is its dimension $\mathbf{A}_r^{(\text{SPOD})} \in \mathbb{R}^{L_r \times N_t} \neq \mathbf{A}_r \in \mathbb{R}^{(L_r \times N_{\text{tr}}) \times N_t}$. It is exactly this aspect that makes SPOD and POD latent spaces different. The former contains frequency information regarding the data, and is able to capture the evolution of coherent structures for finite frequency bands. The latter contains an over-all representation of the data, without the frequency granularity permitted by SPOD. This difference leads to conceptually different emulation strategies, where for SPOD we can control the frequency information we want to retain in the data. This feature allows (1) filtering out unwanted noise, or (2) focusing on specific phenomena arising for given frequency contents. Obviously, the coefficients produced by POD and SPOD have a different time behaviors. Because of this, the neural-network learning task may be more difficult for one with respect to the other data reduction method. In section 4, we will show how the latent space (first two coefficients) for POD and SPOD (Fig. 7), and it will be possible to appreciate how the SPOD coefficients are less stationary if compared to the POD ones, thereby rendering the SPOD emulation task more difficult.

4. Results

4.1. Data layout, neural network configuration and code availability

The emulation strategy proposed in this paper aims to predict future values of the SPOD coefficients $\tilde{\mathbf{A}}_{r,\text{future}} : [\mathbf{A}_r(t+1), \mathbf{A}_r(t+2), \dots, \mathbf{A}_r(t+n_{\text{TF}})]$, provided their past values $\mathbf{A}_{r,\text{past}} : [\mathbf{A}_r(t-1), \mathbf{A}_r(t-2), \dots, \mathbf{A}_r(t-n_{\text{TP}})]$. From these, we can then compute the future data snapshots $\mathbf{Q} : [\mathbf{Q}(t+1), \mathbf{Q}(t+2), \dots, \mathbf{Q}(t+n_{\text{TF}})]$, simply by projecting the predicted coefficients onto the pre-computed modes Φ , that is

$$\mathbf{Q}(t+i) = \Phi \mathbf{A}_r(t+i). \quad (33)$$

The coefficients $\mathbf{A}_r$ are learnt using a recurrent neural network architecture, namely the long-short term memory (LSTM) network. The coefficients learnt via LSTM are denoted by $\tilde{\mathbf{A}}_r$, as opposed to the true coefficients denoted by $\mathbf{A}_r$. The task is to learn the (generally nonlinear) map $\tilde{\mathbf{A}}_{r,\text{future}} \rightarrow \mathcal{N}_{\text{LSTM}}(\mathbf{A}_{r,\text{past}})$ where one seek to minimize the difference between the true $\mathbf{A}_{r,\text{future}}$ and the predicted $\tilde{\mathbf{A}}_{r,\text{future}}$, under a given error norm (L_2 in our specific case). The LSTM neural network $\mathcal{N}_{\text{LSTM}}$ adopted for all the results presented in this section shared the same overall structure, and consisted of a sequential single LSTM layer, with hyperbolic tangent activation function, and dropout of 0.15.




Fig. 2. Learning configuration adopted for the SPOD coefficients. The grey boxes at the top represent the past values of the coefficients. The green boxes at the bottom represent the future values of the coefficients. The light blue boxes are the neural networks $\mathcal{N}$ mapping each SPOD coefficient (for all frequencies) from past to future values. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

For the learning task at hand, we divided the data available into training, validation and test sets. Consequently, the neural network was fed with

- a training input** containing $\mathbf{A}_{r,past}^{(train)}$ of dimension $[n_{samples}^{(train)}, n_{\tau P}, n_{features}]$, with $n_{samples}$ being the number of samples used for the training, $n_{\tau P}$ the number of previous time snapshots used for the future prediction, and $n_{features} = M \times N_{freq}$ the total number of features,
- a training output** containing $\mathbf{A}_{r,future}^{(train)}$ of dimensions $[n_{samples}^{(train)}, n_{\tau F} \times n_{features}]$, where $n_{\tau F}$ is the number of future steps one wants to predict
- a testing input** containing $\mathbf{A}_{r,past}^{(test)}$ of dimension $[n_{samples}^{(test)}, n_{\tau P}, n_{features}]$
- a testing output** containing $\mathbf{A}_{r,future}^{(test)}$ of dimension $[n_{samples}^{(test)}, n_{\tau F} \times n_{features}]$.

The coefficients used for testing, $\mathbf{A}_r^{(test)}$, were computed by projecting the last $n_{\tau P}^{(test)}$ data realizations on the same reduced basis (i.e., SPOD modes) obtained from the training snapshots, that is:

$$\mathbf{A}_r^{(test)} = \Phi_r^{(train)*} \mathbf{Q}^{(test)}, \quad (34)$$

such that information leakages which could inflate the prediction are prevented. We note that there is no validation set, as no hyperparameter optimization was performed. However, we used the test set to evaluate under- and over-fitting issues, by computing the corresponding loss function during training. All the coefficients $\mathbf{A}_r$ were scaled between -0.1 and 0.1 in a feature by feature fashion to avoid output saturation. To speed up neural-network training, we subdivided the coefficient matrix $\mathbf{A}_r$ into L_r distinct sub-matrices, and distributed the computation across L_r identical LSTM neural networks, where the number of features for each of them was equal to the number of frequencies used, as depicted in Fig. 2. We note that this neural network configuration also allows for coefficient-specific hyperparameter optimization, that might be beneficial for complex learning tasks.

We remark that we used the time-domain reconstruction detailed in section 2.2.1 for all results, except for section 4.3, where we compare the frequency-domain reconstruction (detailed in section 2.2.2 vs. the time-domain one).

The results presented in the following sections were obtained using the open-source *PySPOD* [41], where we implemented the neural-network emulation strategy proposed in this paper, using Keras (<https://keras.io>). The interested reader can refer to: <https://github.com/mengaldo/PySPOD>, and in particular to the Git branch *papers/spod-emulation*.

4.2. Engineering problems: a compressible jet

The first test case considered is a subsonic turbulent jet at Mach number 0.4 and Reynolds number of $4.5 \cdot 10^6$. The data, retrieved from https://github.com/SpectralPOD/spod_matlab, consists of 1000 realizations (or snapshots) of a pressure field sampled every 0.2 acoustic time units and corresponds to a coarsened version of a large-eddy simulation dataset computed by Bres et al. [42]. This coarsened dataset interpolates the original data onto a significantly smaller and uniformly spaced cylindrical grid of 22×88 points. The 1000 snapshots representing the pressure field at different time instants were split into a training and a testing set. The training set is composed of $n_{samples}^{(train)} = 800$ snapshots, while the testing set is composed




Fig. 3. SPOD eigenvalues plotted against period: the first few modes are able to capture most of the energy of the original system.




Fig. 4. True pressure field vs. POD reconstructed pressure field (left plot) and SPOD reconstructed pressure field (right plot). The reconstruction used ten modes for both cases (albeit the number of modes for SPOD is larger, and equal to $c = 10(\text{modes}) \times 33(\text{frequencies})$ as we use all 33 frequencies to reconstruct the pressure field). Both figures show time snapshot 500.

of $n_{\text{samples}}^{(\text{test})} = 200$. To compute the SPOD modes, the snapshots of the training set were subdivided into $L = 24$ statistically independent blocks, each of them containing 64 snapshots with a $\ell_{\text{ovlp}} = 50\%$ overlap. No weighting was adopted for the inner product, thus resulting in the weight matrix being equal to the identity matrix $\mathbf{W} = \mathbf{I}$. This choice of parameters provides up to 24 modes and 33 frequencies. The SPOD modes were generated using the methods illustrated in section 2, starting from the snapshots containing the pressure fluctuations computed by subtracting the time mean of the pressure field from each realization. In the following two subsections, we show the comparison between POD and SPOD emulation, and the sensitivity with respect to the number of modes and the number of future values predicted n_{τ_F} – also referred to as forecast horizon – (section 4.2.1), as well as the sensitivity of SPOD emulation to the frequencies retained (section 4.2.2).

4.2.1. POD vs. SPOD emulation: sensitivity to number of modes and forecast horizon

In this section, we compare POD and SPOD emulation for different number of modes retained and different forecast horizons. We retain all 33 frequencies (equivalently periods) that the SPOD provides for each eigenvalue (and associated mode) for this test case, as depicted in Fig. 3. In contrast, the POD yields a set of eigenvalues and associated modes, that do not contain this frequency dimension. From Fig. 3, it is possible to see how the first few eigenvalues (thus associated SPOD modes) are able to capture the majority of the energy contained in the original snapshots. Therefore, as expected, we can reconstruct the dynamics of the jet pressure field by superimposing the first few leading modes multiplied by the associated coefficients computed using the oblique projection introduced in section 2.2. This is true for both POD and SPOD.

An example of pressure field reconstruction for both methods is depicted in Fig. 4. Here, we show a qualitative comparison between the 500-th snapshot of the pressure field and the corresponding reconstructed solution, in which the pressure fluctuations were obtained by superimposing the first ten POD modes (left plot), and the first ten SPOD modes (right plot).

Fig. 5 shows the decay of the averaged L2-norm error against the number of modes for POD (left plot) and SPOD (right plot). We remark again that the SPOD uses all frequencies for the ten modes considered to reconstruct the pressure field. We note how the error of the SPOD reconstruction is generally smaller than the one for the POD reconstruction (note that the y-axis is different for the POD and SPOD). However, the SPOD reconstruction requires a larger amount of data, as we need to consider the frequency dimension for each mode (i.e., 33 frequencies for this test case), that is not present for the




Fig. 5. Decay of the L2-norm error as a function of the number of POD (left plot) and SPOD (right plot) modes used for reconstructing the pressure field.




Fig. 6. Loss histories of the training phase (right figures) and comparison between the first time coefficient and the corresponding neural network prediction (left figures). The two top figures correspond to POD, while the bottom two to the SPOD.

POD reconstruction. Therefore, the data compression achieved for SPOD differs from the one achieved for POD, by a factor associated to the number of frequencies retained for the reconstruction, that is: $r_c = \text{no. modes SPOD} / (\text{no. modes POD} \times N_{\text{fr}}) = 1/N_{\text{fr}}$. The neural-network prediction of the coefficients is depicted in Fig. 6, along with the training and validation (that in this case coincides with testing as no hyperparameter optimization was performed) loss as a function of the number of epochs. The results were obtained using the LSTM configuration described in section 4.1, where the single neural network layer employed contained 25 neurons. The total number of epochs for training was equal to 130. The number of past values adopted to make the predictions was $n_{\tau P} = 60$, and the number of future values predicted was $n_{\tau F} = 1$, with a total number of samples for training equal to $n_{\text{samples}}^{(\text{train})} = 749$. The total number of features for SPOD was $n_{\text{features}}^{(\text{train, SPOD})} = 330$ since the first 10 modes were used, resulting in a data compression of $c = 330 / (88 \times 22) = 0.17$, that is, the latent space had a dimension that was only 17% of the original data. The total number of features for POD was instead $n_{\text{features}}^{(\text{train, POD})} = 10$, achieving a data compression of $c = 10 / (88 \times 22) = 0.005$, corresponding to a latent space containing an amount of data equal to only 0.5% of the original data. To this end, the data compression achieved via POD is larger. However, the granularity provided by the SPOD in terms of frequency space is not available. Therefore, the trade-off is application dependent, and revolves around what the frequency information content can provide in terms of analysis and prediction. Fig. 6 shows how the LSTM prediction for POD (top plots) seems to achieve better performance than the SPOD one (bottom plots). This can be explained by looking at the POD and SPOD coefficients. On the one hand, the POD coefficients are smoother, and resemble a harmonic behavior. On the other hand, the SPOD coefficients are less smooth, possibly less stationary, and more numerous.




Fig. 7. First two POD (on the left) for the same frequency and first two SPOD modes (on the right).




Fig. 8. Comparison among: true solution (top), reconstructed solution obtained by projecting the snapshots on the modal basis (middle), reconstructed solution with the predicted LSTM coefficients (bottom). The figures on the left are for POD, while the figures on the right are for SPOD.

A comparison of the first two POD modes (left plot) and the first two SPOD modes (right plot) is depicted in Fig. 7, where the characteristics described are particularly evident. These different characteristics of the POD and SPOD latent spaces lead to a more challenging neural-network learning task for the SPOD coefficients, thereby negatively affecting the SPOD-LSTM forecasts. One should note that we did not carry out an hyperparameter optimization, nor did we use different neural-network architectures (including attention-mask-based neural network architectures, that can improve the learning performance [43]) as these two aspects are out of the scope of this paper. However, these two points can improve the SPOD-LSTM predictions. In addition, having longer time histories for the coefficients may further improve the predictions. Fig. 8 shows the comparison among the true solution, the solution reconstructed with the coefficients obtained with a direct projection of the snapshots onto the reduced space, and the one with coefficients predicted with the LSTM network. The different errors, described in section 2.4, are reported in Table 1. These are the projection error, due to the projection of the true dynamics onto a reduced space, the learning error, due to the neural-network, and the total error, that is a combination of the previous two errors. For the POD, the projection error is consistently larger than the learning error. This indicates that the dominant part of the total error is related to the projection phase, while the LSTM network was very effective in predicting the future values of the coefficients. For the SPOD, the situation is reversed, with the learning error being dominant. Two main elements can explain this different behavior. The first one is due to the fact that the time series of the SPOD coefficients are more irregular and unsteady with respect to the POD ones and hence more complex to predict (depicted qualitatively in Fig. 7). The second one is connected to the latent SPOD space dimension, which is significantly larger than the POD one and therefore more prone to accumulation of learning errors. The higher number of features impacts negatively the costs of the training phase of the LSTM network and the one of the model inference. This negative impact of the number of modes on the error is confirmed by observing the L_{inf} -norm errors reported in Table 1: the learning error is significant when a high number of SPOD modes are taken into account even if the L_2 -norm error is relatively small, thus indicating that the error is high on a limited number of computational points.




Fig. 9. Projection error, learning error, and total error for POD (left) and SPOD (right). Projection and total error for POD are overlapped.

Table 1

L_1 , L_2 , and L_{inf} norm errors evaluated considering a different number of modes for both POD (top table) and SPOD (bottom table).

POD Nr. of Modes	Projection Error			Learning Error			Total Error		
	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
2	3.04e-03	1.19e-04	2.96e-02	6.36e-05	3.45e-06	8.12e-04	3.04e-03	1.19e-04	2.97e-02
4	2.81e-03	1.11e-04	2.88e-02	9.95e-05	5.02e-06	1.13e-03	2.81e-03	1.11e-04	2.88e-02
6	2.72e-03	1.07e-04	2.82e-02	1.88e-04	8.60e-06	1.82e-03	2.74e-03	1.08e-04	2.82e-02
8	2.65e-03	1.04e-04	2.77e-02	2.40e-04	1.07e-05	2.29e-03	2.68e-03	1.05e-04	2.78e-02
10	2.49e-03	9.77e-05	2.70e-02	3.05e-04	1.28e-05	2.71e-03	2.53e-03	9.89e-05	2.70e-02
12	2.36e-03	9.27e-05	2.69e-02	3.42e-04	1.43e-05	3.18e-03	2.40e-03	9.43e-05	2.71e-02
14	2.26e-03	8.80e-05	2.64e-02	4.06e-04	1.67e-05	3.69e-03	2.31e-03	9.00e-05	2.67e-02

SPOD Nr. of Modes	Projection Error			Learning Error			Total Error		
	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
2	2.81e-03	9.92e-05	2.32e-02	9.50e-04	4.08e-05	1.21e-02	3.01e-03	1.07e-04	2.52e-02
4	2.33e-03	7.93e-05	1.90e-02	1.34e-03	5.64e-05	1.52e-02	2.79e-03	9.75e-05	2.24e-02
6	2.04e-03	6.82e-05	1.45e-02	1.69e-03	7.04e-05	1.86e-02	2.81e-03	9.81e-05	2.17e-02
8	1.79e-03	5.89e-05	1.18e-02	2.45e-03	1.03e-04	2.55e-02	3.26e-03	1.19e-04	2.65e-02
10	1.55e-03	5.07e-05	9.82e-03	3.47e-03	1.47e-04	3.37e-02	4.05e-03	1.55e-04	3.39e-02
12	1.35e-03	4.40e-05	8.56e-03	6.34e-03	2.71e-04	6.45e-02	6.71e-03	2.74e-04	6.47e-02
14	1.14e-03	3.73e-05	7.32e-03	1.97e-01	9.09e-03	2.45e+00	1.97e-01	9.09e-03	2.45e+00

Table 2

Computational times (in seconds) for both POD and SPOD for a different number of modes.

Nr. of Modes	POD		SPOD	
	Offline	Online	Offline	Online
2	68.45	3.92	268.57	14.02
4	69.29	3.82	537.73	28.63
8	71.1	3.95	1166.58	56.20
10	68.08	4.04	1369.86	69.31

For both POD and SPOD, the learning error tends to increase with the number of modes, while the projection error decreases (as also shown in Fig. 9).

The number of modes strongly influences the computational time when SPOD is considered since, as described in Section 4.1, an increasing number of neural networks has to be trained; Table 2 reports the computational times for both POD and SPOD when a different number of modes is considered. Two different phases are considered in the table aforementioned: the Offline phase, which includes both the generation of the latent space and the training of the neural networks, and the Online phase, during which the high dimensional flow fields are predicted.

Fig. 10 shows the projection, learning and total error for a five different forecast horizons $n_{\tau F} = [1, 2, 3, 4, 5]$. The projection error is obviously constant, while the learning error increases as the forecast horizon increases. This is expected, as a longer forecast horizon is typically more difficult to predict than a shorter one. We can also observe how the learning error drives the total error for SPOD, while the projection error drives the total error for POD.




Fig. 10. Projection error, learning error, and total error for POD and SPOD using different values of $n_{\tau F}$.




Fig. 11. Sensitivity to frequency: L2-norm error plot for different numbers of frequency considered and fixed $M=10$: $n_{\tau F} = 1$ (left) and $n_{\tau F} = 3$ (right).

Table 3

L_1 , L_2 , and L_{inf} norm errors evaluated retaining a different number of SPOD frequencies.

SPOD Nr. of Freqs	Projection Error			Learning Error			Total Error		
	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
5	2.70e-03	9.41e-05	2.11e-02	1.50e-03	6.33e-05	1.11e-02	3.19e-03	1.14e-04	2.36e-02
10	2.44e-03	8.36e-05	1.74e-02	2.48e-03	1.05e-04	2.07e-02	3.70e-03	1.36e-04	2.75e-02
15	2.21e-03	7.46e-05	1.48e-02	2.55e-03	1.08e-04	2.30e-02	3.59e-03	1.32e-04	2.76e-02
21	2.02e-03	6.69e-05	1.27e-02	2.73e-03	1.16e-04	2.59e-02	3.65e-03	1.35e-04	2.85e-02
25	1.91e-03	6.27e-05	1.19e-02	3.01e-03	1.27e-04	2.98e-02	3.82e-03	1.42e-04	3.15e-02
33	1.55e-03	5.07e-05	9.82e-03	6.37e-03	2.70e-04	5.77e-02	6.82e-03	2.75e-04	5.81e-02

4.2.2. SPOD emulation: sensitivity to frequency

In this section, we show the sensitivity of the emulation framework to the number of frequencies retained for each SPOD mode. In particular, Fig. 11 shows the projection, learning and total errors for forecast horizon $n_{\tau F} = 1$ (left plot), and $n_{\tau F} = 3$ (right plot). As we increase the number of frequencies retained, the learning error increases, while the projection error decreases, as expected. Table 3 tabulated the errors depicted in Fig. 11. Here, it is important to note that by removing higher frequencies we have a smaller learning error. This aspect can have important applications when we need to perform predictive tasks on noisy data.

4.3. Frequency domain reconstruction

In Section 2.2 two distinct techniques for the evaluation of the expansion coefficients were presented; the aim of this section is to provide a comparison between the frequency-domain approach and the time-domain one used in the previous sections.

One of the main drawback of the frequency-domain approach, when applied to neural network, is that the dimension of the latent space is proportional to the number of blocks L and not to the number of time snapshots. Therefore, in order to have an adequate number of samples for training the LSTM-based neural network, the data set containing the 800 flow realizations of the training set has been split into 687 blocks with an overlap of 98%, i.e. two contiguous blocks differ for just one snapshot. The training phase of the $L \times N_f$ expansion coefficients has been carried out with the aid of the neural




Fig. 12. Loss histories of the training phase (on the left) and comparison between the first two modes series used for the training of the LSTM network.

Table 4

L_1 , L_2 , and L_{inf} norm errors evaluated retaining a different number of SPOD modes for frequency domain reconstruction.

SPODf Nr. of Freqs	Projection Error			Learning Error			Total Error		
	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
2	3.65e-03	1.34e-04	2.79e-02	2.34e-04	1.11e-05	4.21e-03	3.65e-03	1.34e-04	2.79e-02
4	3.58e-03	1.28e-04	2.76e-02	2.60e-04	1.21e-05	4.58e-03	3.58e-03	1.28e-04	2.74e-02
6	3.53e-03	1.27e-04	2.84e-02	2.77e-04	1.27e-05	4.72e-03	3.53e-03	1.27e-04	2.83e-02
8	3.54e-03	1.25e-04	2.64e-02	2.91e-04	1.31e-05	4.67e-03	3.54e-03	1.25e-04	2.63e-02
10	3.53e-03	1.25e-04	2.72e-02	3.01e-04	1.33e-05	4.67e-03	3.53e-03	1.25e-04	2.70e-02
12	3.52e-03	1.25e-04	2.71e-02	3.12e-04	1.36e-05	4.67e-03	3.53e-03	1.25e-04	2.71e-02
14	3.50e-03	1.24e-04	2.65e-02	3.19e-04	1.38e-05	4.70e-03	3.51e-03	1.25e-04	2.64e-02

network configured as described in Section 4.1. The histories of the losses recorded during the training are reported in Fig. 12a and they demonstrate the suitability of the chosen hyper-parameters also for the current task.

The coefficients associated to the first and the second modes are shown in Fig. 12b and their shape is much more regular and sinusoidal-like with respect to the ones obtained in the time-domain (Fig. 7). Following this idea, we expect a lower learning error for the frequency domain case and this is confirmed by the results reported in Table 4 where norm errors are reported when different numbers of modes are considered. On the other hand the oblique projection used in the time-domain approach guarantees the best approximation in a least-square sense and this provide a justification for its smaller projection errors that can be observed. In the present work we have chosen to work mainly with the time domain reconstruction since it leads to a number of training samples which is the same of the POD case and this guarantees a consistent comparison between the two techniques.

4.4. Geophysical problems: the Madden-Julian Oscillation

The second test case which was performed considering a geophysical tropical phenomenon known as the Madden-Julian Oscillation (MJO). This is an intraseasonal phenomenon whose characteristic period varies between 30 and 90 days and it is the result of a large-scale coupling between atmospheric circulation and deep convection. This pattern slowly propagates eastward with a speed of 4 to 8 m s^{-1} . Although it is commonly known as an oscillation, the large ranges of variation of both its local period (30-90 days) and its speed of propagation suggest that MJO is a rather irregular phenomenon which can be seen at a large-scale level as a mix of multiple high-frequency, small-scale convective phenomena. Further information about MJO can be found in [44].

The aim of the present case is to test the capabilities of the SPOD-based approach in isolating and predicting the evolution in time of phenomena characterized by specific subsets of frequencies. The capability of breaking down complex flow fields in simpler structures united by similar characteristic time-scales could be beneficial for applications involving, for instance, data denoising, filtering, and frequency analysis.

In this test case the starting dataset was composed by 5000 time snapshots, each of these providing a representation of the total precipitation with a resolution of 12 hours. The training data set consisted of the first 75% samples (equal to $n_{\text{samples}}^{(\text{train})} = 3750$), while the remaining 25% samples (equal to $n_{\text{samples}}^{(\text{test})} = 1250$) were used for testing and validation purposes. The spatial dimension of the data consists of $480 \times 241 = 115680$ points, where 480 are longitude points and 241 are latitude points.

The snapshots in the training set were split into $L = 5$ statistically independent blocks, each of them containing 730 snapshots with a $\ell_{\text{ovlp}} = 0\%$ overlap. The weighting for the inner product is based on the surface element of a sphere $dS_r = r^2 \sin \theta d\theta d\varphi$ with radius r over $(\theta, \varphi) \in [-\pi/2, \pi/2] \times [-\pi, \pi]$, where we denote the spherical coordinate triplet $\mathbf{x} = (r, \theta, \varphi)$ consisting of radius, geographic latitude and geographic longitude, respectively. The choice of parameters just introduced provides up to 5 modes and 366 frequencies. The SPOD modes were generated using the methods illustrated in




Fig. 13. The 100-th snapshot of the database: total precipitation of rain (left plot). Eigenvalues of the first 5 modes plotted against period obtained using SPOD (right plot).

Table 5

L_1 , L_2 , and L_{inf} norm errors evaluated considering a different number of modes for both POD (top table) and SPOD (bottom table).

POD				Learning Error			Total Error		
Nr. of Modes	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
1	4.86e-04	2.36e-06	3.12e-02	7.49e-06	3.82e-08	9.91e-05	4.86e-04	2.36e-06	3.12e-02
2	4.83e-04	2.34e-06	3.12e-02	1.05e-05	6.07e-08	2.42e-04	4.83e-04	2.35e-06	3.12e-02
3	4.80e-04	2.34e-06	3.12e-02	1.30e-05	7.26e-08	3.90e-04	4.81e-04	2.34e-06	3.12e-02
4	4.79e-04	2.34e-06	3.12e-02	1.55e-05	8.64e-08	4.13e-04	4.80e-04	2.34e-06	3.12e-02
5	4.79e-04	2.33e-06	3.11e-02	3.67e-05	2.37e-07	1.05e-03	4.83e-04	2.35e-06	3.12e-02

SPOD				Learning Error			Total Error		
Nr. of Modes	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
1	3.33e-04	1.91e-06	2.97e-02	1.81e-04	7.54e-07	3.51e-03	3.27e-04	2.05e-06	3.18e-02
2	3.38e-04	1.83e-06	2.81e-02	2.42e-04	1.01e-06	5.58e-03	3.44e-04	2.08e-06	3.17e-02
3	3.38e-04	1.75e-06	2.67e-02	2.86e-04	1.20e-06	7.23e-03	3.63e-04	2.11e-06	3.17e-02
4	3.35e-04	1.69e-06	2.56e-02	3.04e-04	1.29e-06	8.65e-03	3.61e-04	2.11e-06	3.17e-02
5	3.31e-04	1.63e-06	2.45e-02	3.34e-04	1.42e-06	9.84e-03	3.80e-04	2.15e-06	3.17e-02

section 2, starting from the snapshots containing the rain precipitation fluctuations computed by subtracting the time mean from each realization.

The neural network used for learning the dynamics of the SPOD latent space is the same whose configuration has been described in section 4.1.

In the following subsections, we provide a comparison between the results obtained by using the POD-based approach and SPOD-based one, and the sensitivity with respect to some significant parameters such as the number of modes and the number of frequencies retained for the reconstruction is discussed in section 4.4.2 and section 4.4.1, respectively.

4.4.1. POD vs. SPOD emulation: sensitivity to number of modes

In this section the influence on the errors of the numbers of modes is investigated.

First, the SPOD modes are evaluated starting from the flow realizations contained in the training dataset; a visualization of a snapshot belonging to this database is reported in Fig. 13a. The eigenvalues associated to the first five modes at different characteristic periods are represented in the plot shown in Fig. 13b. The influence of the first mode is noticeable in the region associated to low frequencies; in the plot dashed lines help to identify the characteristic period in which the MJO phenomenon is expected to be significant.

Following the procedure described in 2, the coefficients are evaluated and then the latent dynamics are learnt via a neural network for both the SPOD-based approach and the POD-based one. The SPOD coefficients are obtained via time-domain reconstruction, as described in section 2.2.1. The results of the procedure are shown in Fig. 14, where the reconstructions of the solution in time after projection (top plots) and after LSTM learning (bottom plots) are shown. The left plots were obtained using 1 SPOD mode, while the right plots using all the 5 SPOD modes. In all the plots in Fig. 14 all frequencies were retained and a forecast horizon $n_{TF} = 1$ was used. By comparing these figures against the snapshot depicted in Fig. 13a one can observe that, in this specific test case, both the projection- and LSTM-based reconstructions are visibly different from the true solution. Yet, the LSTM-based solution resembles the projection based one. Therefore, the largest source of error in this case is due to the projection, not to the learning task (Table 5).




Fig. 14. Projection solution (top) vs LSTM prediction (bottom) at the same time snapshot = 100-th, retaining 1 SPOD mode (left) and 5 SPOD modes (right). For each mode retained, all frequencies were used. The forecast horizon adopted was $n_{\tau F} = 1$ (corresponding to 12 hours).




Fig. 15. Projection error, learning error, and total error plotted against the number of modes for forecast horizon $n_{\tau F} = 1$ (corresponding to 12 hours): POD (left) and SPOD retaining all frequencies for each mode (right). For POD, projection and total error are overlapped.

Figs. 15 and 16 show the projection, learning and total errors as a function of the number of POD (left plots), and SPOD (right plots) modes, for forecast horizons $n_{\tau F} = 1$ (12 hours) and $n_{\tau F} = 3$ (36 hours), respectively. Similarly to the test case presented in section 4.2, increasing the number of modes for both POD and SPOD increases the learning error, while there is a slight decrease in the projection error for both POD and SPOD. The total error for POD is driven by the projection error, while the total error for SPOD is a combination of the projection and learning errors. The plots for SPOD in both Fig. 15 and 16 are obtained using all the 366 frequencies computed for the SPOD approach. We note that the dimensions of the latent space for POD ranged from 1 (mode) to 5 (modes), achieving a compression of [1 to 5]/115680. The one for SPOD ranged from $1 \times 366 = 366$ to $5 \times 366 = 1830$, resulting in a data compression of [366 to 1830]/115680, that is two orders of magnitude larger than the one obtained using POD.

4.4.2. SPOD emulation: sensitivity to number of frequencies

In this section, we show the sensitivity of the emulation strategy with respect to the number of frequencies retained for the SPOD reduction, similarly to what we showed in section 4.2. In detail, Fig. 17 shows the L2 errors for different numbers of SPOD frequencies retained. As expected (and similar to the jet test case), a higher number of retained frequencies tends




Fig. 16. Projection error, learning error, and total error plotted against the number of modes for forecast horizon $n_{\tau F} = 3$ (corresponding to 36 hours): POD (left) and SPOD retaining all frequencies for each mode (right). For POD, projection and total error are overlapped.




Fig. 17. Projection, learning, and total L_2 errors plotted against the number of frequencies, retaining 3 modes, for forecast horizon $n_{\tau F} = 1$ (left plot), and $n_{\tau F} = 3$ (right plot).

Table 6

L_1 , L_2 , and L_{inf} norm errors evaluated retaining a different number of SPOD frequencies for $n_{\tau F} = 1$ (top) and $n_{\tau F} = 3$.

$n_{\tau F} = 1$	Projection Error			Learning Error			Total Error		
Nr. of Freqs	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
4	3.23e-04	2.05e-06	3.18e-02	3.43e-05	1.46e-07	5.1e-04	3.23e-04	2.05e-06	3.19e-02
12	3.25e-04	2.02e-06	3.16e-02	6.04e-05	2.77e-07	9.12e-04	3.26e-04	2.05e-06	3.19e-02
46	3.31e-04	1.96e-06	3.09e-02	1.41e-04	4.97e-07	1.92e-03	3.41e-04	2.04e-06	3.18e-02
104	3.37e-04	1.91e-06	2.99e-02	2.01e-04	8.01e-07	2.46e-03	3.52e-04	2.08e-06	3.18e-02
146	3.40e-04	1.85e-06	2.92e-02	2.25e-04	9.16e-07	3.50e-03	3.59e-04	2.10e-06	3.18e-02
365	3.41e-04	1.76e-06	2.68e-02	2.67e-04	1.21e-06	5.21e-03	3.68e-04	2.13e-06	3.18e-02
$n_{\tau F} = 3$	Projection Error			Learning Error			Total Error		
Nr. of Freqs	L_1	L_2	L_∞	L_1	L_2	L_∞	L_1	L_2	L_∞
4	3.23e-04	2.05e-06	3.18e-02	5.19e-05	2.21e-07	7.53e-04	3.25e-04	2.06e-06	3.19e-02
12	3.25e-04	2.03e-06	3.16e-02	1.04e-04	4.31e-07	1.42e-03	3.35e-04	2.08e-06	3.19e-02
46	3.31e-04	1.98e-06	3.09e-02	1.80e-04	7.38e-07	2.49e-03	3.56e-04	2.11e-06	3.18e-02
104	3.37e-04	1.92e-06	2.99e-02	2.32e-04	9.55e-07	3.65e-03	3.69e-04	2.14e-06	3.19e-02
146	3.40e-04	1.89e-06	2.92e-02	2.50e-04	1.03e-06	4.38e-03	3.72e-04	2.15e-06	3.19e-02
365	3.41e-04	1.76e-06	2.68e-02	3.14e-04	1.31e-06	7.51e-03	3.88e-04	2.19e-06	3.19e-02

to reduce the projection error while, at the same time, affecting negatively the learning error. The two plots are referred to two different forecast horizons (12 and 36 hours, on the left and on the right, respectively) (Table 6).

The prediction of the 300-th snapshot belonging to the testing set is reported in Fig. 18. In particular, in all plots we retained the same number of frequencies (6), but we used different period intervals: one for low frequencies (70–365 days), one for intraseasonal frequencies (30–50 days), and one for high frequencies (23–30 days). The left plots show the projection-based solutions, while the right plots show the LSTM predictions for forecast horizon $n_{\tau F} = 1$ (12 hours). The LSTM forecast, is qualitatively similar to the projection-based solution, as also highlighted by the relatively small L_2 errors reported in Fig. 17.




Fig. 18. Comparison between the solution obtained with the coefficients projected onto the reduced space (left) and the one evaluated starting from the one learnt by the LSTM using $n_{TF} = 1$ (right) for three period ranges: 23 to 30 (top), 30 to 50 (middle), and 60 to 365 days. The results have been obtained using the same number of frequencies (i.e., 6), retaining 3 modes, and considering the same snapshot.

5. Conclusions

This work has been driven by recent advancement in SPOD-based reduced order modeling for quasi-stationary data. In particular, we presented a surrogate model for the SPOD latent space via recurrent neural networks based on three fundamental blocks: a) extraction of the SPOD modes from a database of previously computed snapshots, b) estimation of coefficients which are able to reconstruct the original data realization when applied to the SPOD modes through an oblique projection, and c) prediction of future evolution in time of the data with the aid of LSTM neural networks. The resulting model is non-intrusive, stable, and it is able to reproduce the dynamics of the underlying data.

Moreover an extensive comparison between a SPOD-based model and its POD-based counterpart has been carried out on two distinct test cases. The numerical results show that SPOD modes provide a larger and more refined basis and this contributes to reduce the projection error with respect to a standard approach which relies on POD. The main drawback is a consistent computational overhead with respect to POD since we consider the evolution in time of multiple dynamics characterized by small frequency ranges. The large size of the latent space affects also the quality of the final predicted solution: the tests we have presented show that the learning error in SPOD tends to grow with the number of modes since numerical accumulation of errors occurs. In addition plots of the evolution of the SPOD coefficients in time show a behavior which is less harmonic-like and regular than the one of POD. The learning error could be probably reduced with a refined tuning of the hyperparameters of the LSTM networks, but the general considerations reported above still hold; the refinement of the parameters was out of purpose of the present work, which is focused on the methodological aspects.

The main advantage of SPOD over the POD model consists into its versatility: in fact the user is allowed to provide a range of frequency and follow the evolution of phenomena which are relevant for the selected time scale. As it has been

pointed out in previous work [26,27], this is fundamental when one has to carry out frequency analysis, data de-noising, or prewhitening, procedures which are common both in industry and in environmental sciences. The dependence of the L2-norm errors for different numbers of retained frequencies have been studied for all the test cases proposed and, as expected, a higher number of frequencies is able to produce accurate reconstructions of the full data field but, as for the number of modes, a bigger latent space implies a bigger learning error.

The framework proposed is general and can be adopted in conjunction with alternative neural-network architectures that may improve the learning error. Extensions of the present work could aim at improving the efficiency and the accuracy of the model; this can be achieved either by investigating alternative networks and strategies for automatic optimization of the hyperparameters or investigating the application of sparse identification of nonlinear dynamics techniques to SPOD-based reduced models. In addition, a comparison against nonlinear methods for seeking latent space with the aid of AI-based architecture (e.g. autoencoders or diffusion maps) will be carried out in upcoming works. Furthermore, understanding the interplay between the projection, learning and total errors constitutes an additional and crucial area of research for non-intrusive surrogate models like the one presented in this paper.

CRediT authorship contribution statement

G. Rozza supervised research activities with G. Mengaldo, as well as revision(s). A. Lario coordinated reduced order modeling activities and simulations. R. Maulik and T. Schmidt took care of the full order part and simulation.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Gianluigi Rozza reports financial support was provided by European Research Council.

Data availability

Data will be made available on request.

Acknowledgements

This work was partially supported by MIUR (Italian ministry for university and research) through FARE-X-AROMA-CFD project, P.I. Prof. Gianluigi Rozza. Gianmarco Mengaldo acknowledges support from NUS startup grant R-265-000-A36-133. This material is partially based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357. Prof. Gianluigi Rozza acknowledges the support by European Union Funding for Research and Innovation – Horizon 2020 Program – in the framework of European Research Council Executive Agency: Consolidator Grant H2020 ERC CoG 2015 AROMA-CFD project 681447 “Advanced Reduced Order Methods with Applications in Computational Fluid Dynamics”.

References

- [1] J.S. Hesthaven, S. Ubbiali, Non-intrusive reduced order modeling of nonlinear problems using neural networks, *J. Comput. Phys.* 363 (2018) 55–78.
- [2] Q. Wang, N. Ripamonti, J.S. Hesthaven, Recurrent neural network closure of parametric POD-Galerkin reduced-order models based on the Mori-Zwanzig formalism, *J. Comput. Phys.* 410 (2020) 109402.
- [3] R. Maulik, A. Mohan, B. Lusch, S. Madireddy, P. Balaprakash, D. Livescu, Time-series learning of latent-space dynamics for reduced-order model closure, *Phys. D: Nonlinear Phenom.* 405 (2020) 132368.
- [4] B. Hamzi, R. Maulik, H. Owhadi, Simple, low-cost and accurate data-driven geophysical forecasting with learned kernels, *Proc. R. Soc. A* 477 (2021) 20210326.
- [5] W. Chen, Q. Wang, J.S. Hesthaven, C. Zhang, Physics-informed machine learning for reduced-order modeling of nonlinear problems, *J. Comput. Phys.* (2021) 110666.
- [6] S. Hijazi, G. Stabile, A. Mola, G. Rozza, Data-driven POD-Galerkin reduced order model for turbulent flows, *J. Comput. Phys.* 416 (2020) 109513.
- [7] F. Pichi, F. Ballarin, G. Rozza, J.S. Hesthaven, An artificial neural network approach to bifurcating phenomena in computational fluid dynamics, *arXiv preprint arXiv:2109.10765*, 2021.
- [8] D. Papapicco, N. Demo, M. Girfoglio, G. Stabile, G. Rozza, The neural network shifted-proper orthogonal decomposition: a machine learning approach for non-linear reduction of hyperbolic equations, *arXiv preprint arXiv:2108.06558*, 2021.
- [9] N. Demo, M. Strazzullo, G. Rozza, An extended physics informed neural network for preliminary analysis of parametric optimal control problems, *arXiv preprint arXiv:2110.13530*, 2021.
- [10] N.V. Shah, M. Girfoglio, P. Quintela, G. Rozza, A. Lengomin, F. Ballarin, P. Barral, Finite element based model order reduction for parametrized one-way coupled steady state linear thermomechanical problems, *arXiv preprint arXiv:2111.08534*, 2021.
- [11] A.T. Mohan, D.V. Gaitonde, A deep learning based approach to reduced order modeling for turbulent flow control using LSTM neural networks, *arXiv preprint arXiv:1804.09269*, 2018.
- [12] F.J. Gonzalez, M. Balajewicz, Deep convolutional recurrent autoencoders for learning low-dimensional feature dynamics of fluid systems, *arXiv preprint arXiv:1808.01346*, 2018.
- [13] K. Hasegawa, K. Fukami, T. Murata, K. Fukagata, CNN-LSTM based reduced order modeling of two-dimensional unsteady flows around a circular cylinder at different Reynolds numbers, *Fluid Dyn. Res.* 52 (2020) 065501.
- [14] R. Maulik, B. Lusch, P. Balaprakash, Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders, *Phys. Fluids* 33 (2021) 037106.

- [15] S.M. Rahman, S. Pawar, O. San, A. Rasheed, T. Iliescu, Nonintrusive reduced order modeling framework for quasigeostrophic turbulence, *Phys. Rev. E* 100 (2019) 053306.
- [16] K.R. dos Santos, D.G. Giovanis, K. Kontolati, D. Loukrezis, M.D. Shields, Grassmannian diffusion maps based surrogate modeling via geometric harmonics, *arXiv preprint arXiv:2109.13805*, 2021.
- [17] S. Lee, K. Jang, H. Cho, H. Kim, S. Shin, Parametric non-intrusive model order reduction for flow-fields using unsupervised machine learning, *Comput. Methods Appl. Mech. Eng.* 384 (2021) 113999.
- [18] G. Berkooz, P. Holmes, J.L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annu. Rev. Fluid Mech.* 25 (1993) 539–575.
- [19] R. Maulik, G. Mengaldo, PyParSVD: a streaming, distributed and randomized singular-value-decomposition library, *arXiv preprint arXiv:2108.08845*, 2021.
- [20] J.L. Lumley, *Stochastic Tools in Turbulence*, Academic Press, New York, 1970.
- [21] A. Towne, O.T. Schmidt, T. Colonius, Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, *arXiv:1708.04393*, 2017, <https://arxiv.org/pdf/1708.04393>.
- [22] O.T. Schmidt, G. Mengaldo, N.P. Wedi, Spectral empirical orthogonal function analysis of weather and climate data, *Mon. Weather Rev.* 147 (2019) 2979–2995.
- [23] B.F. Farrell, P.J. Ioannou, Generalized stability theory. Part I: autonomous operators, *J. Atmos. Sci.* 53 (1996) 2025–2040.
- [24] A.H. Monahan, J.C. Fyfe, M.H.P. Ambaum, D.B. Stephenson, G.R. North, Empirical orthogonal functions: the medium is the message, *J. Climate* 22 (2009) 6501–6514.
- [25] P.D. Sardeshmukh, P. Sura, Reconciling non-Gaussian climate statistics with linear dynamics, *J. Climate* 22 (2009) 1193–1207.
- [26] A. Nekkanti, O.T. Schmidt, Frequency–time analysis, low-rank reconstruction and denoising of turbulent flows using SPOD, *J. Fluid Mech.* 926 (2021).
- [27] T. Chu, O.T. Schmidt, A stochastic SPOD-Galerkin model for broadband turbulent flows, *Theor. Comput. Fluid Dyn.* (2021) 1–24.
- [28] A. Ghate, A. Towne, S. Lele, Broadband reconstruction of inhomogeneous turbulence using spectral proper orthogonal decomposition and Gabor modes, *J. Fluid Mech.* 888 (2020).
- [29] P. Welch, The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms, *IEEE Trans. Audio Electroacoust.* 15 (1967) 70–73.
- [30] L. Sirovich, Turbulence and the dynamics of coherent structures, *Q. Appl. Math.* 45 (1987) 561–571.
- [31] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (1997) 1735–1780.
- [32] J. Chung, C. Gulcehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, *arXiv preprint arXiv:1412.3555*, 2014.
- [33] R.T. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, *arXiv preprint arXiv:1806.07366*, 2018.
- [34] M. Jaderberg, K. Simonyan, A. Zisserman, et al., Spatial transformer networks, *Adv. Neural Inf. Process. Syst.* 28 (2015) 2017–2025.
- [35] C. Lea, M.D. Flynn, R. Vidal, A. Reiter, G.D. Hager, Temporal convolutional networks for action segmentation and detection, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 156–165.
- [36] A. Sherstinsky, Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network, *Phys. D: Nonlinear Phenom.* 404 (2020) 132306.
- [37] H. Mori, Transport, collective motion, and Brownian motion, *Prog. Theor. Phys.* 33 (1965) 423–455.
- [38] R. Zwanzig, Nonlinear generalized Langevin equations, *J. Stat. Phys.* 9 (1973) 215–220.
- [39] C. Ma, J. Wang, et al., Model reduction with memory and the machine learning of dynamical systems, *arXiv preprint arXiv:1808.04258*, 2018.
- [40] D.J. Evans, *Statistical Mechanics of Nonequilibrium Liquids*, 2008.
- [41] G. Mengaldo, R. Maulik, PySPOD: a python package for spectral proper orthogonal decomposition (SPOD), *J. Open Sour. Softw.* 6 (2021) 2862.
- [42] G.A. Brès, S.K. Lele, Modelling of jet noise: a perspective from large-eddy simulations, *Philos. Trans. R. Soc. A, Math. Phys. Eng. Sci.* 377 (2019).
- [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [44] C. Zhang, Madden-Julian oscillation, *Rev. Geophys.* 43 (2005).