



Learning Proper Orthogonal Decomposition of Complex Dynamics Using Heavy-ball Neural ODEs

Justin Baker¹ · Elena Cherkaev¹ · Akil Narayan² · Bao Wang²

Received: 24 February 2022 / Revised: 17 August 2022 / Accepted: 8 March 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

Proper orthogonal decomposition (POD) allows reduced-order modeling of complex dynamical systems at a substantial level, while maintaining a high degree of accuracy in modeling the underlying dynamical systems. Advances in machine learning algorithms enable learning POD-based dynamics from data and making accurate and fast predictions of dynamical systems. This paper extends the recently proposed heavy-ball neural ODEs (HBNODEs) (Xia et al. NeurIPS, 2021] for learning data-driven reduced-order models (ROMs) in the POD context, in particular, for learning dynamics of time-varying coefficients generated by the POD analysis on training snapshots constructed by solving full-order models. HBNODE enjoys several practical advantages for learning POD-based ROMs with theoretical guarantees, including 1) HBNODE can learn long-range dependencies effectively from sequential observations, which is crucial for learning intrinsic patterns from sequential data, and 2) HBNODE is computationally efficient in both training and testing. We compare HBNODE with other popular ROMs on several complex dynamical systems, including the von Kármán Street flow, the Kurganov-Petrova-Popov equation, and the one-dimensional Euler equations for fluids modeling.

Keywords Neural ODE · Momentum · Reduced-order modeling · Deep learning

Mathematics Subject Classification 65P99 · 68T07

✉ Bao Wang
bwang@math.utah.edu

Justin Baker
baker@math.utah.edu

Elena Cherkaev
elena@math.utah.edu

Akil Narayan
akil@sci.utah.edu

¹ Department of Mathematics, University of Utah, Salt Lake City, UT, USA

² Department of Mathematics, and Scientific Computing and Imaging (SCI) Institute, University of Utah, Salt Lake City, UT, USA

1 Introduction

Numerical long-time simulation of full-order models (FOMs) of complex dynamical systems is computationally costly. This is particularly true for physical systems that contain a wide range of spatial and temporal scales, including direct numerical simulation (DNS) [37] or large eddy simulation (LES) in fluid mechanics [12, 17, 60] and chaotic systems [27, 52, 53]. Reduced-order models (ROMs) have been utilized as alternative scientific simulation tools, which are computationally much more efficient than FOMs and retain comparable accuracy for simulating complex dynamical systems. ROMs have played crucial roles in designing, optimizing, and controlling dynamical systems [1, 2, 5, 18].

Several data-driven numerical algorithms have been proposed for reduced-order modeling, including dynamic mode decomposition (DMD) [51] and proper orthogonal decomposition (POD) [6]. These models leverage some FOM simulation data to construct low-dimensional simplified models that describe the underlying dynamics, with the goal of using these simplified models in generalization regimes to predict the unseen dynamics. Classical projection-based reduced-order modeling techniques (of which DMD and POD are examples) are among the most popular approaches for constructing ROMs of dynamical systems. This approach transforms the simulation results of FOM into a suitable low-dimensional subspace that preserves the largest variance of the training data. In, e.g., POD, classical numerical algorithms (such as Galerkin methods), are subsequently used to rewrite the state variable in the governing equation of the underlying dynamics into a system of ODEs, resulting in a substantially reduced degree of freedom for describing the complex dynamics. Both DMD and POD have been widely used in scientific simulations, particularly for fluid simulations.

ROMs generated from projection-based approaches can preserve crucial physical structures of the dynamics system. However, inappropriate truncation of the POD modes in governing equations can severely degrade modeling accuracy and result in unexpected, unphysical predictive results. Moreover, the precise strategy for mode truncation is task-dependent and is typically limited to explicit and closed definitions of the mathematical models [50]. Another drawback of direct projection-based approaches is that they require knowledge of governing equations that model the dynamical system, and this information is often absent for real-world problems. As such, data-driven reduced-order modeling has drawn significant recent attention. For instance, the learning of closure models to compensate for information loss due to mode truncation [36, 38, 48, 49], and data-driven reduced basis representations have been learned from simulation data that provides significantly improved predictive performance of the dynamics compared to classical models [31, 39, 58]. More recently, “vanilla” versions of machine learning approaches such as neural ODEs (NODEs) and recurrent neural networks (RNNs) have been used to learn temporal coefficients of the POD of a given complex dynamical system [24, 25, 45]. A well-known issue of the vanilla NODEs and RNNs is that they lack the capability of capturing long-range dependencies from data, making these machine learning models fail to learn the intrinsic physics of complex systems [4, 59].

1.1 Our Contribution

We employ the recently developed *heavy-ball neural ODE* (HBNODE) [59], an extension of NODE [9], to learn the temporal coefficients of the POD of complex physical systems with a focus on time-dependent simulations in scientific computing. In particular, our examples include the von Kármán Street (VKS) flow, the Kurganov-Petrova-Popov (KPP) equation,

and the one-dimensional Euler equations for fluids modeling. We provide numerical validation on the above three benchmark tasks and a detailed empirical analysis of why HBNODEs are beneficial for learning the dynamics of POD modes. Our numerical results show the adjoint state of HBNODEs does not vanish, confirming that HBNODEs do learn long-range dependencies, which results in remarkable performance gain over the baseline NODEs. Moreover, our experimental results show significant computational advantages in training and testing HBNODEs over the baseline ROM models.

1.2 Related Work

There is a healthy amount of recent work on learning POD mode dynamics using deep neural networks, particularly RNNs and vanilla NODEs. Perhaps the most related papers to this article are [15, 16, 45], which study the NODE framework for learning ROMs. In [45], the authors developed a POD-NODE ROM framework for learning POD coefficients, which starts from FOM snapshots and then uses an autoencoder to encode the POD representations of FOM snapshots, followed by NODE training and forecasting. The POD-NODE ROM framework achieves appealing results for learning reduced dynamics of the VKS model, and it significantly outperforms the direct application of a long short-term memory (LSTM) network for sequential learning. In [15, 16], the authors study the effectiveness of NODEs for reduced-order modeling and predicting environment hydrodynamics. On the one hand, they find that NODEs provide an elegant framework for the stable and accurate evolution of latent-space dynamics with promising generalizability. On the other hand, they noticed that in order to facilitate the widespread adoption of NODEs for large-scale systems, significant effort needs to be directed at accelerating training time. This limitation motivates this article's study and utilization of HBNODEs [59], which is the machine-learning backbone of the reduced-order modeling pipeline proposed in this work. There are three major advantages of learning PODs using HBNODEs over the existing deep learning approaches:

- HBNODEs are a class of continuous-depth neural networks, and they are suitable for learning irregularly-sampled simulation data or physical observations. Hence, observation protocols that entail missing or sparse data are easily tackled in this framework.
- Certain spectral properties of HBNODEs enable them to capture long-range dependencies from sequential data, which is crucial for learning PODs of complex dynamics.
- Both HBNODEs and their adjoint ODEs are computationally much more efficient than baseline NODEs.

In addition to the NODE paradigm of continuous-depth neural networks for reduced-order modeling, the RNN—a natural sequential deep learning model—has also been successfully used for learning-assisted model reduction. Many advanced RNN algorithms can also be leveraged to enhance learning ROMs, e.g., LSTM networks [21]. RNN-based ROMs have achieved remarkable success in many applied domains, including multiphase flow simulation [24, 25], learning advection-dominated systems [35], learning chaotic dynamics [32], and learning nonlinear aeroelastic models [33]. Compared to NODEs for learning ROMs, RNNs cannot learn irregularly-sampled time series effectively and can even depart from the underpinning physics due to their discrete nature.

1.3 Organization

We organize the paper as follows: In Sects. 2 and 3, we briefly review the POD-based reduced-order modeling and HBNODE for continuous-depth deep learning, respectively. We present the benchmark physical models of the complex dynamical systems and full-order modeling for data generation in Sect. 4. Section 5 shows the detailed deep learning model and pipeline for learning POD-based ROMs. We verify the efficacy of our proposed machine learning models and contrast them with several baseline models in Sect. 6, followed by concluding remarks.

2 POD-Based Reduced-Order Modeling

In this section, we briefly review key ideas and procedures of POD-based reduced-order modeling.

2.1 Notation

We denote vectors and matrices by lower- and upper-case boldface letters, respectively. For a vector $\mathbf{x} = (x_1, \dots, x_d)^\top \in \mathbb{R}^d$, where $(x_1, \dots, x_d)^\top$ denotes the transpose of the row vector (x_1, \dots, x_d) , we use $\|\mathbf{x}\| = (\sum_{i=1}^d x_i^2)^{1/2}$ to denote its ℓ_2 norm, and use $\mathbf{0}$ to denote the zero vector. In cases when $d = 2$ and \mathbf{x} is a spatial vector, we will write the components instead as $\mathbf{x} = (x, y)^\top$. For a matrix \mathbf{A} , we use \mathbf{A}^\top , \mathbf{A}^{-1} , and $\|\mathbf{A}\|$ to denote its transpose, inverse, and spectral norm, respectively. We use \mathbf{I} to denote the identity matrix, whose size will be clear based on context.

We will consider the approximation of a space-time function $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ where \mathbf{x} is a spatial vector (typically of 1 or 2 dimensions) and t is a scalar on $[0, T]$ for some fixed and finite terminal time T . The function \mathbf{u} may be vector-valued. In much of our discussion, we will take the concrete example of \mathbf{u} being a solution to a discretized VKS problem, whose details are given in Sect. 4.1. For the VKS problem, $\mathbf{u} \in \mathbb{R}^2$ contains the horizontal (x) and vertical (y) components of a fluid velocity field. We will write $\mathbf{u} = (u_x, u_y)$ to denote these two components.

2.2 POD Snapshots

POD shares a similar spirit and implementation as the celebrated principal component analysis (PCA), the latter of which has been a very popular tool for data analysis [30, 43]. The key idea of PCA is to project high-dimensional data into a lower-dimensional space that is spanned by the eigenvectors corresponding to the leading eigenvalues of the covariance matrix of the data. PCA preserves the largest variance of training data and thus contains the most important information contained in the originally high-dimensional data. POD has been introduced in accelerating fluids simulation and reducing the complexity of fluid models since the pioneering work of Berkooz et al. [7]. Once the leading eigenmodes are obtained via analysis of training data, it is possible to reduce the order (the computational complexity and degrees of freedom) of the complex FOMs. The POD-based dimension reduction approach starts with some training samples of physically observed or numerically simulated snapshots of dynamics. These sample snapshots are aggregated into an ensemble matrix \mathbf{Y} , where each row contains the state of a dynamical system at a fixed time step. Next, we




Fig. 1 POD pipeline: We first pre-process the data, from experimental observation or FOM simulation, by subtracting the mean. Then we apply spectral decomposition of the covariance matrix and only keep the first r eigenmodes, resulting in the reduced representation— $\sum_{i=1}^r \alpha_i(t) \psi_i(x)$ —of $u'(x, t)$, where $\alpha_i(t)$ s are the temporal coefficients and $\psi_i(x)$ s are the eigenmodes

compute the covariance matrix of the rows of the matrix Y , and the eigenvectors—sorted according to the corresponding decreasing-ordered eigenvalues—are used as the new orthogonal basis for representing the ROM. Below we summarize the crucial steps of identifying the low-dimensional representations via the POD approach, which has also been visualized in Fig. 1.

- *Data generation* We simulate the FOM, which is computationally expensive, for a short time to obtain the training data at time steps t_1, t_2, \dots, t_{N_t} . For the VKS problem, when our solution u contains two components $u = (u_x, u_y)$ that depend on the two-dimensional spatial variable $x = (x, y)$, we assume that FOM snapshots $u_x(t_j), u_y(t_j)$ are vectorized representations of N spatial degrees of freedom. Then we have

$$u_x(t_j), u_y(t_j) \in \mathbb{R}^N, \quad j = 1, 2, \dots, N_t.$$

For real-world dynamical systems for which we do not know the exact governing equation, we sample the true dynamics via experimental measurements as training data. In either case, we assume that training data is available to us (as the snapshots above), and our goal is to efficiently leverage this data to learn reduced-order dynamics without recourse to the FOM, which we assume is unknown.

- *Linearly center the data dynamics* With our VKS-centric notation above, according to the Reynolds decomposition of the flow, we have for fixed t_j ,

$$u_x = \bar{u}_x + u'_x; \quad u_y = \bar{u}_y + u'_y, \quad (1)$$

where \bar{u}_x and \bar{u}_y are the temporal mean of the solutions, computed over our N_t snapshots. The components u'_x and u'_y are the *fluctuating components* of the data.

- *Data assembling* Concatenate the simulated and centered FOM snapshots into the following matrix Y ,

$$Y = \begin{pmatrix} - & (u'_x(t_1))^T & - & - & (u'_y(t_1))^T & - \\ - & (u'_x(t_2))^T & - & - & (u'_y(t_2))^T & - \\ & \vdots & & & \vdots & \\ - & (u'_x(t_{N_t}))^T & - & - & (u'_y(t_{N_t}))^T & - \end{pmatrix}, \quad (2)$$

so that row j contains the concatenated snapshot $u'_x(t_j), u'_y(t_j)$, i.e., the two flattened velocity components at time step j . The size of the matrix Y is $N_t \times 2N$.

- *Perform a spectral decomposition of the covariance matrix* We construct the covariance matrix \mathbf{K} of the rows of \mathbf{Y} and compute its eigendecomposition:

$$\mathbf{K} = \mathbf{Y}\mathbf{Y}^\top, \quad \mathbf{K} = \mathbf{A}\mathbf{\Lambda}\mathbf{A}^\top, \quad \mathbf{A} = (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_{N_t}), \quad (3)$$

where $\boldsymbol{\alpha}_j$ is the j th eigenvector, and the matrix \mathbf{A} is diagonal containing entries λ_j , the associated non-negative eigenvalues of \mathbf{K} . We assume the eigenvalues are listed in non-increasing order, $\lambda_j \geq \lambda_{j+1}$.

- *Identify reduced-order modes and truncate* Larger eigenvalues of \mathbf{K} are directly related to the dominant characteristics of the dynamical system, while small eigenvalues correspond to small perturbations of the dynamical behavior. The matrix \mathbf{K} has N_t eigenvalues, and we choose the *order* of the reduced-order model to be $r \ll N_t$ by inspecting a relative information content $I(r)$, defined as follows

$$I(r) = \frac{\sum_{i=1}^r \lambda_i}{\sum_{i=1}^{N_t} \lambda_i}, \quad (4)$$

so that $1 - I(r)$ is a relative Frobenius norm error between \mathbf{K} and its rank- r spectral approximation. As we will see in Sect. 6, $I(r)$ is close to one for practical problems, even for very small r . As output of the procedure, we can construct the following (discretized) ROM of the fluctuating component of the dynamics

$$\mathbf{u}'_*(t_j) \approx \sum_{i=1}^r (\alpha_i)_j \boldsymbol{\psi}_{*,i}, \quad (5)$$

where $*$ $\in \{x, y\}$, and $(\boldsymbol{\psi}_{x,i}^\top, \boldsymbol{\psi}_{y,i}^\top)^\top \in \mathbb{R}^{2N}$ is a vector denoting a discretized spatial function; the entries of $\boldsymbol{\psi}_{*,i}$ correspond to the N degrees of freedom in the snapshots \mathbf{u}_* , and is a subvector of the i -th right-singular vector of \mathbf{Y} . Equivalently, it is defined as,

$$\boldsymbol{\psi}_{*,i} = \frac{1}{\lambda_i} \mathbf{Y}^\top \boldsymbol{\alpha}_i = \frac{1}{\lambda_i} \sum_{j=1}^{N_t} (\alpha_i)_j \mathbf{u}'_*(t_j).$$

A “standard” POD approach would next project the (assumed known) dynamical model onto $\text{span}\{\boldsymbol{\psi}_{*,i}\}_{i=1}^r$. We will proceed to assume that such a dynamical model is unknown to us, and will instead use machine learning models to predict dynamics.

Remark 1 With the training data α_i available at time steps t_1, t_2, \dots, t_{N_t} through the above procedure, extrapolation of the FOM dynamics or experimental measurements amounts to predicting the POD coefficients $\alpha_i(t)$ for future time t accurately, in our case using machine learning models. Notice that $\alpha_i(t)$ is observed sequentially and has a continuous profile, indicating the potential advantages of using NODE for learning $\alpha_i(t)$, as we describe next.

3 Heavy-Ball Neural ODEs

In this section, we briefly review NODE and HBNODE and algorithms for their training and testing. Moreover, we provide some simple analysis of why HBNODE is better for learning POD coefficients for reduced-order modeling leveraging insights from the acceleration theory of the classical momentum methods.

3.1 Neural ODEs

NODEs [9] are a class of continuous-depth (-time) neural networks [11, 46]. The continuous-time nature of NODEs makes them particularly suitable for learning complex dynamics from irregularly-sampled sequential data, see, e.g., [9, 14, 34, 41, 47]. Mathematically, a NODE is formulated as the following first-order ODE:

$$\frac{d\mathbf{h}(t)}{dt} = \mathbf{f}(\mathbf{h}(t), t, \theta), \quad (6)$$

where $\mathbf{f}(\mathbf{h}(t), t, \theta) \in \mathbb{R}^d$ is specified by a neural network parameterized by θ , e.g., a two-layer feed-forward neural network. Starting from the input $\mathbf{h}(0)$, NODEs learn the representation and perform prediction by solving (6) from $t = 0$ to T using a numerical integrator with a given error tolerance, often with adaptive step size solver or adaptive solver for short [13]. Solving (6) from $t = 0$ to T in a single pass with an adaptive solver requires evaluating $\mathbf{f}(\mathbf{h}(t), t, \theta)$ at various timestamps, with computational complexity measured by the number of function evaluations in a time-forward sweep (“forward NFEs”) [9].

The *adjoint sensitivity method*, or the adjoint method [8], is a memory-efficient method for training NODEs through optimization of θ . We regard the output $\mathbf{h}(T)$ as the prediction and denote the loss between the prediction $\mathbf{h}(T)$ and the ground truth as \mathcal{L} . Let $\mathbf{a}(t) := \partial\mathcal{L}/\partial\mathbf{h}(t)$ be the adjoint state, then we have (see [8, 9] for details)

$$\frac{d\mathcal{L}}{d\theta} = \int_0^T \mathbf{a}(t)^\top \frac{\partial \mathbf{f}(\mathbf{h}(t), t, \theta)}{\partial \theta} dt, \quad (7)$$

with $\mathbf{a}(t)$ satisfying the following adjoint ODE

$$\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^\top \frac{\partial}{\partial \mathbf{h}} \mathbf{f}(\mathbf{h}(t), t, \theta), \quad (8)$$

which is solved numerically from $t = T$ to 0 and also requires the evaluation of the right-hand side of (8) at various timestamps, with the number of NFEs during this time-backward sweep (“backward NFEs”) measuring the computational complexity.

There are several critical problems with NODEs, including (i) Given an error tolerance, the NFEs required in a single forward pass can be excessive. Moreover, solving the adjoint ODE (8) often requires more NFEs than solving the forward ODE (6). (ii) In training NODEs, the adjoint state $\mathbf{a}(t)$ often vanishes, i.e., the norm of $\mathbf{a}(t)$ tends to 0, impeding NODEs from learning long-range dependencies [29], resulting in poor predictive performance.

3.2 Heavy-Ball Neural ODEs

The authors of [56, 59] proposed HBNODEs and their generalized version, named generalized HBNODEs (GHBNODEs). HBNODEs are motivated by ideas from momentum-accelerated gradient descent [44] and they can be regarded as the continuous limit of the MomentumRNN model [40]. Mathematically, the HBNODE is a special second-order neural ODE of the following form

$$\frac{d^2\mathbf{h}(t)}{dt^2} + \gamma \frac{d\mathbf{h}(t)}{dt} = \mathbf{f}(\mathbf{h}(t), t, \theta), \quad (9)$$

where $\gamma \geq 0$ is the damping parameter, which can be set as a tunable or a learnable hyperparameter with positivity constraint. In the trainable case, we adopt the one used in [59], that

is $\gamma = \epsilon \cdot \text{sigmoid}(\omega)$ for a trainable $\omega \in \mathbb{R}$ and a fixed tunable upper bound ϵ , e.g., $\epsilon = 1$. The HBNODE (9) can be rewritten as the following system of first-order NODEs

$$\frac{d\mathbf{h}(t)}{dt} = \mathbf{m}(t); \quad \frac{d\mathbf{m}(t)}{dt} = -\gamma\mathbf{m}(t) + \mathbf{f}(\mathbf{h}(t), t, \theta). \quad (10)$$

3.2.1 Computational advantages of HBNODE vs. NODE

To show why HBNODE enjoys computational efficiency in training and testing, let us first consider the adjoint equation of (9), which will again be solved using adaptive numerical ODE solvers. First, the following theoretical result [59] shows that the adjoint of an HBNODE is also an HBNODE.

Proposition 1 (Adjoint equation for HBNODE [59]) *The adjoint state $\mathbf{a}(t) := \partial\mathcal{L}/\partial\mathbf{h}(t)$ for the HBNODE (9) satisfies the following HBNODE with the same damping parameter γ as that in (9),*

$$\frac{d^2\mathbf{a}(t)}{dt^2} - \gamma \frac{d\mathbf{a}(t)}{dt} = \mathbf{a}(t) \frac{\partial \mathbf{f}}{\partial \mathbf{h}}(\mathbf{h}(t), t, \theta). \quad (11)$$

Notice that we solve the adjoint equation (11) from $t = T$ to 0 via backward propagation. By letting $\tau = T - t$ and $\mathbf{b}(\tau) = \mathbf{a}(T - \tau)$, we can rewrite (11) as follows,

$$\frac{d^2\mathbf{b}(\tau)}{d\tau^2} + \gamma \frac{d\mathbf{b}(\tau)}{d\tau} = \mathbf{b}(\tau) \frac{\partial \mathbf{f}}{\partial \mathbf{h}}(\mathbf{h}(T - \tau), T - \tau, \theta). \quad (12)$$

Therefore, the adjoint of the HBNODE is also an HBNODE and they have the same damping parameter.

The above result indicates that the adjoint problem for HBNODE is of the same type as the forward problem, accelerating backward propagation provided the forward propagation is accelerated.

Next, we provide theoretical insights into the computational efficiency of HBNODEs. Leveraging the acceleration theory of the heavy-ball momentum method in taming the oscillation of the optimization trajectory by reducing the condition number of the underlying problem. Furthermore, it is worth mentioning that the heavy-ball method only provably accelerates strongly convex quadratic optimization problems, while in practice, it accelerates general optimization problems. In analogy to the acceleration theory of the heavy-ball method, we consider the following two linearized high-dimensional ODE systems

$$\frac{d\mathbf{h}(t)}{dt} = \mathbf{A}\mathbf{h}(t), \quad (13)$$

and

$$\begin{aligned} \frac{d\mathbf{h}(t)}{dt} &= \mathbf{m}(t) \\ \frac{d\mathbf{m}(t)}{dt} &= -\gamma\mathbf{m}(t) + \mathbf{A}\mathbf{h}(t) \end{aligned} \quad \Leftrightarrow \quad \frac{d}{dt} \begin{pmatrix} \mathbf{h}(t) \\ \mathbf{m}(t) \end{pmatrix} = \underbrace{\begin{pmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{A} & (-\gamma\mathbf{I}) \end{pmatrix}}_{:=\mathbf{B}} \begin{pmatrix} \mathbf{h}(t) \\ \mathbf{m}(t) \end{pmatrix}, \quad (14)$$

where we assume \mathbf{A} is positive definite to simplify our analysis and reveal intuition of the advantages of HBNODE over NODE. Let the eigenvalues and eigenvectors of \mathbf{A} be given by λ_i , \mathbf{v}_i respectively. Following the proof of the acceleration of heavy-ball momentum,¹ we can show that $|\tilde{\lambda}_{\max}|/|\tilde{\lambda}_{\min}| \leq \sqrt{|\lambda_{\max}|/|\lambda_{\min}|}$, where λ_{\max} and λ_{\min} are the largest and

¹ see <http://www.math.utah.edu/~bwang/mathds/Lecture8.pdf> for details.

smallest eigenvalues. Similarly, $\tilde{\lambda}_{\max}$ and $\tilde{\lambda}_{\min}$ are the largest and smallest eigenvalues, in magnitude, of \mathbf{B} .

Notice that the ratio $|\lambda_{\max}|/|\lambda_{\min}|$ and $|\tilde{\lambda}_{\max}|/|\tilde{\lambda}_{\min}|$ are the stiffness of the linear ODE model (13) and the corresponding linear HBNODE counterpart (14). Thus the heavy-ball NODE can be much less stiff than the original NODE. If the stiffness of the ODE model is κ , using the heavy-ball model results in stiffness of at most $\sqrt{\kappa}$, which is a substantial reduction.

Recall that we use the adaptive step size explicit solver to solve both forward and backward ODEs, from $t = 0$ to T , in training NODEs and HBNODEs. A less stiff model allows the adaptive solver to use a much large step size and thus can significantly reduce NFEs. Moreover, Proposition 1 indicates that the adjoint equation of an HBNODE is also an HBNODE, and therefore we reap the computational advantages of relaxed stiffness in both forward and backward propagation phases. Our previous analysis only considers very simple linear ODE models. How to extend the analysis to the neural network is a very interesting future direction. One particular idea is analyzing the NODE and HBNODE when they are overparameterized, in which case one could leverage neural tangent kernel theory [23].

3.2.2 Generalized HBNODEs (GHBNODEs)

Compared to vanilla NODEs, high-order NODEs include HBNODEs usually suffer from the uncontrolled aggregation of the hidden state, deteriorating model performance at best, and blowing up training at worst. To alleviate this issue, in [59] the authors propose the following generalized HBNODE

$$\frac{d\mathbf{h}(t)}{dt} = \sigma(\mathbf{m}(t)); \quad \frac{d\mathbf{m}(t)}{dt} = -\gamma\mathbf{m}(t) + \mathbf{f}(\mathbf{h}(t), t, \theta) - \xi\mathbf{h}(t), \quad (15)$$

where $\sigma(\cdot)$ is a nonlinear activation, which is set as \tanh by default. The positive hyperparameters $\gamma, \xi > 0$ are two tunable or learnable hyperparameters. In the trainable case, we let $\gamma = \epsilon \cdot \text{sigmoid}(\omega)$ as in HBNODE, and $\xi = \text{softplus}(\chi)$ to ensure that $\gamma, \xi \geq 0$. Compared to HBNODEs, GHBNODEs integrate two ideas to improve the neural network architecture design: (i) Incorporating the gating mechanism σ used in LSTM [22] and GRU [10], which can suppress the aggregation of $\mathbf{m}(t)$; (ii) Following the idea of skip connections [20], HBNODEs add the term $\xi\mathbf{h}(t)$ into the governing equation of $\mathbf{m}(t)$, which benefits training and generalization of GHBNODEs. It has been extensively verified that GHBNODE can indeed control the growth of $\mathbf{h}(t)$ effectively, which significantly improve the performance of machine learning models on various sequential learning tasks.

Another interesting result is that though the adjoint state of the GHBNODE does not satisfy the exact heavy-ball ODE, it also significantly reduces the backward NFEs in practice. We observe that sometimes GHBNODEs are computationally more efficient than HBNODEs.

3.2.3 (G)HBNODEs learn long-range dependencies effectively

Learning long-range dependencies is crucial for the success of deep learning for sequential data, and vanishing and exploding gradients are two bottlenecks for training RNNs to learn long-range dependencies [4, 42]. The exploding gradients issue can be effectively resolved via gradient clipping, training loss regularization, etc [42]. The vanishing gradient phenomenon in training RNNs materializes in continuous-depth neural networks as vanishing of the adjoint state [59]. In particular, we consider $\mathbf{a}(t) := \partial\mathcal{L}/\partial\mathbf{h}(t)$, and when the vanishing gradient

phenomenon occurs, $\mathbf{a}(t)$ goes to $\mathbf{0}$ quickly as $T - t$ increases, so that $d\mathcal{L}/d\theta$ in (7) will be essentially independent of $\mathbf{a}(t)$ for larger $T - t$. We have the following expressions for the adjoint states of the NODE and HBNODE (see [59] for details):

– For NODE, we have

$$\frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \exp \left\{ - \int_T^t \frac{\partial \mathbf{f}}{\partial \mathbf{h}}(\mathbf{h}(s), s, \theta) ds \right\}. \quad (16)$$

– For GHBNODE,² we have

$$\begin{aligned} \left[\frac{\partial \mathcal{L}}{\partial \mathbf{h}} \frac{\partial \mathcal{L}}{\partial \mathbf{m}_t} \right] &= \left[\frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \begin{bmatrix} \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_t} \frac{\partial \mathbf{h}_T}{\partial \mathbf{m}_t} \\ \frac{\partial \mathbf{m}_T}{\partial \mathbf{h}_t} \frac{\partial \mathbf{m}_T}{\partial \mathbf{m}_t} \end{bmatrix} \\ &= \left[\frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \exp \left\{ - \underbrace{\int_T^t \left[\begin{pmatrix} \mathbf{0} & \frac{\partial \sigma}{\partial \mathbf{m}} \\ \left(\frac{\partial \mathbf{f}}{\partial \mathbf{h}} - \xi \mathbf{I} \right) - \gamma \mathbf{I} \end{pmatrix} ds \right]}_{:= \mathbf{M}} \right\}. \end{aligned} \quad (17)$$

For the matrix \mathbf{M} , we have the following useful property about its spectrum.

Proposition 2 ([59]) *The eigenvalues of $-\mathbf{M}$ can be paired so that the sum of each pair equals $(t - T)\gamma$.*

Following the argument in [59], Proposition 2 can be used to show that the adjoint state of NODE in (16) may vanish when $T - t$ is large, but the adjoint state of (G)HBNODEs in (17) will not vanish. This property supports the claim that HBNODEs benefit in learning long-range dependencies, which in turn further boosts the accuracy in learning POD of complex dynamical systems.

In Sect. 6, we will validate the above theoretical merits of HBNODEs over NODEs using the benchmark problems listed in Sect. 4 below.

4 Benchmarks and Data Preparation

In this section, we will present some details of the three benchmark physical models—VKS, KPP, and Euler equations—used for validating the efficacy of learning POD with HBNODEs. Our training data for the KPP and Euler equations are generated by solving FOMs; the training data for the VKS dataset is adopted from a publicly available dataset.

4.1 VKS Model

The von Kármán vortex street (VKS) is a fluid dynamics phenomenon where vortices appear in a periodic fashion in the wake of flow past a blunt object, frequently a cylinder. A very small Reynolds number results in a laminar smooth flow past the cylinder, and very large Reynolds numbers result in a turbulent flow. In an appropriate middle regime, the VKS phenomenon appears and can be simulated. The associated dynamical model is the two-dimensional Navier-Stokes equations; with $\mathbf{u} = (u_x, u_y)$ the fluid velocity, these equations read,

$$\frac{\partial \mathbf{u}}{\partial t} = \nu \nabla^2 \mathbf{u} - (\mathbf{u} \cdot \nabla) \mathbf{u} - \frac{1}{\rho_0} \nabla p,$$

² HBNODE can be seen as a special GHBNODE with $\xi = 0$ and σ be the identity map.

where ρ_0 is the spatially uniform pressure, and ν is the kinematic viscosity, which is inversely related to the Reynolds number. Our experimental setup concerns flow past a cylinder in two spatial dimensions with conditions that result in steady-state VKS flow after an initial transient period. We follow the experimental setting used in [45] to acquire simulation data.

4.2 KPP Model

The Kurganov-Petrova-Popov (KPP) model is a scalar, two-dimensional conservation law, first proposed in [28]. This system is difficult to simulate since it features a non-convex flux, and is given by,

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathbf{f}(u) = 0, \quad t > 0, \quad x \in [-2, 2], \quad y \in \left[-\frac{5}{2}, \frac{3}{2}\right],$$

where $\mathbf{f}(u) = (\sin u, \cos u)^T$ and $\nabla := \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}\right)$. Our setup mirrors that in [28], so that we use the following initial data

$$u(x, y, 0) = \begin{cases} \frac{14\pi}{4}, & x^2 + y^2 < 1, \\ \frac{\pi}{4}, & \text{else} \end{cases}$$

We employ a finite volume scheme utilizing a Lax-Friedrichs flux with a 5th-order WENO reconstruction over the two-dimensional rectangular domain with a Cartesian mesh up to time $T = 10$. The simulation uses a tensorial grid with $N_x = 50$, $N_y = 50$ (corresponding to $N = N_x N_y = 2500$ total spatial degrees of freedom), and $N_t = 1250$.

4.3 Euler Equations for Fluids Modeling

The one-dimensional Euler equations of gas dynamics are a system of conservation laws. We consider the simulation of a *parameterized* shock-entropy problem from this differential equation, whose setup is given by,

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{u})}{\partial x} = 0, \quad t > 0, \quad x \in [-5, 5],$$

$$\text{with } \mathbf{f}(\mathbf{u}) = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ (E + p)u \end{pmatrix}^T,$$

where $\mathbf{u} := (\rho \ \rho u \ E)^T \in \mathbb{R}^3$ is the unknown with (ρ, u, p, E) denoting the gas density, velocity, pressure, and energy, respectively. The system is closed via the following relationship between E and p :

$$p = (\gamma - 1) \left(E - \frac{1}{2} \rho u^2 \right),$$

where γ is the heat capacity ratio, a gas-dependent constant.³ We take boundary conditions at $x = \pm 5$ as those given by the initial data. The shock-entropy problem features smoothly oscillating as well as discontinuous features.

We again employ a finite volume scheme to solve the Euler equations, using a Harten-Lax-van Leer (HLL) flux, which is an approximated Riemann solver [19]. Our simulations

³ This γ is distinct from the γ discussed in Sect. 3.2.

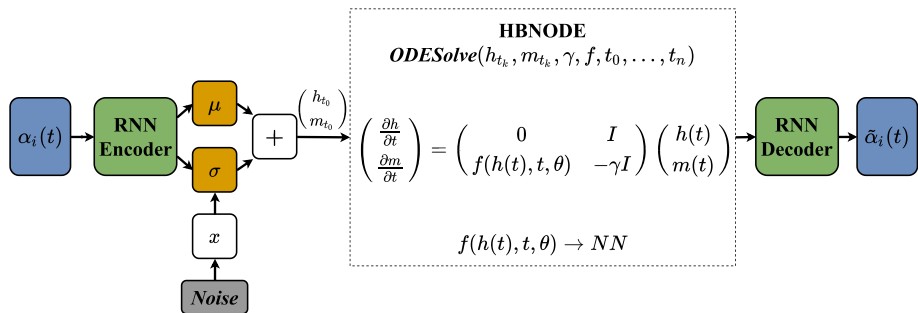


Fig. 2 The pipeline of predicting the temporal coefficients for a single step forward using HBNODE leveraging a variational autoencoder. We first use a RNN encoder to encode the input data and then sample the states \mathbf{h} and \mathbf{m} and evolve them using an HBNODE. Finally, we apply an RNN decoder to the final representation to get the prediction

integrate up to terminal time $T = 1.8$, with a uniform grid having $N = 1000$ degrees of freedom in the scalar spatial variable x .

This last example differs from the previous two in that we consider this a *parametric* equation, where $\boldsymbol{\eta} = (\eta_u, \eta_\rho) \in \mathbb{R}^2$ is a parameter for the initial conditions. We initialize the dynamics using the parameter $\boldsymbol{\eta}$ as follows, where η_u varies on the interval $[2, 3]$ and η_ρ varies on the interval from $[3, 4]$. The parametric initial data $(u(x, 0), \rho(x, 0), p(x, 0)) = (u_0, \rho_0, p_0)$ are given by

$$u_0 = \begin{cases} \eta_u & x < -4 \\ 0 & \text{else} \end{cases}, \quad \rho_0 = \begin{cases} \eta_\rho & x < -4 \\ 1 + 0.2 \sin(\pi x) & \text{else} \end{cases}, \quad p_0 = \begin{cases} \frac{31}{3} & x < -4 \\ 1 & \text{else} \end{cases}$$

We generate training data by gathering an ensemble of trajectories for the above problem over a grid of $\boldsymbol{\eta}$ values and attempt to learn dynamics on unseen values of $\boldsymbol{\eta}$. Thus, in this example we not only seek to predict to future times, but also trajectories on parameter values not in the training set.

5 Learning Pipeline

In this section, we describe the detailed pipeline of using deep learning for reduced-order modeling accompanied by the baseline ROMs.

5.1 Learning-Based Reduced-Order Modeling

Our machine learning-based reduced-order modeling framework is flexible for machine learning model selection, e.g., using either HBNODE or NODE as shown in Figs. 2 and 3, respectively. In our learning-based reduced-order modeling framework, we first apply POD outlined in Sect. 2 on the training data to extract (discretized) temporal coefficients $\boldsymbol{\alpha}(t)$'s and the eigenmodes $\boldsymbol{\psi}(\mathbf{x})$'s following (5). Next, we will use machine learning models to predict future dynamics $u(\mathbf{x}, t)$ leveraging these coefficients and modes. In particular, the main task is an extrapolation of the temporal coefficients $\boldsymbol{\alpha}(t)$'s using NODEs or HBNODEs.

To predict future values of the POD data, we consider two different machine learning architectures, shown in Figs. 2 and 3, respectively. The first architecture is a one-to-one




Fig. 3 The pipeline of predicting the temporal coefficients for multi-steps ahead. First, we encode the input sequence $\{\alpha_i(t_j)\}_{j=0}^{n-1}$ using a RNN encoder to obtain the latent sequence $\{z_i(t_j)\}_{j=0}^{n-1}$. Second, we use a NODE to learn a “good” representation of the input sequence by evolving the latent sequence $\{z_i(t_j)\}_{j=0}^{n-1}$. Third, we apply a RNN decoder to the “good” representation to get the final prediction. Notice that NODE can be replaced with (generalized) HBNODE, in which case we need to obtain another sequence of momentum states from the RNN encoder

architecture that predicts the value at t_{k+1} based on the data at t_k . The second architecture is a sequence-to-sequence architecture that uses sequence data points to predict the following sequence of data points. The overlap in the sequence prediction can be adjusted so that the predicted sequence is entirely new or that only one new data point is predicted.

The first architecture under our study is adapted from [45], which was originally used to compare the performance of NODE and LSTM in model reduction. We replace the vanilla NODE used in [45] with the HBNODE, and we depict the modified architecture in Fig. 2. Compared to the pipeline used in [45], after the RNN encoding of the temporal coefficients we have to sample both \mathbf{h} and \mathbf{m} to accommodate learning using HBNODE. In contrast, the vanilla NODE used in [45] only needs to sample the state \mathbf{h} . The above encoding and sampling procedure is accomplished via a variational autoencoder [26].

We plot the second architecture in Fig. 3, where the vanilla NODE can be replaced with (generalized) HBNODE. For the second architecture, i.e., the sequence-to-sequence architecture, takes a sequence of length n inputs and predicts a sequence of outputs, we encode the input sequence $\{\alpha_i(t_j)\}_{j=0}^{n-1}$ into the latent sequence $\{z_i(t_j)\}_{j=0}^{n-1}$ using an RNN encoder, then we use NODE or HBNODE to evolve the latent sequence to get the desired representation, followed by an RNN decoder to get the final long-term prediction $\{\alpha_i(t_j)\}_{j=n}^N$.

5.2 A Baseline Comparison: Dynamic Mode Decomposition (DMD)

We employ DMD as another baseline model reduction method to demonstrate the effectiveness of learning-based model reduction using HBNODEs. In this part, we briefly review the idea of DMD for reduced-order modeling. To compare DMD to the learning-based reduced-order modeling using HBNODE, we consider only modeling the fluctuating components \mathbf{u}' of the snapshots, see (1). The predictions of DMD are generated by a linear operator \mathbf{A} corresponding to a linear difference equation $\mathbf{u}'_{k+1} = \mathbf{A}\mathbf{u}'_k$, where \mathbf{A} must be learned. In DMD, dominant eigenvalues and eigenvectors of \mathbf{A} are computed via the singular value decomposition (SVD). Although the true underlying dynamics may be nonlinear, the Koopman operator formalism concludes that a *lifted* version of the dynamics is indeed linear. For nonlinear problems, DMD attempts to learn these lifted linear dynamics.

Let $\mathbf{U}^{(k)}$ be the snapshot matrix for the time interval $t_0, \dots, t_{\text{train}-1}$ and $\mathbf{U}^{(k+1)}$ be the snapshot matrix for the time interval $t_1, \dots, t_{\text{train}}$, i.e., column j of $\mathbf{U}^{(k)}$ corresponding to




Fig. 4 DMD pipeline: The data is pre-processed by subtracting the mean to capture the fluctuating components of the data. In addition, the following lifts $\{\cos(x), \sin(x), x^2, x^3\}$ were applied to the data, and then vectorized along the snap-shot axis. We then generate the full-order model according to the spectral decomposition of the linear transformation between the two snapshot matrices at subsequent time intervals. To reduce the order only the dominant r eigenmodes are selected, resulting in a representation $u'(x, t) = \sum_{i=1}^r \phi_i \lambda_i^{k-1} b_i$ for the lifted data $u'(x, t)$

time snapshot t_{j-1} . In particular, let $U^{(k+1)} \approx AU^{(k)}$ where $U^{(k)}$ is given by the SVD $U^{(k)} = X\Sigma V^*$. We further denote \tilde{X} , \tilde{V} , and $\tilde{\Sigma}$ as the rank- r truncation of X , V , and Σ , respectively. Then we may compute an approximation \tilde{A} directly from $U^{(k+1)}$ by the following,

$$\tilde{A} = \tilde{X}^* U^{(k+1)} \tilde{V} \tilde{\Sigma}^{-1} \quad (18)$$

The reduced matrix \tilde{A} is composed of the dominant r eigenvalues $\lambda_1, \dots, \lambda_r$ and eigenvectors ϕ_1, \dots, ϕ_r . These eigenvectors are also known as the DMD modes. Given training data on the training interval $t_0, \dots, t_{\text{train}}$, the matrix \tilde{A} is formulated by partitioning the snapshot matrix u' into two time intervals. Validation data on the interval $t_{\text{train}+1}, \dots, t_{\text{valid}}$ is generated by solving $u'(t_{\text{train}+k}) = \tilde{A}^k u'(t_{\text{train}})$. We depict DMD-based reduced-order modeling in Fig. 4. More details of DMD can be found at e.g., [51].

6 Experimental Results

In each experiment below, we contrast the performance of HBNODE-based ROM to two baseline ROMs, namely, NODE-based and DMD-based ROMs. We observe consistently improved predictive performance of HBNODE over baseline ROMs. We interpret the improved performance using HBNODEs by inspecting the stiffness and adjoint state of HBNODEs, confirming the theoretical results. Animated comparisons of the data reconstructions can be found at [3].

6.1 Transient and Steady-State VKS

The VKS dataset is obtained by simulating the FOM in Sect. 4.1 on the time interval $[0, 400]$, containing two different regimes. When $t < 100$, the dynamics lie in the transient state and approach the steady state as t increases; while the dynamics maintain a steady state when $t \geq 100$, as shown in Fig. 5.

ROMs for Steady-State Dynamics We contrast different ROMs for simulating VKS flow in the steady-state regime. In particular, both the DMD and POD training is taken over the time interval from $t = 100$ to 400. The POD modes for the steady-state flow oscillate quasi-periodically, see Fig. 6, and the relative information content $I(r)$ in (4) decays rapidly in r .

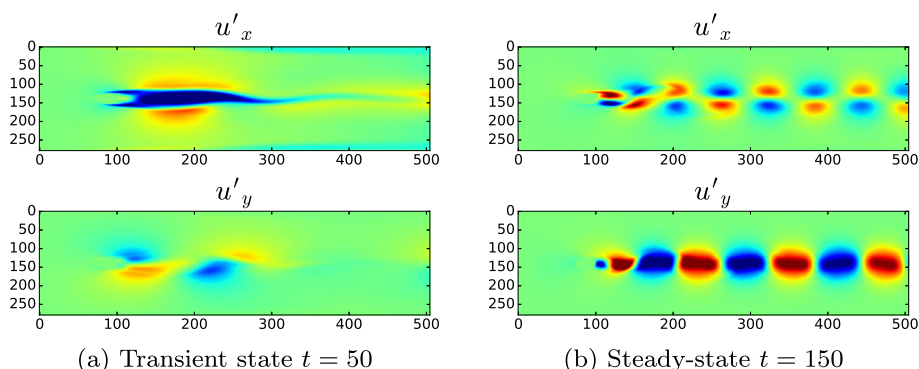


Fig. 5 Comparison of transient and steady-state phases of the VKS dataset. The steady-state phase contains quasi-periodic solutions conducive to machine learning. The transient phase does not contain such well-behaved dynamics




Fig. 6 The steady-state POD modes are highly oscillatory with quasi-periodic patterns

The POD relative information content for 8 leading modes is $\sim 99\%$, as illustrated in Fig. 7a. In contrast, the lifted DMD model, using the lifts $\{\cos(\mathbf{x}), \sin(\mathbf{x}), \mathbf{x}^2, \mathbf{x}^3\}$ with $\mathbf{x} = (x, y)$, requires 24 modes to achieve $\sim 99\%$ relative information content, shown in Fig. 7b. The quasi-periodic nature of the POD modes indicates that a model with high training accuracy will continue to perform well on the validation data.

ROMs for Transient to Steady State Dynamics ROMs behave very differently over the entire time interval from $t = 0$ to 400. The POD modes do not oscillate over the entire interval but only over the steady-state phase. For both POD and DMD, the relative information decays much slower. The POD relative information decreases to $\sim 96\%$ for the dominant 8 modes. While for the dominant 24 lifted DMD modes, the relative information is reduced to $\sim 94\%$ over the full dynamics. This suggests that a machine learning-based ROM which is able to train on the transient phase to predict the steady-state phase accurately captures the intrinsic patterns of the underlying dynamics.




Fig. 7 Comparison of the relative information decay for POD and DMD over the steady-state VKS. The POD modes decay far more rapidly than those of DMD. Therefore, one can expect better results with a smaller order of ROM using POD than DMD

Table 1 The hyperparameters for the VAE architecture—shown in Fig. 2—for NODE and HBNODE-based ROMs

Hyperparameter	Value
Latent dimension	6
Layers encoder	4
Units encoder	10
Layers ODE	12
Units decoder	41
Layers decoder	4
Learning rate	.00153
Epochs	2000


The parameters are tuned to the best NODE specification

Learning Steady-State Dynamics In this task, we train the pipeline shown in Fig. 2 for single-input-single-output dynamics prediction. Following the baseline in [45], we train over the steady-state dynamics starting from $t = 100$ using the dominant 8 POD modes. The training data consists of the POD modes from $t = 100$ to 174, and the training labels consist of the POD modes from $t = 101$ to 175. The validation data consists of the POD modes from $t = 175$ to 199, with the objective to predict the POD modes at time steps from $t = 176$ to 200. We use the mean squared error to measure the loss between the labeled data and the predictions. We utilized an AdamW optimizer to train the network based on this loss criteria. For the black-box integration method, we selected DOPRI-5 [13] with a relative tolerance of $1e-8$. The model's hyperparameters are tuned to best the NODE as outlined in [45] and restated in Table 1.

Learning Transient to Steady-State Dynamics In this task, we train the pipeline outlined in Fig. 3 for multi-input-single-output dynamics prediction. The objective of this task is to capture the phase transition at $t = 100$. The data consisted of the dominant 8 POD modes for the time interval from $t = 0$ to 400. The data was sequenced in a multi-input-single-output structure so that 9 preceding time steps were used to predict the 10-th time step. The training data consists of the POD modes for the transient time interval from $t = 0$ to 79. The training labels consisted of the POD modes from $t = 10$ to 80. The validation data utilizes the POD modes from steady-state time interval $t = 80$ to 119, and the validation labels consist of

Table 2 Hyperparameters of NODE and HBNODE for learning ROMs from transient to steady-state VKS dynamics

Hyper-parameter	Value
Layers	12
Hidden layers	64
Sequence length	9
Learning rate	.001
Epochs	500

**Fig. 8** Contrasting NODE and HBNODE-based ROMs for learning steady-state VKS dynamics. HBNODE outperforms NODE in both training and validation loss

data from $t = 90$ to 120. The other experimental settings follow the above single-input-single-output scenario. The model's hyperparameters are the same for NODE and HBNODE components and are given in Table 2.

6.1.1 Results and comparison to existing ROMs

Results of Learning Steady-State Dynamics We contrast HBNODE and NODE-based ROMs in Figs. 8 and 9. Figure 8 shows that HBNODE-based ROM not only achieves remarkably smaller training loss but also significantly smaller validation loss than NODE-based ROM. In terms of the predictive performance, we see that HBNODE performs better at capturing several of the peaks of the oscillatory modes as shown in Fig. 9.

Results of Learning Transient to Steady-State Dynamics Compared to learning steady-state VKS dynamics, HBNODE achieves more significant performance gain over NODE for learning transient to steady-state dynamics in terms of training and validation loss, as shown in Fig. 10. Since we are doing sequential learning, one interpretation of the improvement in learning dynamics is the effective learning of long-range dependencies. Indeed, the criterion of learning long-range dependencies has been widely used in measuring the efficacy of sequential learning models [42, 59]. In NODE and HBNODE, the effectiveness of learning long-range dependencies can be measured by whether the adjoint state vanishes quickly or not. We visualize the evolution of the magnitude of the adjoint states of NODE and HBNODE in Fig. 11, which support the theoretical result in Sect. 3.2.3. In particular, we see that the adjoint state of NODE vanishes much more rapidly than that of HBNODE as $T - t$ increases. A more detailed connection between the adjoint state and learning long-range dependencies is provided in [59].

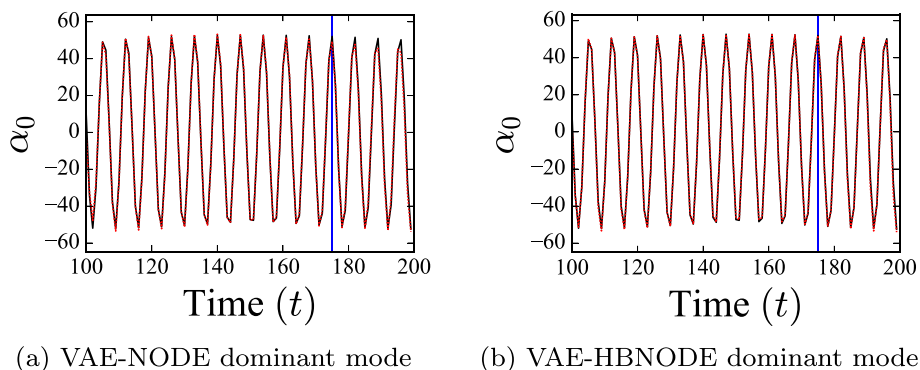


Fig. 9 Comparison of modes reconstruction for NODE and HBNODE-based ROMs for learning steady-state VKS dynamics. HBNODE captures the peaks of the dominant POD modes better than NODE. Before and after the vertical blue line stands for training and validation, respectively

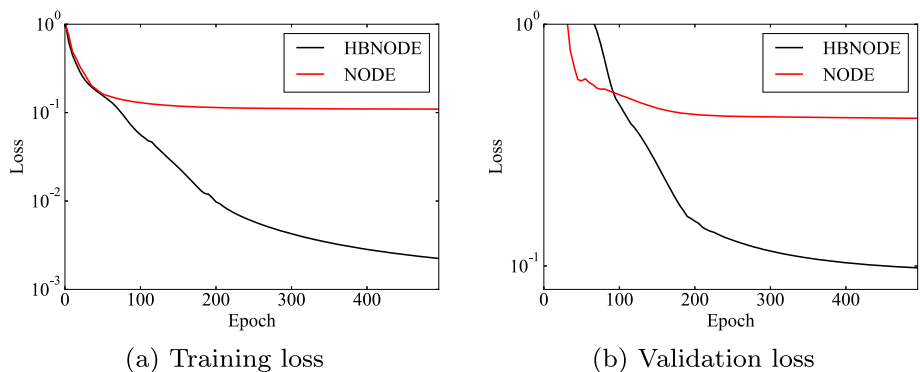


Fig. 10 Contrasting training and validation loss of NODE and HBNODE for learning ROMs of transient to steady-state VKS dynamics. The progress made by the NODE is significantly reduced by the vanishing gradient. HBNODE has a much slower decaying gradient and is able to continue to make progress in both training and validation sets

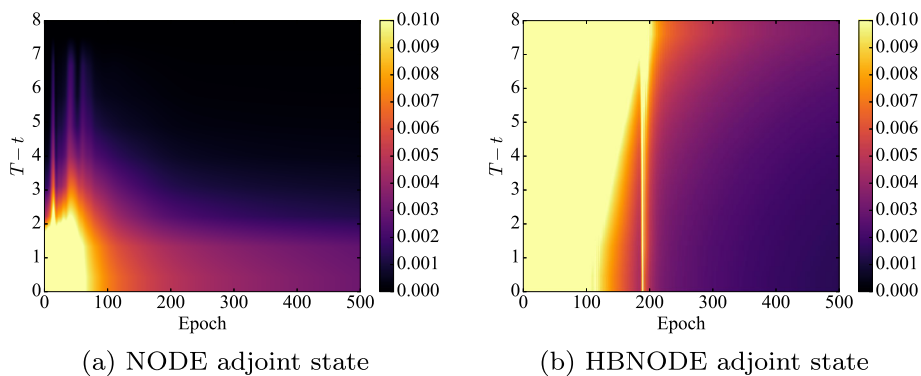


Fig. 11 Comparison of the adjoint states for the NODE and HBNODE in learning multi-input-single-output. The NODE adjoint state vanishes substantially faster than HBNODE

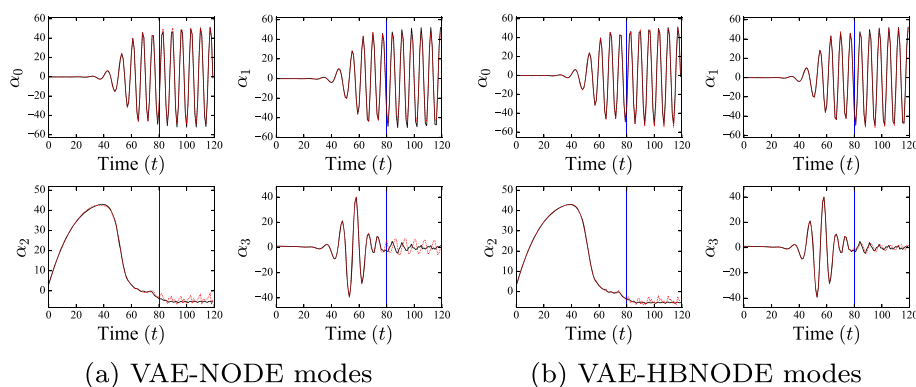


Fig. 12 Comparison of modes reconstruction for NODE and HBNODE-based ROMs for learning multi-input-single-output task. NODE is able to reliably learn the first two dominant modes but is unable to capture the steady-state dynamics, having a much larger frequency and introducing substantial lag into the oscillation frequency. HBNODE is able to more reliably capture the steady-state dynamics with slightly larger frequency and developing lag much later than the NODE component. Before and after the vertical blue line stands for training and validation, respectively

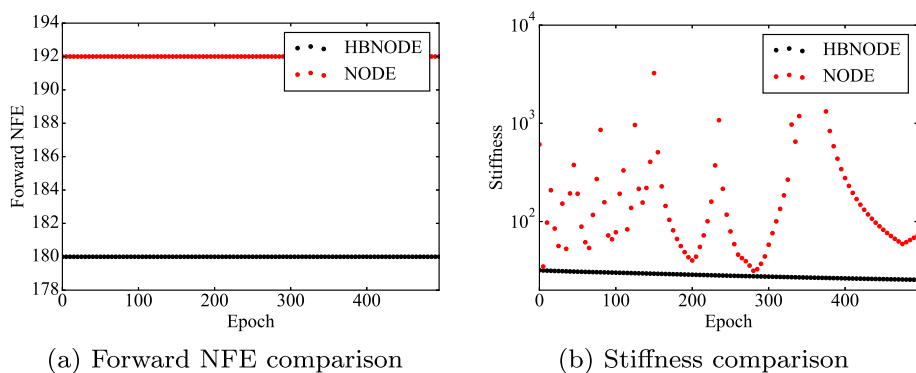


Fig. 13 Comparison of NFEs and stiffness for the NODE and HBNODE in learning transient to steady-state VKS dynamics. NODE requires more NFEs in each forward pass than HBNODE as the NODE is much stiffer than HBNODE. The stiffness of NODE varies sharply as training goes on, while the stiffness of HBNODE decays during the training

In terms of the predictive performance, as shown in Fig. 12, the HBNODE predictor captures the peaks of the oscillatory dynamics better than NODEs, especially in the first two modes α_1 and α_2 . Moreover, the prediction error using NODE is much larger than that of HBNODE, and the prediction error amplifies as the prediction time goes, in particular, for modes α_3 and α_4 .

Another primary advantage of HBNODE over NODE-based ROMs lies in computational efficiency, which is theoretically supported by the discussion in Sect. 3.2.1. As shown in Fig. 13a, the forward NFE required in each forward pass by HBNODE is consistent smaller than that of NODE. We also monitor the stiffness of both NODE and HBNODE during the learning process, and Fig. 13b shows that the stiffness of NODE oscillates and maintains much larger than HBNODE.




Fig. 14 The decay of relative information content for POD and DMD of the KPP dataset. The rapid decay in the modes indicates the problem is suitable for POD and DMD




Fig. 15 The dominant POD modes for the KPP data. The oscillations of the modes decay rapidly and have a very low frequency, making it challenging for learning compared to the quasi-periodic oscillations of the VKS dataset

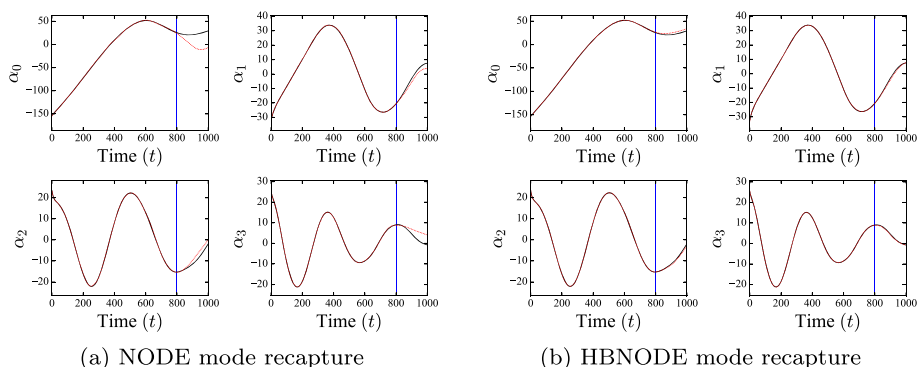
6.2 KPP Model

We obtain the KPP dataset by simulating the FOM presented in Sect. 4.2 for 1000 timesteps ($N_t = 1000$). The KPP model is well-suited for reduced-order modeling due to the rapidly decaying eigenvalues in both POD and DMD, seeing Fig. 14. However, we found in our experiments that it is particularly difficult to capture the dynamics using machine learning architectures due to the slow decaying ROM dynamics depicted in Fig. 15.

In our experiments, we note that the non-lifted DMD continuously deforms the center of mass in a way that defies the physical constraints of the system. A comparison of lifted and non-lifted DMD predictions and POD predictions are shown in [3]. To lift DMD, we utilized the lifting functions $\{\cos(\mathbf{x}), \sin(\mathbf{x}), \mathbf{x}^2, \mathbf{x}^3\}$ with $\mathbf{x} = (x, y)$. Figure 14 shows that the POD

Table 3 The hyperparameters of NODE and NODE for the learning ROMs of KPP model

Hyper-parameter	Value
Layers	2
Hidden layers	64
Sequence length	4
Learning rate	.01
Epochs	500

**Fig. 16** Comparison of the KPP modes prediction using NODE and HBNODE. HBNODE is better at predicting the dynamics for the validation set than NODE. Before and after the vertical blue line stands for training and validation, respectively

modes decay faster than the lifted DMD modes. The dominant 24 DMD modes correspond to 97% of the relative information content; in contrast, the dominant 8 POD modes correspond to 99% of the relative information value.

We train the pipeline depicted in Fig. 3 for learning multi-input-single-output dynamics. The data is constructed from the 8 dominant POD modes on the time interval from $t = 0$ to 1000. The data is sequenced so that every 4 preceding time step is used to predict the 5-th time step. The training data consists of the POD modes from $t = 0$ to 799 and the training labels consist of POD modes from $t = 5$ to 800. The validation data utilizes data from $t = 800$ to 999 and the validation labels consist of data from $t = 805$ to 1000. The model's hyperparameters are the same for NODE and HBNODE components and are given in Table 3.

6.2.1 Results and comparison to existing ROMs

We compare the prediction of NODE and HBNODE against ground truth in Fig. 16, and we see that HBNODE performs remarkably better than NODE in predicting the dynamics. In particular, HBNODE is able to properly capture the oscillation dynamics of the modes, unlike NODE. Figure 17 shows that HBNODE has a much smaller training and validation loss than that of NODE.

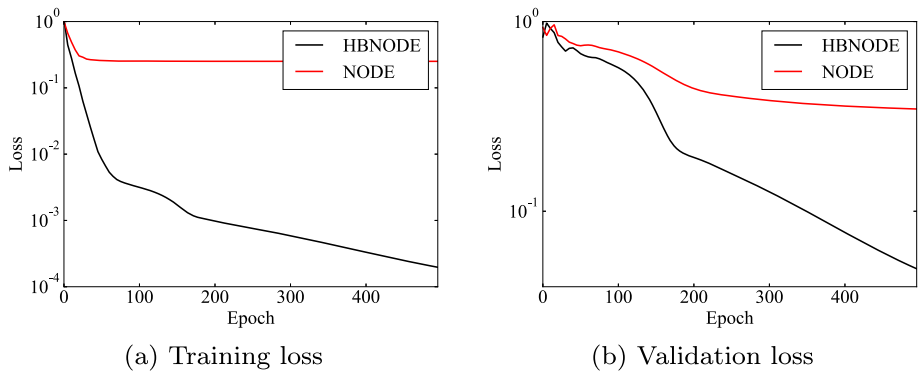


Fig. 17 Comparison of the training and validation loss of NODE and HBNODE for learning ROMs for the KPP model. NODE is unable to make progress due to a rapidly vanishing gradient, impeding learning long-range dependencies

Table 4 The hyperparameters of NODE and HBNODE for the learning ROMs of Euler equations

Hyper-parameter	Value
Layers	6
Hidden layers	16
Learning rate	.01
Epochs	100

6.3 Euler Equations for Fluids Modeling

We further consider learning reduced-order models for simulating the Euler equations, where the dataset is obtained by simulating the full-order model presented in Sect. 4.3 with a discrete ensemble of parameters η_1, \dots, η_M with $M = 100$, over 180 timesteps. Two different parameters η values can produce rather different dynamics, as evidenced in Fig. 18. The Euler equations data is unique in the sense that it may be segmented based on these initial conditions. The ROM is generated by taking the dominant 8 POD modes for each parameter η_i on the time interval from $t = 0$ to 180. This data is shuffled randomly among the initial parameter η_i to no longer increase sequentially. The average relative information content across all 100 values of η is $\sim 95\%$.

In this task, we train the machine learning pipeline shown in Fig. 3 for learning multi-input-multi-output dynamics. The training dataset comprises the dominant 8 POD modes for each of the training parameters among $\eta_1, \dots, \eta_{\text{train}}$. We use 90 of the 100 parameters, η_1, \dots, η_{90} , for training and the rest for validation. The training input consists of the dominant 8 POD modes for each of the 90 training parameters on the time interval from $t = 0$ to 150. The training labels consist of the dominant 8 POD modes for each of the 90 training parameters time steps from $t = 151$ to 180. The validation dataset is composed of the validation parameters $\eta_{91}, \dots, \eta_{100}$. The validation input and labels are segmented using the same intervals as the training data.

The model uses a GHBNODE component with a hyperbolic tangent activation function. All other experimental settings are the same as in the KPP dataset, and the tuned hyperparameters are listed in Table 4. The NODE and HBNODE models are trained and validated over the same data shuffling.




Fig. 18 The Euler Equations data at time step $t = 50$ with two different initial parameters, η_0 and η_{50} . Variations in the parameter η produce widely varying dynamics and as a result varying POD modes. The objective of this task is to predict the dynamics for unseen parameters η using a set of training parameters η_{train}

6.3.1 Results and comparison to existing models

We compare the prediction of NODE and GHBNODE for a randomly selected parameter η_{train} from the training set and a randomly selected parameter η_{valid} from the validation set. The modes for the training parameter η_{train} are shown in Fig. 19, and the modes for the validation parameter η_{valid} are shown in Fig. 20. We observe that the POD modes for the parameter η_{train} in Fig. 19 differ from those for η_{valid} primarily in amplitude rather than shape. As a result, a poor prediction model will have a sudden discontinuity between the input and prediction values. The transition point between the input and the prediction is indicated in Figs. 19 and 20 by the vertical blue line.

Figure 19 shows the dominant 4 POD modes for η_{train} from the training set. The predictive capabilities of GHBNODE significantly outperform the NODE model. In particular, for α_2 and α_3 , we observe that the NODE has a large jump discontinuity at the transition between the input and prediction. The GHBNODE modes are smoother in the transition region, which indicates the ability of the GHBNODE model to distinguish between separate parameters.

In Fig. 20, we observe the same characteristics for NODE and GHBNODE. NODE is unable to accurately predict the output for the parameter η_{valid} from the validation set. In particular, for α_3 of the validation parameter, the NODE prediction is even less smooth than that of the training parameter shown in Fig. 19. In this experiment, we observe that NODE is unable to distinguish data with varying parameters as accurately as GHBNODE.

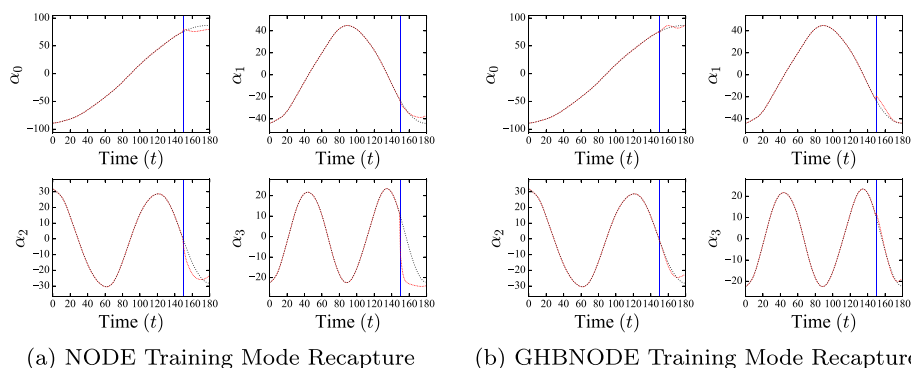


Fig. 19 For a randomly selected initial parameter η_{train} , we plot the dominant four modes and their predicted data using NODE and GHBNODE. The blue vertical line separates the input and output data. The ground truth is in black, while the prediction data is in dashed red. GHBNODE can learn the dynamics of α_2 and α_1 for the parameterized data significantly better than NODE. This is evidenced by the fact that NODE predicts the inflection point of α_2 too early

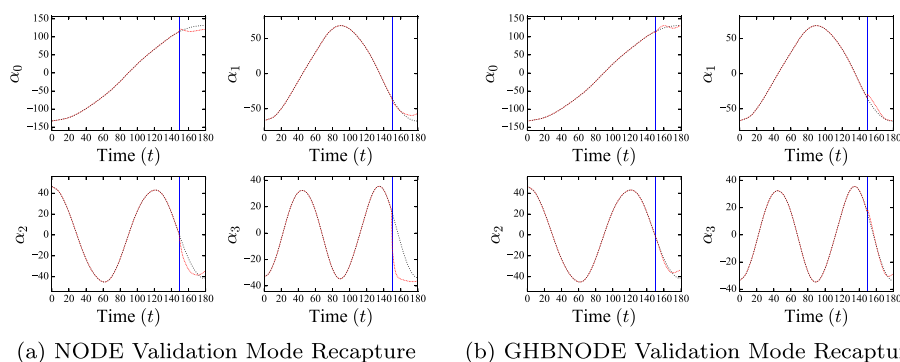


Fig. 20 For the randomly selected initial parameter η_{valid} , we plot the dominant four modes and their predicted data using NODE and GHBNODE. The blue vertical line separates the input and output data. The ground truth is in black, while the prediction data is in dashed red. HBNODE can predict the data for unseen parameterizations much better than NODE

7 Concluding Remarks

This paper employs the recently developed HBNODEs and their generalization for learning POD coefficients for model reduction. We analyze through simple linearized models and empirically verify the advantages of HBNODEs over existing NODEs. In particular, HBNODEs enjoy the following advantages that imply practical benefits for learning POD-based ROMs, including 1) The deep learning model is continuous-depth, providing flexibility in learning irregularly-sampled time series and faithful to the continuous profiling of the underlying physical models. 2) Both the forward and adjoint ODEs of HBNODEs are of the heavy-ball style, accelerating both training and testing of the machine learning procedure. And 3) HBNODEs can learn long-range dependencies effectively, capturing intrinsic patterns from data. There are numerous avenues for future works, and two particular interesting directions in our mind are 1) Improving HBNODEs, particularly replacing the fine-tuned or learned damping parameter with an adaptive one that are motivated by certain optimization

algorithms with adaptive momentum [54, 55, 57], and 2) Applying HBNODE-based ROMs to model reduction arising from scientific challenges, especially when we do not have the ground truth governing equation of the dynamical systems.

Acknowledgements This material is based on research sponsored by NSF grants DMS-1848508, DMS-1924935, DMS-1952339, DMS-2110145, DMS-2111117, DMS-2152762, and DMS-2208361, DOE grant DE-SC0021142 and DE-SC0023490, and AFOSR FA9550-20-1-0338. We also acknowledge support from a seed grant from the College of Science at the University of Utah.

Data Availability All data and code related to this paper are available at https://github.com/JustinBakerMath/pod_hbnode/.

References

1. Antoulas, A.: Approximation of Large-Scale Dynamical Systems. Advances in Design and Control. Society for Industrial and Applied Mathematics (2005)
2. Antoulas, A.C., Beattie, C.A., Gugercin, S.: Interpolatory Model Reduction of Large-Scale Dynamical Systems. In: Mohammadpour, J., Grigoriadis, K.M. (eds.) Efficient Modeling and Control of Large-Scale Systems, pp. 3–58. Springer, Boston, MA (2010)
3. Baker, J., Cherkaev, E., Narayan, A., Wang, B.: Learning pod of complex dynamics using heavy-ball neural odes: Animations. https://www.github.com/JustinBakerMath/pod_hbnode/blob/master/README.md#animations
4. Bengio, Y., Simard, P., Frasconi, P.: Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Netw.* **5**(2), 157–166 (1994)
5. Benner, P., Gugercin, S., Willcox, K.: A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev.* **57**(4), 483–531 (2015)
6. Benner, P., Gugercin, S., Willcox, K.: A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev.* **57**(4), 483–531 (2015)
7. Berkooz, G., Holmes, P., Lumley, J.L.: The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* **25**(1), 539–575 (1993)
8. Bittner, L., Pontryagin, L.S., Boltyanskii, V.G., Gamkrelidze, R.V., Mishechenko, E.F.: the mathematical theory of optimal processes. VIII + 360 S. New York/London 1962. Wiley. Preis 90/-. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 43(10–11):514–515, (1963)
9. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 31. Curran Associates, Inc., New York (2018)
10. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014)
11. Cohen, M.A., Grossberg, S.: Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Trans. Syst. Man Cybern.* **SMC-13**(5), 815–826 (1983)
12. Craster, R.V., Matar, O.K.: Dynamics and stability of thin liquid films. *Rev. Mod. Phys.* **81**(3), 1131 (2009)
13. Dormand, J.R., Prince, P.J.: A family of embedded runge-kutta formulae. *J. Comput. Appl. Math.* **6**(1), 19–26 (1980)
14. Dupont, E., Doucet, A., Teh, Y.W.: Augmented neural odes. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc. (2019)
15. Dutta, S., Rivera-Casillas, P., Cecil, O.M., Farthing, M.W., Perracchione, E., Putti, M.: Data-driven reduced order modeling of environmental hydrodynamics using deep autoencoders and neural odes. *arXiv preprint arXiv:2107.02784*, (2021)
16. Dutta, S., Rivera-Casillas, P., Farthing, M.W.: Neural ordinary differential equations for data-driven reduced order modeling of environmental hydrodynamics. *arXiv preprint arXiv:2104.13962* (2021)
17. Germano, M., Piomelli, U., Moin, P., Cabot, W.H.: A dynamic subgrid-scale eddy viscosity model. *Phys. Fluids A* **3**(7), 1760–1765 (1991)
18. Gugercin, S., Antoulas, A.C.: A Survey of Model Reduction by Balanced Truncation and Some New Results. *Int. J. Control* **77**(8), 748–766 (2004)

19. Harten, A., Lax, P.D., van Leer, B.: On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Rev.* **25**(1), 35–61 (1983)
20. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pp. 630–645 (2016)
21. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
22. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
23. Jacot, A., Gabriel, F., Hongler, C.: Neural tangent kernel: convergence and generalization in neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 8580–8589, (2018)
24. Kani, J.N., Elsheikh, A.H.: Dr-rnn: A deep residual recurrent neural network for model reduction. *arXiv preprint arXiv:1709.00939*, (2017)
25. Nagoor Kani, J., Elsheikh, A.H.: Reduced-order modeling of subsurface multi-phase flow models using deep residual recurrent neural networks. *Transp. Porous Media* **126**(3), 713–741 (2019)
26. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
27. Kuramoto, Y.: Diffusion-Induced Chaos in Reaction Systems. *Prog. Theor. Phys. Suppl.* **64**, 346–367 (1978). (02)
28. Kurganov, A., Petrova, G., Popov, B.: Adaptive Semidiscrete Central-Upwind Schemes for Nonconvex Hyperbolic Conservation Laws. *SIAM J. Sci. Comput.* **29**(6), 2381–2401 (2007)
29. Lechner, M., Hasani, R.: Learning long-term dependencies in irregularly-sampled time series. *arXiv preprint arXiv:2006.04418*, (2020)
30. Liang, Y.C., Lee, H.P., Lim, S.P., Lin, W.Z., Lee, K.H., Wu, C.G.: Proper orthogonal decomposition and its applications-part i: Theory. *J. Sound Vib.* **252**(3), 527–544 (2002)
31. Lui, H.F.S., Wolf, W.R.: Construction of reduced-order models for fluid flows using deep feedforward neural networks. *J. Fluid Mech.* **872**, 963–994 (2019)
32. Ma, C., Wang, J., et al.: Model reduction with memory and the machine learning of dynamical systems. *arXiv preprint arXiv:1808.04258*, (2018)
33. Mannarino, A., Mantegazza, P.: Nonlinear aeroelastic reduced order modeling by recurrent neural networks. *J. Fluids Struct.* **48**, 103–121 (2014)
34. Massaroli, S., Poli, M., Park, J., Yamashita, A., Asama, H.: Dissecting neural odes. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, pp. 3952–3963. Curran Associates, Inc., New York (2020)
35. Maulik, R., Lusch, B., Balaprakash, P.: Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders. *Phys. Fluids* **33**(3), 037106 (2021)
36. Mohebbujaman, M., Rebholz, L.G., Iliescu, T.: Physically constrained data-driven correction for reduced-order modeling of fluid flows. *Int. J. Numer. Meth. Fluids* **89**(3), 103–122 (2019)
37. Moin, P., Mahesh, K.: Direct numerical simulation: a tool in turbulence research. *Annu. Rev. Fluid Mech.* **30**(1), 539–578 (1998)
38. Mou, C., Liu, H., Wells, D.R., Iliescu, T.: Data-driven correction reduced order models for the quasi-geostrophic equations: A numerical investigation. *Int. J. Comput. Fluid Dyn.* **34**(2), 147–159 (2020)
39. Murata, T., Fukami, K., Fukagata, K.: Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *J. Fluid Mech.* **882**, A13 (2020)
40. Nguyen, T., Baraniuk, R., Bertozzi, A., Osher, S., Wang, B.: MomentumRNN: Integrating momentum into recurrent neural networks. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, pp. 1924–1936. Curran Associates, Inc., New York (2020)
41. Norcliffe, A., Bodnar, C., Day, B., Simidjievski, N., Lió, P.: On second order behaviour in augmented neural odes. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, pp. 5911–5921. Curran Associates, Inc., New York (2020)
42. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pp. 1310–1318, (2013)
43. Pearson, K.: Liit on lines and planes of closest fit to systems of points in space. *London Edinburgh Dublin Philosoph. Magaz. J. Sci.* **2**(11), 559–572 (1901)
44. Polyak, B.T.: Some methods of speeding up the convergence of iteration methods. *USSR Comput. Math. Math. Phys.* **4**(5), 1–17 (1964)
45. Rojas, C.J.G., Dengel, A., Ribeiro, M.D.: Reduced-order Model for Fluid Flows via Neural Ordinary Differential Equations. *arXiv:2102.02248 [physics]*, February (2021). *arXiv: 2102.02248*
46. Rosenblatt, F.: Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Technical report, Cornell Aeronautical Lab Inc Buffalo NY, (1961)

47. Rubanova, Y., Chen, R.T.Q., Duvenaud, D.K.: Latent ordinary differential equations for irregularly-sampled time series. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., (2019)
48. San, O., Maulik, R.: Neural network closures for nonlinear model order reduction. *arXiv preprint [arXiv:1705.08532](https://arxiv.org/abs/1705.08532)*, (2017)
49. San, O., Maulik, R.: Machine learning closures for model order reduction of thermal fluids. *Appl. Math. Model.* **60**, 681–710 (2018)
50. San, O., Maulik, R., Ahmed, M.: An artificial neural network framework for reduced order modeling of transient flows. *Commun. Nonlinear Sci. Numer. Simul.* **77**, 271–287 (2019)
51. Schmid, P.J.: Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28 (2010)
52. Sivashinsky, G.I.: On flame propagation under conditions of stoichiometry. *SIAM J. Appl. Math.* **39**(1), 67–82 (1980)
53. Sivashinsky, G.I.: Nonlinear analysis of hydrodynamic instability in laminar flames-i. derivation of basic equations. *Acta Astronaut.* **4**(11), 1177–1206 (1977)
54. Sun, T., Ling, H., Shi, Z., Li, D., Wang, B.: Training deep neural networks with adaptive momentum inspired by the quadratic optimization. *arXiv preprint [arXiv:2110.09057](https://arxiv.org/abs/2110.09057)*, (2021)
55. Wang, B., Nguyen, T.M., Bertozzi, A.L., Baraniuk, R.G., Osher, S.J.: Scheduled restart momentum for accelerated stochastic gradient descent. *arXiv preprint [arXiv:2002.10583](https://arxiv.org/abs/2002.10583)*, (2020)
56. Wang, B., Xia, H., Nguyen, T., Osher, S.: How does momentum benefit deep neural networks architecture design? a few case studies. *arXiv preprint [arXiv:2110.07034](https://arxiv.org/abs/2110.07034)*, (2021)
57. Wang, B., Ye, Q.: Stochastic gradient descent with nonlinear conjugate gradient-style adaptive momentum. *arXiv preprint [arXiv:2012.02188](https://arxiv.org/abs/2012.02188)*, (2020)
58. Wang, M., Li, H.-X., Chen, X., Chen, Y.: Deep learning-based model reduction for distributed parameter systems. *IEEE Trans. Syst. Man Cybern. Syst.* **46**(12), 1664–1674 (2016)
59. Xia, H., Suliafu, V., Ji, H., Nguyen, T., Bertozzi, A., Osher, S., Wang, B.: Heavy ball neural ordinary differential equation. In *Advances in Neural Information Processing Systems*, volume 34. Curran Associates, Inc., (2021)
60. You, D., Moin, P.: A dynamic global-coefficient subgrid-scale eddy-viscosity model for large-eddy simulation in complex geometries. *Phys. Fluids* **19**(6), 065110 (2007)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.