
NONLINEAR PROPER ORTHOGONAL DECOMPOSITION FOR CONVECTION-DOMINATED FLOWS

A PREPRINT

Shady E. Ahmed

School of Mechanical & Aerospace Engineering,
Oklahoma State University,
Stillwater, OK 74078, USA.
shady.ahmed@okstate.edu

Omer San

School of Mechanical & Aerospace Engineering,
Oklahoma State University,
Stillwater, OK 74078, USA.
osan@okstate.edu

Adil Rasheed

Department of Engineering Cybernetics,
Norwegian University of Science and Technology,
N-7465, Trondheim, Norway.
Department of Mathematics and Cybernetics,
SINTEF Digital,
7034 Trondheim, Norway.
adil.rasheed@ntnu.no

Traian Iliescu

Department of Mathematics,
Virginia Tech,
Blacksburg, VA 24061, USA.
iliescu@vt.edu

ABSTRACT

Autoencoder techniques find increasingly common use in reduced order modeling as a means to create a latent space. This reduced order representation offers a modular data-driven modeling approach for nonlinear dynamical systems when integrated with a time series predictive model. In this letter, we put forth a nonlinear proper orthogonal decomposition (POD) framework, which is an end-to-end Galerkin-free model combining autoencoders with long short-term memory networks for dynamics. By eliminating the projection error due to the truncation of Galerkin models, a key enabler of the proposed nonintrusive approach is the kinematic construction of a nonlinear mapping between the full-rank expansion of the POD coefficients and the latent space where the dynamics evolve. We test our framework for model reduction of a convection-dominated system, which is generally challenging for reduced order models. Our approach not only improves the accuracy, but also significantly reduces the computational cost of training and testing.

Keywords Reduced order models, data-driven modeling, autoencoders, long short-term memory networks

1 Introduction

Full order models (FOMs) based on the solution of the Navier-Stokes equations, e.g., direct numerical simulations, large eddy simulations, and Reynolds averaged Navier-Stokes simulations, have made a tremendous impact in the numerical simulation of high Reynolds number fluid flows. However, current FOMs cannot be used effectively in multiple-query simulations (e.g., uncertainty quantification, optimization, and control) because of their prohibitively high computational cost. Reduced order models (ROMs), on the other hand, are efficient low-dimensional models created from available data. ROMs have been often used as surrogate models for structure-dominated problems. However, although traditional projection-based ROMs work in simple, canonical test problems, they generally fail (or require a significantly large number of basis functions) for convection-dominated flows because a low-dimensional ROM basis cannot accurately represent the complex dynamics.

Among fluid dynamicists, projection-based truncated methods are quite popular. Practitioners often utilize the proper orthogonal decomposition (POD) method to generate a set of orthonormal basis functions, or a reduced order representation [1]. After a suitable basis is chosen, the Galerkin projection is generally used to project the dynamics of the equation onto the subspace spanned by a truncated set of basis functions (e.g., see the Galerkin POD (GPOD)

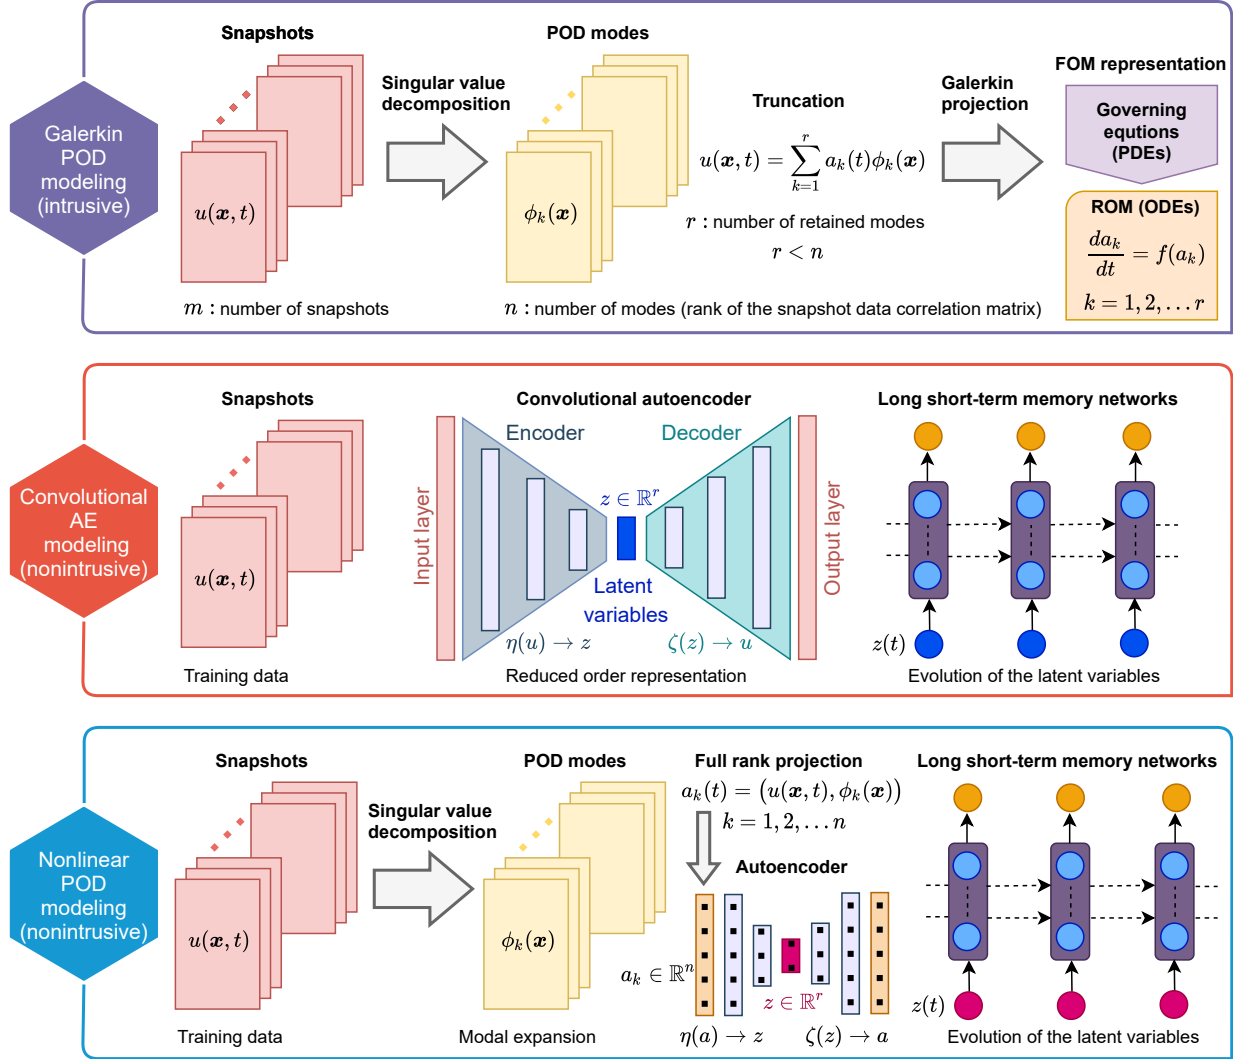


Figure 1: Illustration of the research methodology: (i) Galerkin POD modeling – an intrusive approach, which has been the workhorse for most projection-based ROM applications, (ii) convolutional autoencoder modeling – a modular nonintrusive ROM framework that can be built from data without requiring access to the underlying governing equations, and (iii) the proposed nonlinear POD modeling approach that utilizes the autoencoder technology on a manifold defined by the full-rank POD coefficients (i.e., considering 99.9% of the relative information content). The elimination of the subspace projection error stemming from the truncation is a major advantage of the proposed nonlinear POD method over the Galerkin POD approach, especially for the nonlinear dynamical systems with slow decay in the Kolmogorov N-width, such as convection-dominated flows studied in this paper. The proposed nonintrusive approach has also two other advantages over the convolutional autoencoder model. First, it can reduce the computational cost of the training process since it involves significantly fewer trainable parameters. Second, it ensures an accurate implementation of the boundary conditions since the spatial variations are naturally embedded in the reduced order representation.

approach in the top panel of Figure 1). Although time-periodic or quasi-time-periodic dynamical systems can be easily represented by a small number of POD modes, convection-dominated unsteady flows might require a large number of POD modes. Therefore, a significant projection error is often introduced by a truncation process. This projection error increases for dynamical systems when the Kolmogorov N-width [2, 3] increases. The Kolmogorov N-width is an approximation theory concept that determines the linear reducibility of the underlying systems, which can be connected to the POD spectrum [4].

The emerging autoencoder (AE) technology provides a powerful alternative way of generating a reduced order representation, often called latent space or finite-dimensional manifold. In recent years, there has been an ever-

increasing interest in discovering inertial manifolds of partial differential equations (PDEs) using autoencoders [5–9]. Convolutional AE (CAE) is a promising approach that has been utilized in fluid dynamics applications [10–19]. CAEs are often combined with a time series prediction approach to model the evolving dynamics of the latent space [20]. A self-attention mechanism is also used within the CAEs to enhance the feature extraction ability of the network [21, 22]. A combination of principal component analysis (PCA) with Gaussian process regression has been investigated in the development of digital twins of reacting flow applications [23]. These approaches bypass the Galerkin projection and hence become fully nonintrusive [24–27], building an end-to-end ROM solely from available data without requiring any knowledge of the underlying evolution equations. However, the number of trainable parameters might quickly reach millions even for two-dimensional PDE settings [28]. Moreover, the padding approaches that are available in most CAE architectures might pose additional challenges in representing the desired boundary conditions unless specific customization is embedded.

In this letter, we introduce a simple and modular nonintrusive Galerkin-free ROM framework for unsteady flows by employing ideas from POD-based and AE-based ROM approaches (see the bottom panel in Figure 1). In our methodology, we first apply the POD procedure to generate a set of orthonormal basis functions. Instead of truncating the number of modes, we consider an almost full-rank POD expansion defined by 99.9% of the relative information content (RIC) measure, and perform an inner product between snapshots and the POD modes to obtain the set of POD coefficients. We then construct a plain multilayer perceptron AE for finding the embedding of these POD coefficients (which requires substantially fewer trainable parameters compared to those required by the convolutional AEs) to generate a nonlinear mapping between these POD coefficients and a latent space constructed with only a few parameters in the bottleneck layer. Finally, we utilize a recurrent neural network approach for integrating evolution dynamics of the latent space. In this modular way, we utilize the best features of all relevant approaches. Specifically, POD is utilized in ranking the structural content using a linear spatiotemporal decomposition, AE is used in nonlinear dimension reduction, and LSTM is integrated for the time series prediction. Following the established name convention of nonlinear PCA (NLPCA) [29–32], we call our approach nonlinear POD (NLPOD).

We note that previous studies aimed at leveraging the synergy between POD and AEs. One notable example is their use together with dynamic mode decomposition approaches [33–35]. In particular, the combination of POD with AE has been exploited to efficiently reveal Koopman invariant subspaces [36]. Furthermore, the idea of using POD and AE to develop a constrained Koopman neural network framework has been proposed by Puligilla and Jayaraman [37]. Although most ROM efforts aim at providing a predictive model trained with ample amount of data with carefully-designed experiments to be operated in real-time, ROM can also act as a method for data compression to mitigate challenges related to data storage and transfer. For instance, Carlberg *et al.* [38] applied PCA to reduce the dimensionality of the vector of local encodings across the entire spatial domain to facilitate data I/O in large-scale simulations.

The combination of POD/PCA and AE has been also utilized for model order reduction, similar to the NLPOD methodology. Casas, Arcucci, and Guo [39] retained all the principal components (PCs) of urban air pollutant data and applied a fully connected AE onto these PCs for further reduction. They found that the PC-AE requires substantially fewer trainable parameters than an equivalent fully connected AE on the full-space. This approach was adopted by Phillips *et al.* [40] with an additional step of linearizing the decoder mapping from the latent variables to the scalar fluxes to reduce the computational costs of solving eigenvalue problems using the power iteration method. This combination was also employed by Kosut, Ho, and Rabitz [41] in the quantum physics community as a better alternative than using a plain-vanilla AE approach.

In the present study, we apply AE onto the coefficients of a RIC-guided POD expansion of parameterized convection-dominated flows. The truncation of basis function beyond a prespecified RIC value helps to filter out redundant and noisy signals before passing them to the AE. It also provides an informed trade-off between accuracy and storage requirements. Furthermore, we discuss in detail the benefits of using POD as an upper layer for the interface with the physical space, compared to CAEs. This includes computational aspects (e.g., training and deployment costs) and physical insights (e.g., respecting boundary conditions, symmetries, and conservation laws). Furthermore, we explore the applicability of NLPOD for time-dependent parametric systems, where we augment the bottleneck layer of the autoencoder with the operating parameters to enhance the latent space discovery. We highlight that NLPOD provides a robust end-to-end ROM data compression framework built using significantly fewer trainable parameters than CAEs. Moreover, NLPOD can easily handle data on unstructured and non-uniform grids. Our approach could be well suited for the digital twin applications, where fast transfer learning procedures are often desired when new training data streams are incorporated [42].

2 GPOD modeling

In standard Galerkin-based ROMs, a handful of orthonormal basis functions are predetermined with the assumption that the solution approximately lives in the subspace spanned by these bases. POD provides a systematic way for building such basis functions from collected data sets (often called snapshots). Arguably, the simplest way to perform POD involves the singular value decomposition (SVD). For instance, assuming that the state variable is denoted by $u(\mathbf{x}, t) \in \mathbb{R}^N$ (where N is the spatial resolution), a data matrix $\mathbf{A} \in \mathbb{R}^{N \times m}$ can be formed by stacking m temporal snapshots as $\mathbf{A} = [u(\mathbf{x}, t_1), u(\mathbf{x}, t_2), \dots, u(\mathbf{x}, t_m)]$. Then, an SVD can be applied either directly onto \mathbf{A} (or onto its mean-subtracted/shifted version) as follows:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*, \quad (1)$$

where \mathbf{U} and \mathbf{V} are orthogonal matrices representing the left and right singular vectors of \mathbf{A} , respectively, while $\mathbf{\Sigma}$ is a diagonal matrix containing the corresponding singular values, σ_k , in descending order. The columns of \mathbf{U} define the sought POD modes, $\mathbf{U} = [\phi_1, \phi_2, \dots]$. Moreover, the inherent sorting of the singular values provides a valuable feature of POD since the most important modes are placed first and the RIC values of the leading r modes can be defined as

$$\text{RIC}(\%) = \frac{\sum_{k=1}^r \sigma_k^2}{\sum_{k=1}^m \sigma_k^2} \times 100, \quad (2)$$

where σ_k is the k^{th} singular value. Then, a rank- r approximation of the state variable can be written as a linear superposition of the contributions of these first r modes as follows:

$$u(\mathbf{x}, t) = \sum_{k=1}^r a_k(t) \phi_k(\mathbf{x}). \quad (3)$$

In order to compute (evolve) the time-dependent coefficients $a_k(t)$, Galerkin-based methods perform a projection of the governing equations onto the corresponding POD modes to yield a system of r ordinary differential equations for the vector of coefficients \mathbf{a}_k as follows:

$$\frac{d\mathbf{a}}{dt} = f(\mathbf{a}). \quad (4)$$

Due to the quadratic nonlinearity in most fluid flow systems, the *online* computational cost of the Galerkin POD (GPOD) approach scales cubically with the number of retained modes. Thus, a modal truncation is performed to achieve a computational gain from GPOD. Nonetheless, most convection-dominated flows cannot be simply described using Eq. 3 if r is small, which leaves us working in under-resolved regimes, where there are both a representation error and a closure error [43].

3 CAE modeling

In order to mitigate the deficiency of POD in providing a *compact* set of basis function, CAE frameworks have shown substantial success in providing a compressed latent space that defines a nonlinear manifold on which the system's dynamics evolves more accurately than on the linear manifold defined by a similar POD compression. The CAE starts with an encoding process that involves applying a series of convolutions and nonlinear mappings onto the input snapshot data to shrink the dimensionality down to a bottleneck layer representing the low rank or latent space embedding. An inverse mapping from the latent space variables to the physical space is performed by another set of deconvolutions and nonlinear mappings, defining the decoding part. For example, denoting the encoder function as η and the decoder as ζ , we can represent the manifold learning through autoencoder as follow,

$$\eta, \zeta = \arg \max_{\eta, \zeta} \|\mathbf{A} - (\eta \circ \zeta)(\mathbf{A})\|, \quad (5)$$

$$\eta : u(\mathbf{x}, t) \in \mathbb{R}^N \rightarrow \mathbf{z} \in \mathbb{R}^r, \quad (6)$$

$$\zeta : \mathbf{z} \in \mathbb{R}^r \rightarrow u(\mathbf{x}, t) \in \mathbb{R}^N, \quad (7)$$

where \mathbf{z} represents the low order code at the bottleneck. Note that η and ζ are parameterized by the neural network weights and biases, which are learned in the training process.

For the temporal dynamics, a surrogate model emulator is constructed to evolve the latent variables onto the manifold revealed by the CAE. In the present study, we utilize the capabilities of the long short-term memory (LSTM) networks in sequential data prediction to propagate the latent variables in time. Despite the success and popularity of CAEs in the past few years, there still exist potential limitations to their applications in fluid flows. Below, we highlight a few of them.

- The number of trainable parameters quickly explodes with the dimensionality of input data and complexity of the CNN architecture. As a result, proper training requires a prohibitively large number of training samples and a long training time.
- The design of CNNs often involves padding operations, and physically-consistent boundary conditions are not guaranteed unless custom padding is adopted (e.g., for periodic boundary conditions). Thus, it is more challenging to impose general boundary conditions in CAEs.
- It is rarely possible to physically and/or mathematically interpret the CAE latent space. Therefore, the use of CAE to extract and analyze the underlying coherent structures is limited.
- Plain CAEs do not guarantee conservation properties (e.g., continuity) unless additional constraints are explicitly imposed [44].
- There are no clear estimates for error bounds given a latent space dimension, and the output totally depends on the architecture hyperparameters and the training process. Thus, it is not possible to evaluate how much accuracy is gained/lost by varying the dimensionality of the latent space. Similarly, there is no information regarding the relevant importance of different latent variables.
- CNNs have been intrinsically designed to process data that are uniformly sampled in the spatial domain. However, in computational disciplines, we often encounter data that are defined on non-uniform grids. Plain-vanilla CAEs cannot be applied in these cases unless a mapping of the original data onto a uniform grid is employed [45].

4 NLPOD modeling

In the proposed NLPOD framework, we aim at utilizing the capabilities of both the GPOD and CAE modeling approaches while mitigating their potential limitations. In particular, we are interested in promoting the following key benefits of the POD methodology:

- There is a rich history of POD developments for large scale problems, which makes the compression step more computationally efficient than training CAE, e.g., by using randomization and streaming algorithms.
- By construction, the POD basis functions respect the underlying boundary conditions prevailing in the input data sets. Therefore, the reconstructed fields from POD are supported by physically consistent boundary conditions.
- The flow fields reconstructed using POD satisfy the conservation laws, which improves the physical soundness of the resulting ROM.
- There exists a clear mathematical and physical definition of the resulting POD modes, which renders POD as a useful tool for further analysis and interpretation of the dominant patterns [46].
- Each mode is associated with a metric of its relative importance and contribution to the overall information/variance in the given data sets. Therefore, error bounds have been proved for a given level of data compression.
- POD is not restricted to specific representation of the flow field variables. For example, POD can be easily extended to data defined on non-uniform or unstructured grids by employing different types of inner products [47] and numerical integration schemes.

However, in order to accurately describe convection-dominated flows, a large number of POD modes is often required, which increases the computational cost of solving the ROM if a Galerkin projection is used. We mitigate this limitation by inserting a second layer of compression that uses a standard feed-forward neural network AE for the POD coefficients. In particular, we first employ the standard POD algorithm to efficiently decrease the dimensionality of the input field $u(\mathbf{x}, t)$ to near full-rank approximation as follows:

$$u(\mathbf{x}, t) = \sum_{k=1}^n a_k(t) \phi_k(\mathbf{x}), \quad (8)$$

where n is the number of modes that capture the required amount of information (e.g., RIC = 99.9%). Then, we train the AE to learn the underlying nonlinear correlations between the coefficients $\{a_k(t)\}_{k=1}^n$ to yield a further encoded latent space $\mathbf{z} \in \mathbb{R}^r$, where $r \ll n$. Next, we exploit a simple LSTM architecture to evolve \mathbf{z} forward. Finally, we utilize the decoder part of the trained AE to recover the full-rank POD expansion coefficients. With this, we provide an end-to-end Galerkin-free ROM that enables us to efficiently preserve the aforementioned merits of POD. Moreover,

the AE component of the NLPOD is trained onto the POD coefficients, so its architecture becomes independent of the dimensionality of the full order data. Therefore, training NLPOD is significantly more efficient than training CAE directly onto the input field data.

5 Results and discussion

We showcase the performance of the NLPOD framework using the Marsigli flow problem, where a fluid is placed into two partitions with different temperatures. The separating barrier is removed instantaneously to allow the fluids to slide over each other in a convection-dominated and buoyancy-driven manner. More details on the problem setting can be found in Refs. [48, 49]. In order to decrease the computational cost of training the CAE, in the present study, we consider a uniform Cartesian grid of 512×64 . We collect spatio-temporal data corresponding to different values of Reynolds number (Re) to provide a parameterized time-dependent setup. In particular, we store 200 snapshots at each value of $Re \in \{700, 900, 1100, 1300\}$ and use them for the offline training phase (i.e., $m = 800$). We then test the performance of GPOD, CAE, and NLPOD at $Re = 1000$, which is not included in the training data set. In order to improve the accuracy of AEs, we augment the bottleneck layer with the Reynolds number as an additional input feature. In other words, the decoder function for the CAE (defined in Eq. 7) becomes $\zeta : z \in \mathbb{R}^r \cup Re \in \mathbb{R} \rightarrow u(\mathbf{x}, t) \in \mathbb{R}^N$. For the NLPOD, it is rather defined as $\zeta : z \in \mathbb{R}^r \cup Re \in \mathbb{R} \rightarrow \mathbf{a} \in \mathbb{R}^n$. We train the AEs and LSTMs separately for the sake of simplicity and to facilitate the CAE and NLPOD combinations with different time series prediction tools.




Figure 2: The relative information content for different numbers of retained modes. While the first 2 POD modes have only a 76.76% RIC, the leading 74 POD modes have a 99.9% RIC of the spatiotemporal temperature field data. Therefore, the proposed autoencoder is built to learn a latent space from 74 POD coefficients.

We aim at approximating the temperature fields and assume that our target low order compression is $r = 2$. For the given problem, Figure 2 shows that 2 POD modes capture only 76.76% of the information, and at least 74 POD modes are required to attain a RIC value of 99.9% (i.e., rank- n approximation with $n = 74$). In the top panels of Figure 3, we plot the temperature fields at $t = 6$ and $t = 8$ from the FOM solver and the GPOD results for $r = 2$ and $r = 74$. Although the GPOD provides good accuracy with $r = 74$, the online computational cost is relatively high (around 10 minutes, compared to a fraction of a second for GPOD for $r = 2$). Therefore, we train AE onto the 74 POD expansion coefficients $\mathbf{a}(t) \in \mathbb{R}^{74}$ to learn a zipped latent space for $\mathbf{z}(t) \in \mathbb{R}^2$. Similarly, we train a CAE to reduce the dimensionality of the temperature field from $u(\mathbf{x}, t) \in \mathbb{R}^{512 \times 64}$ to a latent space for $\mathbf{z}(t) \in \mathbb{R}^2$.

The reconstructed temperature fields from NLPOD and CAE frameworks are shown in Figure 3, where we can see that the NLPOD predictions are more accurate than the CAE predictions. In order to provide an estimate of the model uncertainty in the output results, we utilize an ensemble-based framework, where we train different networks with different weight initializations (i.e., by changing the seed number). We find that the latent space revealed by CAE varies significantly with the initial weights specification. Therefore, we perform the training with 30 different initializations, and select 10 cases, corresponding to the highest performance on the validation data sets, to build our ensemble. We can observe the non-smoothness of the contour lines at different places in the CAE results, which indicates a non-physical behavior and/or insufficient latent space. Moreover, there is an evident problem in reconstructing the boundary conditions with CAE. Although this boundary condition issue can be reduced with proper padding methods, it requires additional case-specific customization of the neural network architecture. On the other hand, these complications are naturally mitigated in the NLPOD results since the restriction step and the prolongation process back to the FOM space

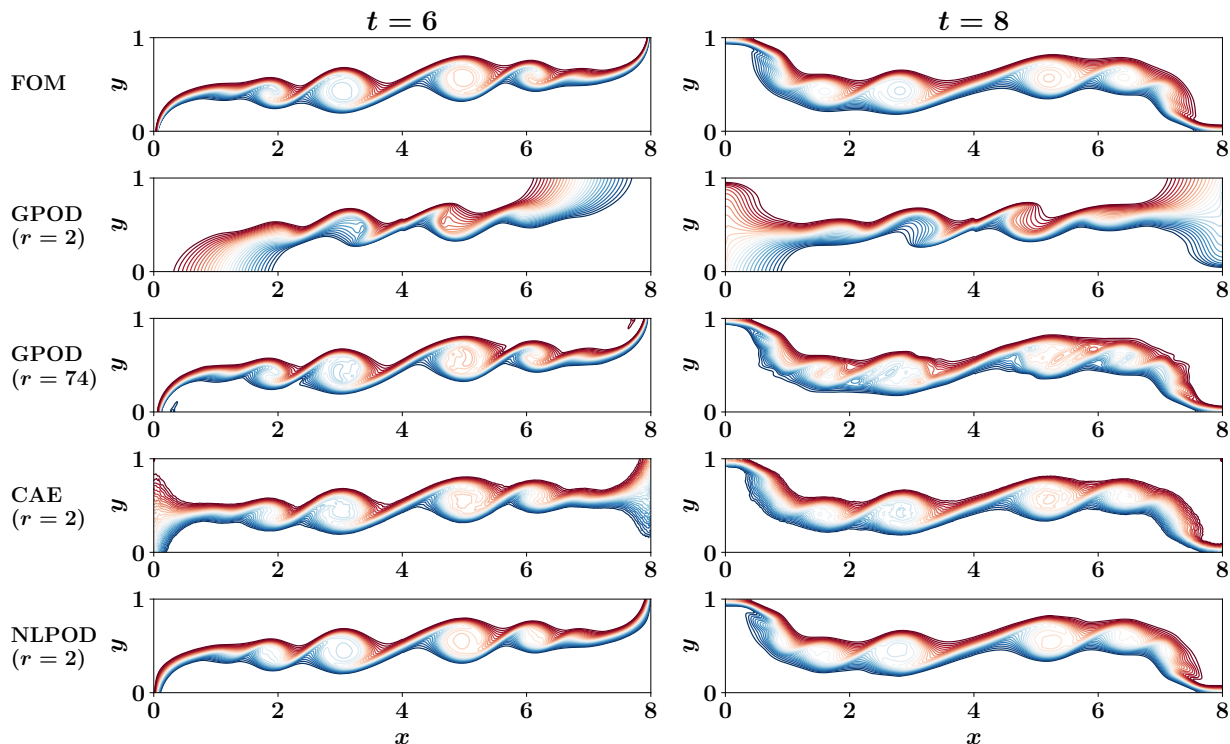


Figure 3: Predictions of temperature fields using FOM, compared to GPOD, CAE, and NLPOD modeling approaches. Training NLPOD is significantly faster than training CAE while the computational efficiency (running time) of NLPOD becomes in the same order with GROM ($r = 2$), and it is significantly faster than GROM ($r = 74$). Note that CAE can yield physically-inconsistent results and fails to capture the accurate boundary conditions. Outputs of CAE and NLPOD represent the ensemble-averaged predictions from 10 different neural network weight initializations.

are performed under the physically accurate POD umbrella. We also emphasize that in our experiments the training time of the CAE is significantly larger than the training time of the NLPOD (even using GPU capability), while the online deployment time for both cases is around half a minute. Thus, as shown in Table 1, NLPOD is both more accurate and more efficient than the CAE. Moreover, standard deviation data obtained from an ensemble of 10 different runs (i.e., runs trained using different weight initializations) clearly show that the uncertainty of the reconstruction is significantly lower in the NLPOD approach than in the CAE approach. Quantitatively, the standard deviation of the CAE reconstructions is about 55% of their mean value, while the standard deviation of the NLPOD reconstructions is only 4% of their mean value.

Table 1: Comparison of performance metrics for the NLPOD and CAE methodologies. Training times are reported only for autoencoder part since LSTM training and deployment times are similar in both cases. We perform numerical experiments using two different versions of tensorflow without* and with** GPU support. The data in the table correspond to the mean value \pm the standard deviation, obtained from an ensemble of 10 different runs.

Model	CAE	NLPOD	Gain [CAE/NLPOD]
Number of trainable parameters	2,194,499	40,516	54.16
CPU training time [s]*	1239.05 ± 32.30	14.16 ± 0.44	87.48
GPU training time [s]**	123.30 ± 2.09	56.2 ± 0.80	2.19
MSE of reconstruction (without LSTM)	$(7.18 \pm 3.93) \times 10^{-4}$	$(1.25 \pm 0.05) \times 10^{-4}$	5.75
MSE of prediction (with LSTM)	$(1.53 \pm 0.75) \times 10^{-3}$	$(3.61 \pm 0.74) \times 10^{-4}$	4.24

6 Conclusions

In conclusion, CAE technology has been shown to be a crucial dimensionality reduction approach for nonintrusive modeling and prediction of fluid flows. However, the CAE training time could be significant for high-dimensional systems often encountered in fluid dynamics applications. On the other hand, the intrusive projection-based approach has been the workhorse principle for reduced order modeling of fluid flows. The governing equations of fluid flows possess a quadratic nonlinearity. Therefore, POD-based methods, often combined with the Galerkin projection, yield a coupled dynamical system with computational complexity of $O(r^3)$, where r is the number of retained modes. To become computationally tractable, these projection-based methods truncate the reduced order representation ($r < n$), where n is the full rank of snapshot data matrix. Unfortunately, this truncation leads to a large projection error, especially for convection-dominated flows. In this work, we introduce a robust nonintrusive method that combines POD and multilayer perceptron autoencoders to generate a projection-error-free reduced order representation (with significantly reduced training time), and integrate this latent space with LSTM recurrent neural networks for its dynamics. Our results for modeling a lock-exchange density current problem show a substantial performance improvement over both nonintrusive convolutional AEs and intrusive Galerkin ROMs.

This is the first step toward building robust end-to-end Galerkin-free nonintrusive models for convection-dominated flows. Our future efforts will aim at extending this autoencoder-based framework to more complex higher-dimensional multiphysics problems in fluid dynamics. We highlight that advanced hyperparameter selection studies may be conducted for performance boost. Furthermore, instead of a deterministic autoencoder, one can utilize a variational autoencoder to perform probability distribution modeling of the latent space. It is also worth noting that the NLPOD approach is not restricted to a specific methodology of time series prediction. Although we introduce our results using LSTMs, NLPOD is equally applicable with other techniques [50], including sparse regression [51], Gaussian process regression [52], Seq2seq algorithms [53], temporal fusion transformers [54], and auto-regression methods. One can also replace LSTM with neural ODEs [55] to accommodate varying time stepping discretization. Lastly, the proposed NLPOD ROM framework can be seamlessly integrated with the physics-guided machine learning [56] to reduce the uncertainty of the nonintrusive model.

Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Award Number DE-SC0019290. O.S. gratefully acknowledges the Early Career Research Program (ECRP) support of the U.S. Department of Energy. O.S. also gratefully acknowledges the financial support of the National Science Foundation under Award Number DMS-2012255. T.I. acknowledges support through National Science Foundation Grant Number DMS-2012253.

Data availability

The data that supports the findings of this study are available within the article. Implementation details and Python scripts can be accessed from the Github repository [57].

References

- [1] Shady E Ahmed, Suraj Pawar, Omer San, Adil Rasheed, Traian Iliescu, and Bernd R Noack. On closures for reduced order models – a spectrum of first-principle to machine-learned avenues. *Physics of Fluids*, 33:091301, 2021.
- [2] Constantin Greif and Karsten Urban. Decay of the Kolmogorov N-width for wave problems. *Applied Mathematics Letters*, 96:216–222, 2019.
- [3] Shady E Ahmed and Omer San. Breaking the Kolmogorov barrier in model reduction of fluid flows. *Fluids*, 5(1):26, 2020.
- [4] Seddik M Djouadi. On the connection between balanced proper orthogonal decomposition, balanced truncation, and metric complexity theory for infinite dimensional systems. In *Proceedings of the 2010 American Control Conference*, pages 4911–4916. IEEE, 2010.
- [5] Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019.

- [6] Alec J Linot and Michael D Graham. Deep learning to discover and predict dynamics on an inertial manifold. *Physical Review E*, 101(6):062209, 2020.
- [7] David Sondak and Pavlos Protopapas. Learning a reduced basis of dynamical systems using an autoencoder. *Physical Review E*, 104(3):034202, 2021.
- [8] Rambod Mojjani and Maciej Balajewicz. Low-rank registration based manifolds for convection-dominated PDEs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 399–407, 2021.
- [9] Youngkyu Kim, Youngsoo Choi, David Widemann, and Tarek Zohdi. Efficient nonlinear manifold reduced order model. *arXiv preprint arXiv:2011.07727*, 2020.
- [10] Francisco J Gonzalez and Maciej Balajewicz. Learning low-dimensional feature dynamics using deep convolutional recurrent autoencoders. *arXiv preprint arXiv:1808.01346*, 2018.
- [11] Steffen Wiewel, Moritz Becher, and Nils Thuerey. Latent space physics: Towards learning the temporal evolution of fluid flow. *Computer Graphics Forum*, 38(2):71–82, 2019.
- [12] Arvind Mohan, Don Daniel, Michael Chertkov, and Daniel Livescu. Compressed convolutional LSTM: An efficient deep learning framework to model high fidelity 3D turbulence. *arXiv preprint arXiv:1903.00033*, 2019.
- [13] Saakaar Bhatnagar, Yaser Afshar, Shaowu Pan, Karthik Duraisamy, and Shailendra Kaushik. Prediction of aerodynamic flow fields using convolutional neural networks. *Computational Mechanics*, 64(2):525–545, 2019.
- [14] Takaaki Murata, Kai Fukami, and Koji Fukagata. Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *Journal of Fluid Mechanics*, 882:A13, 2020.
- [15] Kai Fukami, Taichi Nakamura, and Koji Fukagata. Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data. *Physics of Fluids*, 32(9):095110, 2020.
- [16] Jiayang Xu and Karthik Duraisamy. Multi-level convolutional autoencoder networks for parametric prediction of spatio-temporal dynamics. *Computer Methods in Applied Mechanics and Engineering*, 372:113379, 2020.
- [17] Kookjin Lee and Kevin T Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.
- [18] Lionel Agostini. Exploration and prediction of fluid dynamical systems using auto-encoder technology. *Physics of Fluids*, 32(6):067103, 2020.
- [19] Masaki Morimoto, Kai Fukami, Kai Zhang, Aditya G Nair, and Koji Fukagata. Convolutional neural networks for fluid flow analysis: toward effective metamodeling and low-dimensionalization. *arXiv preprint arXiv:2101.02535*, 2021.
- [20] Romit Maulik, Bethany Lusch, and Prasanna Balaprakash. Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders. *Physics of Fluids*, 33(3):037106, 2021.
- [21] Pin Wu, Siqun Gong, Kaikai Pan, Feng Qiu, Weibing Feng, and Christopher Pain. Reduced order model using convolutional auto-encoder with self-attention. *Physics of Fluids*, 33(7):077107, 2021.
- [22] R Fu, D Xiao, IM Navon, and C Wang. A data driven reduced order model of fluid flow by auto-encoder and self-attention deep learning methods. *arXiv preprint arXiv:2109.02126*, 2021.
- [23] Gianmarco Aversano, Aurelie Bellemans, Zhiyi Li, Axel Coussement, Olivier Gicquel, and Alessandro Parente. Application of reduced-order models based on PCA & Kriging for the development of digital twins of reacting flow applications. *Computers & Chemical Engineering*, 121:422–441, 2019.
- [24] Suraj Pawar, SM Rahman, H Vaddireddy, Omer San, Adil Rasheed, and Prakash Vedula. A deep learning enabler for nonintrusive reduced order modeling of fluid flows. *Physics of Fluids*, 31(8):085101, 2019.
- [25] Sk M Rahman, Suraj Pawar, Omer San, Adil Rasheed, and Traian Iliescu. Nonintrusive reduced order modeling framework for quasigeostrophic turbulence. *Physical Review E*, 100(5):053306, 2019.
- [26] Jian Yu, Chao Yan, and Mengwu Guo. Non-intrusive reduced-order modeling for fluid problems: A brief review. *Journal of Aerospace Engineering*, 233(16):5896–5912, 2019.
- [27] S Dutta, P Rivera-Casillas, OM Cecil, and MW Farthing. pyNIROM—A suite of Python modules for non-intrusive reduced order modeling of time-dependent problems. *Software Impacts*, 10:100129, 2021.
- [28] Suraj Pawar, Omer San, Adil Rasheed, and Prakash Vedula. A priori analysis on deep learning of subgrid-scale parameterizations for Kraichnan turbulence. *Theoretical and Computational Fluid Dynamics*, 34(4):429–455, 2020.
- [29] Mark A Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2):233–243, 1991.

- [30] Adam H Monahan. Nonlinear principal component analysis by neural networks: Theory and application to the Lorenz system. *Journal of Climate*, 13(4):821–835, 2000.
- [31] William W Hsieh. Nonlinear principal component analysis by neural networks. *Tellus A*, 53(5):599–615, 2001.
- [32] William W Hsieh. Nonlinear principal component analysis of noisy data. *Neural Networks*, 20(4):434–443, 2007.
- [33] Samuel E Otto and Clarence W Rowley. Linearly recurrent autoencoder networks for learning dynamics. *SIAM Journal on Applied Dynamical Systems*, 18(1):558–593, 2019.
- [34] N Benjamin Erichson, Michael Muehlebach, and Michael W Mahoney. Physics-informed autoencoders for Lyapunov-stable fluid flow prediction. *arXiv preprint arXiv:1905.10866*, 2019.
- [35] Tomoharu Iwata and Yoshinobu Kawahara. Neural dynamic mode decomposition for end-to-end modeling of nonlinear dynamics. *arXiv preprint arXiv:2012.06191*, 2020.
- [36] Shaowu Pan and Karthik Duraisamy. Physics-informed probabilistic learning of linear embeddings of nonlinear dynamics with guaranteed stability. *SIAM Journal on Applied Dynamical Systems*, 19(1):480–509, 2020.
- [37] Shivakanth Chary Puligilla and Balaji Jayaraman. Deep multilayer convolution frameworks for data-driven learning of fluid flow dynamics. In *2018 Fluid Dynamics Conference*, page 3091, 2018.
- [38] Kevin T Carlberg, Antony Jameson, Mykel J Kochenderfer, Jeremy Morton, Liqian Peng, and Freddie D Witherden. Recovering missing CFD data for high-order discretizations using deep neural networks and dynamics learning. *Journal of Computational Physics*, 395:105–124, 2019.
- [39] César Quilodrán Casas, Rossella Arcucci, and Yike Guo. Urban air pollution forecasts generated from latent space representation. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.
- [40] Toby RF Phillips, Claire E Heaney, Paul N Smith, and Christopher C Pain. An autoencoder-based reduced-order model for eigenvalue problems with application to neutron diffusion. *International Journal for Numerical Methods in Engineering*, 122(15):3780–3811, 2021.
- [41] Robert L Kosut, Tak-San Ho, and Herschel Rabitz. Quantum system compression: A Hamiltonian guided walk through Hilbert space. *Physical Review A*, 103(1):012406, 2021.
- [42] Adil Rasheed, Omer San, and Trond Kvamsdal. Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 8:21980–22012, 2020.
- [43] Shady E Ahmed, Omer San, Adil Rasheed, and Traian Iliescu. A long short-term memory embedding for hybrid uplifted reduced order models. *Physica D: Nonlinear Phenomena*, 409:132471, 2020.
- [44] Arvind T Mohan, Nicholas Lubbers, Daniel Livescu, and Michael Chertkov. Embedding hard physical constraints in neural network coarse-graining of 3D turbulence. *arXiv preprint arXiv:2002.00021*, 2020.
- [45] Han Gao, Luning Sun, and Jian-Xun Wang. Phygeonet: physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state pdes on irregular domain. *Journal of Computational Physics*, 428:110079, 2021.
- [46] Muruhan Rathinam and Linda R Petzold. A new look at proper orthogonal decomposition. *SIAM Journal on Numerical Analysis*, 41(5):1893–1925, 2003.
- [47] Xuping Xie, Peter J Nolan, Shane D Ross, Changhong Mou, and Traian Iliescu. Lagrangian reduced order modeling using finite time Lyapunov exponents. *Fluids*, 5(4):189, 2020.
- [48] Omer San and Jeff Borggaard. Principal interval decomposition framework for POD reduced-order modeling of convective Boussinesq flows. *International Journal for Numerical Methods in Fluids*, 78(1):37–62, 2015.
- [49] Shady E Ahmed, Omer San, Kursat Kara, Rami Younis, and Adil Rasheed. Multifidelity computing for coupling full and reduced order models. *PLoS ONE*, 16(2):e0246092, 2021.
- [50] Zhe Bai and Liqian Peng. Non-intrusive nonlinear model reduction via machine learning approximations to low-dimensional operators. *arXiv preprint arXiv:2106.09658*, 2021.
- [51] Kai Fukami, Takaaki Murata, Kai Zhang, and Koji Fukagata. Sparse identification of nonlinear dynamics with low-dimensionalized flow representations. *Journal of Fluid Mechanics*, 926, 2021.
- [52] Romit Maulik, Themistoklis Botsas, Nesar Ramachandra, Lachlan R Mason, and Indranil Pan. Latent-space time evolution of non-intrusive reduced-order models using gaussian process emulation. *Physica D: Nonlinear Phenomena*, 416:132797, 2021.
- [53] Sandeep B Reddy, Allan Ross Magee, Rajeev K Jaiman, J Liu, W Xu, A Choudhary, and AA Hussain. Reduced order model for unsteady fluid flows via recurrent neural networks. In *International Conference on Offshore Mechanics and Arctic Engineering*, volume 58776, page V002T08A007. American Society of Mechanical Engineers, 2019.

- [54] Bryan Lim, Sercan Ö Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):1748–1764, 2021.
- [55] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [56] Suraj Pawar, Omer San, Aditya Nair, Adil Rasheed, and Trond Kvamsdal. Model fusion with physics-guided machine learning: Projection-based reduced-order modeling. *Physics of Fluids*, 33(6):067123, 2021.
- [57] Shady E Ahmed. NLPOD. <https://github.com/Shady-Ahmed/NLPOD>, 2021.