







RESEARCH ARTICLE | JULY 16 2024


## Reduced-order modeling of unsteady fluid flow using neural network ensembles

Rakesh Halder  ; Mohammadmehdi Ataei ; Hesam Salehipour ; Krzysztof Fidkowski ; Kevin Maki 



*Physics of Fluids* 36, 077140 (2024)

<https://doi.org/10.1063/5.0207978>



## Physics of Fluids

Special Topic:

Flow and Lipid Nanoparticles

Guest Editors: Richard Braatz and Mona Kanso

[Submit Today!](#)

# Reduced-order modeling of unsteady fluid flow using neural network ensembles

Cite as: Phys. Fluids **36**, 077140 (2024); doi: [10.1063/5.0207978](https://doi.org/10.1063/5.0207978)

Submitted: 12 March 2024 · Accepted: 22 June 2024 ·

Published Online: 16 July 2024



View Online



Export Citation



CrossMark

Rakesh Halder,<sup>1,2,a)</sup>  Mohammadmehdi Ataei,<sup>2</sup>  Hesam Salehipour,<sup>2</sup>  Krzysztof Fidkowski,<sup>1</sup>   
and Kevin Maki<sup>3</sup> 

## AFFILIATIONS

<sup>1</sup>Department of Aerospace Engineering, University of Michigan, Ann Arbor, Michigan, USA 48109

<sup>2</sup>Autodesk Research, Toronto, Ontario, Canada M5G 1M1

<sup>3</sup>Department of Naval Architecture and Marine Engineering, University of Michigan, Ann Arbor, Michigan, USA 48109

<sup>a)</sup> Author to whom correspondence should be addressed: [rhald@umich.edu](mailto:rhald@umich.edu)

## ABSTRACT

The use of deep learning has become increasingly popular in reduced-order models (ROMs) to obtain low-dimensional representations of full-order models. Convolutional autoencoders (CAEs) are often used to this end as they are adept at handling data that are spatially distributed, including solutions to partial differential equations. When applied to unsteady physics problems, ROMs also require a model for time-series prediction of the low-dimensional latent variables. Long short-term memory (LSTM) networks, a type of recurrent neural network useful for modeling sequential data, are frequently employed in data-driven ROMs for autoregressive time-series prediction. When making predictions at unseen design points over long time horizons, error propagation is a frequently encountered issue, where errors made early on can compound over time and lead to large inaccuracies. In this work, we propose using bagging, a commonly used ensemble learning technique, to develop a fully data-driven ROM framework referred to as the CAE-eLSTM ROM that uses CAEs for spatial reconstruction of the full-order model and LSTM ensembles for time-series prediction. When applied to two unsteady fluid dynamics problems, our results show that the presented framework effectively reduces error propagation and leads to more accurate time-series prediction of latent variables at unseen points.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0207978>

## I. INTRODUCTION

Physics-based numerical simulation has become an indispensable tool in engineering and scientific applications, allowing for accurate representations of complex physical phenomena. Physical simulations are formulated as sets of governing equations, typically in the form of parameterized partial differential equations (PDEs) that are discretized over a computational domain. Simulations often involve a set of design parameters  $\mu$  that govern aspects such as boundary conditions, the geometry of the computational domain, and physical properties. In tasks like design optimization, where numerous simulations are run for different designs, it is important to achieve high accuracy or fidelity. However, high-fidelity simulation comes at a large computational cost, and this can lead to a bottleneck in the design optimization process.

Employing reduced-order models (ROMs) is a prevalent strategy to mitigate this high computational cost.<sup>1</sup> By utilizing a limited set of training data from computed high-fidelity simulations, ROMs construct a low-dimensional surrogate model that can be evaluated in real-time and provide accurate full-order model (FOM) approximations at unseen designs within the parameter space covered by the

training data. ROMs significantly reduce the degrees of freedom of the FOM through the use of a low-dimensional embedding that can be effectively mapped back to the full-order state, leading to a substantial decrease in the computational cost. ROMs consist of two stages: an off-line stage, which is computationally intensive, involving the computation of high-fidelity solutions by solving the FOM to generate data snapshots and train the low-dimensional surrogate model, and an online stage, where the model is used to approximate solutions at desired points. Running the FOM requires numerically solving the governing equations over the computational domain, which comes at a much larger computational cost than the ROM, which represents the FOM using a small number of variables.

Most ROMs utilize the proper orthogonal decomposition,<sup>2</sup> a linear method that uses the singular value decomposition (SVD) to obtain a low-rank subspace consisting of a number of linearly independent basis vectors. A linear combination of these basis vectors is computed using a set of expansion coefficients to approximate full-order states within the solution space. The POD basis vectors are interpretable, allowing for visualization of dominant physical features.

Although POD is widely used, it encounters difficulties when applied to highly nonlinear problems, often requiring a large number of basis vectors to provide reasonable accuracy.<sup>3</sup>

More recently, machine learning and artificial intelligence (AI) methods have been used to provide nonlinear mappings between the low-dimensional embedding and high-fidelity solution space. In particular, deep learning<sup>4</sup> approaches have been used to develop ROMs that provide efficient nonlinear manifolds of physical systems. Convolutional autoencoders (CAEs), a type of artificial neural network, have been used to build ROMs and shown to outperform POD-based methods.<sup>5,6</sup> Artificial neural networks that use convolutional layers efficiently learn structures and patterns present in spatially distributed data, including the solutions to PDEs. Autoencoders consist of two individual neural networks: an encoder, which maps high-dimensional inputs to a low-dimensional latent space, and a decoder, which maps the low-dimensional latent space to a reconstruction of the high-dimensional input. Although autoencoders can effectively learn nonlinear relationships, they lack interpretability. Recent works have used variational autoencoders<sup>7,8</sup> to incorporate interpretable nonlinear modes into ROMs, although this results in lower reconstruction accuracy when compared to vanilla autoencoders. Another approach for simulating systems governed by PDEs at a reduced cost is the use of neural operators,<sup>9,10</sup> which implicitly learn the governing equations using deep neural networks. However, this approach typically requires very large amounts of training data to cover parameter spaces effectively, limiting their use in design optimization.

ROMs are categorized as either non-intrusive or projection-based; non-intrusive ROMs use a fully data-driven approach, while projection-based ROMs incorporate the governing equations to solve a low-dimensional version of the FOM. While projection-based ROMs can provide better accuracy and more physically consistent results, they incur a larger computational cost and often exhibit stability issues<sup>11</sup> that inhibit convergence. Additionally, they are generally not portable between different solvers.

ROMs are developed for both steady and unsteady physics problems. The first approach involves computing a single point estimate of the low-dimensional embedding, while the latter involves making predictions over a prescribed time horizon. Unsteady non-intrusive ROMs combine either POD or CAE for spatial reconstruction of full-order states with a model used to make time-series predictions of the low-dimensional embeddings. Deep learning methods are popular for this as well, including long short-term memory (LSTM) neural networks<sup>12</sup> and transformer neural networks.<sup>8</sup> Both LSTMs and transformers are powerful models for handling data that are sequential in nature, such as time-series data. Transformers utilize an architecture that is much more complex than the one LSTMs use, which can lead to better performance for complex time-series problems and other tasks such as developing large language models.<sup>13</sup> However, this comes at an increased cost for both training and inference when compared to LSTMs, which use a significantly lower number of parameters. The use of LSTMs for time-series forecasting is well-established,<sup>14,15</sup> making them a popular choice.

When making time-series predictions at unseen datasets over a long time horizon, error propagation is a common issue. Errors made in early predictions can accumulate and compound over time, leading to substantial inaccuracies. Unseen datasets are particularly suspect to this, as the model may not account for shifting data patterns.

Additionally, there is no feedback from previously seen data to correct errors. This phenomenon is commonly observed in data-driven computational fluid dynamics (CFD) applications.<sup>16,17</sup> As a result, most studies in the literature focus on applying unsteady non-intrusive ROMs to single-parameter problems,<sup>8,12,18</sup> where the ROM is both trained on and used for a single design. We have identified two previously published studies that combine CAEs and LSTMs for non-intrusive ROMs and apply them to unseen designs.<sup>19,20</sup> The work by Maulik *et al.* does not mention the error propagation issue, while the work by Hasegawa *et al.* presents results of low accuracy. A recent paper<sup>21</sup> by Jeon *et al.* introduced a hybrid AI-CFD method using flow residuals to address the issue of error propagation. However, this method is intrusive and requires the ROM to have access to the CFD solver for computing the residuals, which leads to a more computationally intensive online stage.

In this work, we propose the use of ensemble learning,<sup>22</sup> a machine learning technique for improving the stability and lowering the variance of predictive models. Ensemble learning involves combining multiple base models, referred to as weak learners, to create a composite model that offers greater accuracy. To this end, we employ bootstrap aggregating (bagging) as the ensemble learning method for the temporal portion of the ROM, which involves training the weak learners on subsets of the dataset chosen randomly through sampling with replacement. The fully data-driven framework is referred to as the CAE-eLSTM ROM, and our results show that using ensembles leads to significantly improved stability and predictive performance when applied to two unsteady, incompressible, laminar fluid dynamics problems using different CFD solvers.

## II. METHODS

In this section, we give an overview of both CAEs and LSTMs and how they are combined to develop an unsteady ROM using bagging as an ensemble learning method. CAEs are used to both provide a low-dimensional latent representation of the solution space and provide spatial reconstructions of full-order states. A temporal forecasting model of the latent variables is also required when using unsteady ROMs. LSTMs, a type of recurrent neural network, are used in this work. Although we do not provide direct comparisons to POD-based ROMs, a brief introduction to POD is given in this section. Comparisons of reconstruction errors using POD and CAE are given in [Appendix A](#).

### A. Proper orthogonal decomposition

Using training data obtained from a set of  $n$  solution snapshots calculated at chosen design points in the parameter space, a snapshot matrix,  $\mathbf{X} \in \mathbb{R}^{N \times n}$ , can be assembled.

$$\mathbf{X} \in \mathbb{R}^{N \times n} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n] = [\mathbf{x}(\boldsymbol{\mu}^1), \mathbf{x}(\boldsymbol{\mu}^2), \dots, \mathbf{x}(\boldsymbol{\mu}^n)]. \quad (1)$$

A subspace  $\mathcal{V}$  associated with  $\mathbf{X}$  exists such that  $\mathcal{V} = \text{span}(\mathbf{X})$ . It is assumed that  $\mathcal{V}$  provides a good approximation of the solutions for  $\boldsymbol{\mu}$  within the parameter space if there is a sufficient variety of training data in  $\mathbf{X}$ . A rank  $k$  set of orthonormal basis vectors  $[\boldsymbol{\psi}^1, \boldsymbol{\psi}^2, \dots, \boldsymbol{\psi}^k] \in \mathbb{R}^N$ , where  $k \ll N$ , is associated with  $\mathcal{V}$  such that each solution  $\mathbf{x}^i$  in  $\mathbf{X}$  can be reconstructed as a linear combination of the basis vectors

$$\mathbf{x}^i \approx a_1^i \boldsymbol{\psi}^1 + a_2^i \boldsymbol{\psi}^2 \dots + a_k^i \boldsymbol{\psi}^k, \quad (2)$$

where  $\mathbf{a}^i$  are the low-dimensional expansion coefficients. The truncated singular value decomposition (SVD) of  $\mathbf{X}$  decomposes the matrix into two orthonormal matrices  $\mathbf{U} \in \mathbb{R}^{N \times n}$  and  $\mathbf{V} \in \mathbb{R}^{n \times n}$ , and a diagonal matrix  $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$  such that

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T. \quad (3)$$

$\mathbf{U}$  contains a set of  $n$  left singular vectors that span the column space of  $\mathbf{X}$ ,  $\mathbf{V}$  contains a set of  $n$  right singular vectors that span the row space of  $\mathbf{X}$ , and  $\text{diag}(\mathbf{\Sigma}) \in \mathbb{R}^n = [\sigma_1, \sigma_2, \dots, \sigma_n]$  contains the singular values in the decreasing order,  $\sigma_1 \geq \dots \geq \sigma_n \geq 0$ . The first  $k$  vectors of  $\mathbf{U}$  form the POD basis,  $\mathbf{\Psi} \in \mathbb{R}^{N \times k} = [\psi^1, \psi^2, \dots, \psi^k]$ . The singular values are a measure of the amount of information represented by each vector. Often, the singular values decay rapidly, and only the first  $k$  singular vectors are chosen to form the POD basis to preserve the most important features of the solution space. The relative information content  $E$  of the subspace is used in practice to choose a value of  $k$ ,

$$E(k) = \frac{\sum_{j=1}^k \sigma_j^2}{\sum_{j=1}^n \sigma_j^2}, \quad (4)$$

and  $k$  is chosen such that  $E(k) \geq \gamma$ , where  $\gamma \in [0,1]$  is usually set to a value  $\gamma \geq 0.95$ .<sup>23</sup> Approximations of full-order solutions at unseen design parameters  $\mathbf{x}(\mu^*)$  are obtained as follows:

$$\mathbf{x}(\mu^*) \approx \mathbf{\Psi}\mathbf{a}^* = a_1^*\psi^1 + a_2^*\psi^2 \dots + a_k^*\psi^k. \quad (5)$$

### B. Convolutional autoencoders

Convolutional neural networks were initially developed for computer vision applications, where they have been shown to vastly outperform traditional statistical image processing techniques.<sup>24,25</sup> They have also demonstrated high predictive performance when used in ROMs,<sup>5,26</sup> highlighting their versatility. Autoencoders are a type of feedforward neural network that aim to accurately reconstruct inputs in the output layer,  $g : \mathbf{x} \rightarrow \tilde{\mathbf{x}}$  where  $\mathbf{x} \approx \tilde{\mathbf{x}}$ . Autoencoders are composed of two individual feedforward neural networks. The encoder  $g_{\text{enc}} : \mathbb{R}^N \rightarrow \mathbb{R}^k$  where  $k \ll N$  maps a high-dimensional input  $\mathbf{x}$  into the low-dimensional latent space  $\mathbf{a}$ , and the decoder  $g_{\text{dec}} : \mathbb{R}^k \rightarrow \mathbb{R}^N$  maps the latent variables back to an approximation of the high-dimensional input  $\tilde{\mathbf{x}}$ . The combination of the two results in

$$g : \tilde{\mathbf{x}} = g_{\text{dec}} \circ g_{\text{enc}}(\mathbf{x}). \quad (6)$$

Deep neural networks provide nonlinear function maps through the use of activation functions. During training, the neural network parameters are tuned using backpropagation,<sup>27</sup> an algorithm utilizing automatic differentiation that aims to minimize a differentiable loss function that measures the discrepancy between the computed and true outputs. This allows for discovering highly complex and abstract functional relationships that do not follow any strong model assumptions. The number of trainable parameters in vanilla artificial neural networks grows large with the number of hidden layers and neurons and can lead to a very large computational training cost.

Convolutional neural networks effectively implement parameter sharing to limit the total number of trainable parameters in the network, where rather than weight combinations existing for each pair of neurons between layers, multiple neurons share a single weight. Convolutional layers use a number of filters that convolve over the input data, with each filter having its own set of weights. Pooling layers are also used in

convolutional networks to summarize the features in input layers through operations including averaging and maximization. The input layer to a convolutional neural network is composed of a number of channels, with each representing a different state. In images, this is generally the levels of red, green, and blue present in each pixel, while for physical simulation data the states can represent components of the normalized velocity, pressure, density, etc. Once an autoencoder is sufficiently trained and  $g(\mathbf{x}) \approx \mathbf{x}$  for all inputs over the training dataset  $\mathbf{X}$ , the corresponding latent variables  $\mathbf{a}$  can be passed to the decoder  $g_{\text{dec}}(\mathbf{a})$  to obtain accurate approximations  $\tilde{\mathbf{x}}$  for all data in  $\mathbf{X}$ . High-dimensional data existing outside of the training set  $\mathbf{x}^*$  may also be well-approximated if an accurate approximation of the latent variables  $\mathbf{a}^*$  is obtained. Figure 1 shows an example of the encoder section of a convolutional autoencoder, which is composed of convolutional, pooling, and fully connected layers. A similar architecture in reverse would form the decoder.

The input and output layers of CAEs are often structured as two-dimensional states within each channel when used for two-dimensional physical simulations. To accommodate this structure, training data needs to be reshaped before being input into the network through the use of a reshape operator  $\mathbf{R}$ ,

$$\mathbf{R} : \mathbb{R}^{N \times n_c} \rightarrow \mathbb{R}^{n_y \times n_x \times n_c}, \quad (7)$$

where  $n_y$  refers to the number of data points in the vertical direction and  $n_x$  refers to the number of data points in the horizontal direction. The reshape operator is applied to each separate state that occupies the  $n_c$  input channels.

### C. Long short-term memory (LSTM) neural networks

Traditional recurrent neural networks are well-suited for handling sequential data; however, they have difficulties learning and capturing long-range dependencies. A primary reason for this is the vanishing gradient problem,<sup>28</sup> which occurs when gradients shrink significantly as they are back-propagated through time, which can cause the model to forget early sequence information. LSTMs were introduced<sup>29</sup> to address this issue by using a more complex network architecture that includes a series of gates which effectively control the flow of information. An individual LSTM cell consists of a cell state  $\mathbf{c}_t$  which acts as an internal memory, and a hidden state  $\mathbf{h}_t$ , which serves as an output. The cell state is responsible for selectively retaining or forgetting information from previous steps, while the hidden state represents the model's cumulative output. LSTM cells contain a forget gate  $\mathbf{f}_t$ , input gate  $\mathbf{i}_t$ , and output gate  $\mathbf{o}_t$ . Forget gates determine what information from previous states should be forgotten. Input gates decide what new pieces of information should be stored in the current state. Output gates use the current and previous cell states to determine what information is retained in the model. Using this architecture, the model is able to retain information that is relevant for long-range sequential dependencies and discard information that becomes irrelevant over time, making LSTMs a popular choice for sequential modeling. For a given input  $\mathbf{a}_t \in \mathbb{R}^k$ , the LSTM equations are given as follows:

$$\begin{aligned} \mathbf{f}_t &= \sigma(\mathcal{F}_f(\mathbf{a}_t)), \\ \mathbf{i}_t &= \sigma(\mathcal{F}_i(\mathbf{a}_t)), \\ \mathbf{o}_t &= \sigma(\mathcal{F}_o(\mathbf{a}_t)), \\ \tilde{\mathbf{c}}_t &= \tanh(\mathcal{F}_c(\mathbf{a}_t)), \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t, \\ \mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t), \end{aligned} \quad (8)$$




FIG. 1. Architecture of the encoder of a convolutional autoencoder (CAE) consisting of convolutional, pooling, and fully connected layers.

where  $\sigma$  is the sigmoid function,  $\tanh$  is the hyperbolic tangent function, and  $\odot$  is the Hadamard product.  $\mathcal{F}$  is a linear function of the weights  $W_a \in \mathbb{R}^{n_h}$  and  $W_h \in \mathbb{R}^{n_h}$  and biases  $b \in \mathbb{R}^{n_h}$ , where  $n_h$  is the number of neurons in the hidden layer of each LSTM cell and is given as follows:

$$\mathcal{F} = W_a a_t + W_h h_{t-1} + b. \tag{9}$$

A unique set of weights and biases belongs to each gate or cell state. When using LSTMs for time-series prediction, an autoregressive prediction is used, often referred to as the sliding window approach. For a given input sequence containing multivariate data for  $w$  timesteps, the next step in the sequence is predicted. When making the next prediction, the current prediction is incorporated into the input sequence by removing the first element and shifting the rest to the prior position. Eventually, the model inputs will consist of only previously computed predictions, which makes model performance sensitive to error propagation. The input sequence length, or window size  $w$ , is an important hyperparameter to consider when training LSTMs. Although LSTMs are designed to effectively learn long-range dependencies, using an input sequence length that is too long can lead to the inclusion of outdated and irrelevant information. On the other hand, using an input sequence length that is too small can ignore important long-range information. The optimal choice is highly problem-dependent; similar problems found in the literature can be consulted, or methods such as cross-validation can be used to find an optimal value. A diagram of a single-layer LSTM making a prediction one time step ahead is shown in Fig. 2.

#### D. Ensemble learning

While LSTMs are powerful models for time-series prediction, the quality of predictions over long time horizons is highly sensitive to the weights and biases of the model. Additionally, the total number of weights and biases is usually very large, which makes finding an optimal configuration very difficult, even when using state of the art optimization algorithms.

To mitigate the issue of high model variance, ensemble learning is a commonly used approach. By leveraging multiple base models that exhibit high variance, ensemble methods significantly reduce errors and improve robustness. Boosting and bootstrap aggregating (bagging) are the two main types of ensemble learning methods. Boosting<sup>30</sup> trains individual models sequentially, where each subsequent weak learner focuses on correcting prediction errors from the previous ones. While boosting is widely used and offers good performance, the base models must be trained iteratively, which incurs a large computational cost, especially in the context of neural networks.

Bagging<sup>31</sup> is an algorithm that consists of two stages: bootstrapping and aggregation. Bootstrapping is a resampling technique where multiple random subsets of the dataset, chosen through sampling with

replacement, are constructed. The subsets of the data typically contain the same number of data points as the original dataset. This approach leads to individual data points being present multiple times in the individual subsets. An example of bootstrapping is shown in Fig. 3, where the bootstrapped datasets contain more or less instances of the original data points. Multiple base models are trained individually on each of the bootstrapped datasets, which can be done in parallel. The aggregation stage creates an ensemble model by taking an average of the predictions given by each weak learner. Given  $m$  weak learners  $f_i$ , the aggregate prediction  $\bar{f}$  is given as follows:

$$\bar{f} = \frac{1}{m} \sum_i^m f_i. \tag{10}$$

Bagging reduces the issue of error propagation by averaging the errors of multiple weak learners, which will become increasingly small as the number of learners grows, leading to much lower variance. Bagging also offers a considerable gain in robustness and generalization by leveraging the diversity of the weak learners to increase the capability of handling varying patterns outside of the training data. As an increasing number of weak learners are added, the reduction in




FIG. 2. Diagram of a single-layer LSTM neural network making a prediction one time step ahead.




FIG. 3. An example of bootstrapping, where random subsets of the original dataset are chosen through sampling with replacement.

variance plateaus, leading to diminishing returns in performance gains. After a certain point, the computational cost of adding more weak learners outweighs the marginal improvement in predictive performance. As the complexity of the problem increases, the number of weak learners required to make accurate predictions also generally increases. While theoretical considerations can be made to analyze the tradeoffs between bias, variance, and covariance to determine the optimal number of weak learners, this number is difficult to choose *a priori*. A reasonable choice is based on both the problem's complexity and similar problems found in the literature, which are not yet present in our case and are established in this work.

### E. CAE-eLSTM ROM

This section describes the non-intrusive ROM framework combining convolutional autoencoders for spatial reconstruction and LSTM ensembles utilizing bagging for temporal prediction, which we will refer to as a CAE-eLSTM ROM. The offline stage is computationally intensive, requiring multiple solves of the FOM in addition to training multiple neural networks. First, solutions to the FOM for designs  $\mu$  in the training set of design parameters  $\mathcal{U}_{\text{train}} \in \mathbb{R}^{n \times p}$  are obtained for all timesteps and are assembled into a matrix  $\mathbf{X}$ . Next, a convolutional autoencoder with latent dimension  $k$  is trained on  $\mathbf{X}$  for a sufficient number of epochs such that inputs are accurately reconstructed. Then, the expansion coefficients for the training data  $\mathbf{A}_{\text{train}}$  are computed by feeding the full-order states from  $\mathbf{X}$  through the encoder  $g_{\text{enc}}$ . Sequences of length  $w$  are then generated from  $\mathbf{A}_{\text{train}}$  and  $m$  LSTMs using the same set of initial weights and biases are trained on individual data subsets chosen randomly through sampling with replacement. PyTorch,<sup>32</sup> an open source deep learning framework available for Python, is used to implement the autoencoder and LSTM using default weight and bias initialization settings.

A schematic describing the online stage of the ROM is shown in Fig. 4. The online stage involves executing the ROM to compute approximate solutions at a point  $\mu^*$ . First, the FOM is run for  $T_i$  timesteps, as an initial sequence of latent variables is required for the LSTM ensemble. The computed full-order states are fed through the encoder to obtain the initial sequence of latent variables  $\mathbf{A}_{T_i}^* \in \mathbb{R}^{k \times w}$ . The FOM must be run for at least  $T_i \geq w$  timesteps, and potentially longer depending on the how useful simulation data from the initial timesteps are for the ROM. The latent variables for the rest of the prediction horizon are calculated autoregressively using the LSTM ensemble by taking an average of the predictions from the individual weak learners.

After the predicted latent variables are found for each time step  $t$ , they are fed through the decoder to obtain approximate full-order fields. The online and offline stages are outlined in Algorithm 1.

ALGORITHM 1. Offline and online stages of CAE-eLSTM ROM.

```

1: function CAE_eLSTM_OFFLINE( $\mathcal{U}_{\text{train}}, k, m, w$ )
2:   Compute high-fidelity solutions for  $\mu \in \mathcal{U}_{\text{train}}$  by solving FOM and assemble into  $\mathbf{X}$ .
3:   Train convolutional autoencoder with latent dimension  $k$  on  $\mathbf{X}$ .
4:   Calculate expansion coefficients for training data  $\mathbf{A}_{\text{train}} = g_{\text{enc}}(\mathbf{X})$ .
5:   Train  $m$  LSTMs  $\mathcal{E} = [\text{LSTM}_1(\mathbf{A}), \text{LSTM}_2(\mathbf{A}), \dots, \text{LSTM}_m(\mathbf{A})]$  using a window size of  $w$  on sequences sampled randomly with replacement from  $\mathbf{A}_{\text{train}}$ .
6:   return ( $g_{\text{enc}}, g_{\text{dec}}, \mathcal{E}$ )
7: end function

1: function CAE_eLSTM_ONLINE( $\mu^*, g_{\text{enc}}, g_{\text{dec}}, \mathcal{E}$ )
2:   Run FOM for  $T_i$  timesteps at  $\mu^*$  and compute the expansion coefficients  $\mathbf{A}_{T_i}^* = g_{\text{enc}}([X_{T_i}^*, X_{T_i+1}^*, \dots, X_{T_i+w}^*])$ .
3:   for  $t \in \{T_i, T_i+1, \dots, T-1\}$  do
4:     Compute  $\mathbf{a}_{t+1}^* = \mathcal{E}(\mathbf{A}_t^*)$  by using an average of the  $m$  individual LSTM predictions.
5:     Incorporate  $\mathbf{a}_{t+1}^*$  into the current window  $\mathbf{A}_t^*$ .
6:   end for
7:   Predict full-order fields  $\tilde{\mathbf{X}}^* = g_{\text{dec}}(\mathbf{A}^*)$ 
8:   return  $\tilde{\mathbf{X}}^*$ 
9: end function

```

When using a ROM with a single LSTM as in previous works,<sup>19,20</sup> the LSTM is trained on the entire time-series dataset and is used alone for time-series prediction. When compared to these works, the main difference in our presented ROM framework is the use of an ensemble model for time-series prediction. Although this results in an increased offline training cost, using an ensemble allows for better accuracy and stability over long prediction horizons. While the use of an ensemble LSTM is simple, it can greatly improve ROM performance without requiring additional training data.

21 July 2024 19:03:28

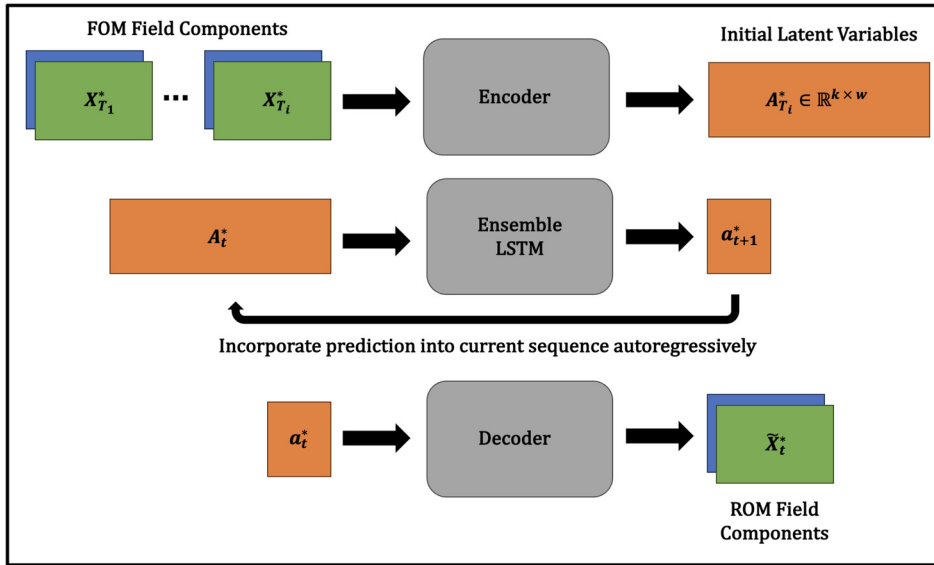


FIG. 4. Schematic of the online stage of the CAE-eLSTM ROM.

### III. RESULTS

The two test cases used to demonstrate the performance of the CAE-eLSTM ROM are a lid-driven cavity and the flow over a cylinder. Both cases use two-dimensional computational domains. The lid-driven cavity case consists of three design parameters controlling the geometry of the domain, while the cylinder case consists of two design parameters controlling geometric and physical properties of the simulation. Different CFD solvers are used for each case; this is done to illustrate the portability of the ROM and allow for the use of a structured grid for the cylinder case. Latin hypercube sampling (LHS),<sup>33</sup> a popular statistical method for generating samples from a multi-dimensional parameter space, is used to generate sets of design parameters that are split into training and test sets. LHS aims to maximize the distance and minimize the correlation among samples produced by using uniform sampling. The ROM is used to predict the vertical and horizontal components of the velocity. The training data for each velocity component are scaled to a range [0,1] using min-max scaling before input into the CAE. Similarly, the CAE latent variables are also scaled using min-max scaling before being used for the LSTM. Data normalization improves performance and allows the optimizer to learn optimal network parameters at a much faster rate.<sup>34</sup> Both the CAE and LSTM are trained using the mean squared error loss function. An NVIDIA DGX system, consisting of eight NVIDIA A100 GPUs, is used to train the CAE and LSTMs, run the cylinder simulations, and for ROM inference. The lid-driven cavity simulations are run on a local workstation using a single CPU core.

The performance of the ensemble ROM is tested against a CAE-LSTM ROM that uses a single LSTM model. Using five different random initializations of the weights and biases prescribed by different seeds, the time-series of latent variables are visually compared in addition to performance metrics averaged over the seeds. When training multiple LSTMs on bootstrapped datasets for a single seed  $s$ , the same initial weights and biases are used for each individual weak learner. Evaluating the ROM's performance using different seeds allows for testing its sensitivity to the initial weights and biases.  $N_s = 5$  different

seeds are used to initialize the LSTM model, with  $s = [1-5]$ . For a single test point  $\mu^*$  and seed  $s$ , an error metric of interest is the relative  $L^2$  error in a field component averaged over the prediction time horizon,

$$\epsilon_s = \frac{1}{T - T_i} \sum_{t=T_i+1}^T \frac{\|x^t - \tilde{x}^t\|^2}{\|x^t\|^2}. \quad (11)$$

The errors are then averaged over each seed to give a single error metric  $\bar{\epsilon}$ ,

$$\bar{\epsilon} = \frac{1}{N_s} \sum_{s=1}^{N_s} \epsilon_s. \quad (12)$$

We are also interested in how sensitive the error is to the initial weights and biases of the model. The sample standard deviation  $\sigma_s$  of  $\epsilon_s$  is used as a measure of this

$$\sigma_s = \text{std}(\epsilon_s). \quad (13)$$

A lower standard deviation in the error term indicates that the model performance over the prediction horizon does not vary much with the initial weights and biases.

#### A. Lid-driven cavity

The first test case is a geometrically parameterized two-dimensional lid-driven cavity flow, a popular benchmark problem for CFD solvers. Unsteady, incompressible, laminar flow is simulated by solving the Navier–Stokes equations, given as follows:

$$\int_S \vec{v} \cdot d\vec{n} \, dS = 0, \quad (14)$$

$$\int_{\Omega} \frac{\partial \vec{v}}{\partial t} \, d\Omega + \int_S \vec{v} \vec{v} \cdot d\vec{n} \, dS + \int_{\Omega} \nabla p \, d\Omega - \int_S \nu (\nabla \vec{v} + \nabla \vec{v}^T) \cdot d\vec{n} \, dS = 0, \quad (15)$$




FIG. 5. Schematics describing the lid-driven cavity problem.

where  $\vec{V} = [u, v]$  is the velocity vector,  $\Omega$  is the fluid domain,  $S$  is the face-area vector,  $\vec{n}$  is the outward-pointing normal,  $\nu$  is the kinematic viscosity, and  $p$  is the pressure. OpenFOAM,<sup>35</sup> an open-source toolbox for multiphysics simulation, is used. The boundary conditions and design parameters are shown in Fig. 5. On each edge  $\Gamma_i$ ,  $i \in [1, 2, 3, 4]$  of the computational domain  $\Omega$ , the pressure gradient  $\nabla p$  is set to 0. No slip conditions exist on all of the edges except the top, where  $u = 1, v = 0$ . The reference pressure is set to  $p = 0$  at the bottom left corner. The first design parameter  $\mu_1$  controls the horizontal length, the second  $\mu_2$  controls the slanting length, while the third  $\mu_3$  controls the slanting angle. The Reynolds number  $Re$  is set to 400 and is related to the kinematic viscosity as follows:

$$Re = 400 = \frac{\max(\mu_1, \mu_2)}{\nu(\boldsymbol{\mu})}. \quad (16)$$

The bounds for the design parameters are given as follows:

$$\begin{aligned} \mu_1 &\in [1.2, 1.8], \\ \mu_2 &\in [1.2, 1.8], \\ \mu_3 &\in \left[-\frac{\pi}{6}, \frac{\pi}{6}\right]. \end{aligned}$$

The computational mesh consists of  $128 \times 128$  cells distributed uniformly in the  $x$  and  $y$  directions. A single cell exists in the spanwise direction  $z$ , resulting in  $N = 16\,384$ . The initial condition is the solution to steady, incompressible, laminar flow at  $Re = 20$  which is computed using the standard OpenFOAM solver simpleFoam. Unsteady flow is simulated using the OpenFOAM solver icoFoam for  $T = 5$  s with data being saved every 0.025 s, leading to 200 time snapshots for a single simulation. Both solvers utilize the finite volume method (FVM). Figure 6 shows contours of  $u$  and  $v$  at  $T = 5$  at three different sets of design parameters. A vortex that varies in shape, size, and location with the design parameters is shown for both velocity components. The variation of the flow with  $\boldsymbol{\mu}$  is highly nonlinear, making this a difficult prediction problem for ROMs.

In total, 100 sets of design parameters are generated and randomly split into 90 training samples and 10 test points, which are given in Table I. The CAE-eLSTM ROM uses  $m = 64$  bagged LSTMs and a window size  $w = 20$ . These parameters were chosen through a trial-and-error process with the goal of maximizing predictive accuracy while trying to keep the computational cost of training the LSTM

ensemble relatively low. This was done by directly evaluating the performance on the test set, while in practice methods like cross-validation would be used. Initial guesses for the window size were based on similar problems found in the literature. An initial number of approximately 50 weak learners was used, which was increased in increments of 8 until the performance gains plateaued. The CAE latent dimension is set to  $k = 4$ ; below this value, the CAE reconstruction errors were found to be worse, and increasing  $k$  offered no improvement. The LSTM architecture consists of two hidden layers consisting of 50 neurons each and a dropout<sup>36</sup> rate of 0.1. The output layer contains a sigmoid activation function so the outputs are constrained to a range of  $[0, 1]$ . The CAE architecture is given in Appendix B. The Adam optimizer<sup>37</sup> is used to train both the CAE and LSTM, with an initial learning rate of  $\eta = 5 \times 10^{-4}$  and weight decay of  $\lambda = 1 \times 10^{-6}$ . Dropout and weight decay are used to prevent overfitting. The CAE is trained for a total of 200 epochs, while an individual LSTM is trained for 250 epochs. At the test points, the FOM is simulated for  $T_i = 0.75$  s (30 snapshots) or 15% of the total time horizon. The final 20 of these snapshots are used as the initial input sequence for the LSTM ensemble.

Figure 7 shows the trajectories of the latent variables computed using ensemble and single LSTMs at the test point  $\boldsymbol{\mu}^* = [1.299, 1.689, -0.0367]$  using five different seeds which control the initial weights and biases of the LSTM. It is shown that for all of the latent variables, using LSTM ensembles leads to higher prediction accuracy and significantly lower variance. When using a single LSTM, the different predictions quickly diverge from each other and the ground truth. Due to their high capacity to learn, neural networks often exhibit high variance. Since neural networks contain a large number of parameters, the effects of error propagation between different trained models are heightened due to relatively large differences in autoregressive predictions being fed back into the model. In contrast, the ensemble model is much less sensitive to the initial seed, and the predicted trajectories do not differ much. The second latent variable exhibits diverging trajectories when using the ensemble model, but the effect is much less pronounced than when using a single LSTM. Figure 8 shows contours of  $u$  and  $v$  at the test point as well as absolute ROM errors averaged over the last 10% of the simulation ( $t \in [4.5, 5]$  s) for a single seed. The errors are considerably larger when using a single LSTM throughout the computational domain for both  $u$  and  $v$ . Table II lists the seed-averaged relative errors  $\bar{\epsilon}$  in  $u$  and  $v$  for the test points. The ensemble




FIG. 6. Snapshots at  $T=5$  of  $u$  (top) and  $v$  (bottom) for the lid-driven cavity problem at three different sets of design parameters.

method offers better performance in predicting both fields at all of the test points, usually by wide margins. Table III lists the standard deviation of these errors; at all of the test points, the standard deviation is also smaller. As there is a lower spread amongst the different seeds in the time-averaged error metrics, this shows that using the ensemble method leads to much greater stability and reliability in predictions. Table IV lists the computational costs associated with neural network training, CFD simulation, and ROM inference. The given CFD wall time is the portion of the simulation over the prediction time horizon; using the ROM for inference over this period of time offers a speed-up of approximately  $9.4\times$  over CFD. The ROM inference costs include

TABLE I. Test case design parameters for the lid-driven cavity problem.

Test case index	$\mu_1$	$\mu_2$	$\mu_3$
1	1.473	1.701	0.288
2	1.659	1.653	-0.372
3	1.593	1.611	0.455
4	1.209	1.443	0.371
5	1.299	1.689	-0.0367
6	1.371	1.311	-0.351
7	1.443	1.617	-0.487
8	1.232	1.335	-0.330
9	1.719	1.785	-0.455
10	1.281	1.449	-0.340

autoregressive prediction using the ensemble LSTM and obtaining full-order approximations using the decoder. The CAE is much more expensive to train than a single LSTM as it has a larger number of trainable parameters. While a multi-GPU architecture was used to train the LSTMs in parallel, the cost to train a single LSTM is given.

### B. 2D cylinder

The next test case involves two-dimensional incompressible, unsteady, laminar flow over a cylinder between two solid walls. Eventually, the lid-driven cavity flow from the previous test case reaches steady-state and does not exhibit long-term transient behavior that is commonly found in fluid dynamics problems. Laminar flow over a cylinder is a well-studied problem in fluid dynamics, with both experimental and computational results present in the literature.<sup>38,39</sup> Unsteady cylinder flow is characterized by the presence of vortices that separate from the surface and form in the wake. This distinctive pattern is known as the Von Kármán vortex street, where alternating vortices of a regular pattern are shed downstream of the cylinder. The unsteady Navier–Stokes equations found in Eqs. (14) and (15) are solved using XLB,<sup>40</sup> a lattice Boltzmann method<sup>41</sup> (LBM) library utilizing the JAX framework<sup>42</sup> available for Python, which allows for effective scaling across multiple CPUs, GPUs, and distributed multi-GPU systems. The cited works can be referred to for an overview of the Lattice Boltzmann method and its implementation in XLB.

No-slip boundary conditions are applied to the cylinder’s surface and top and bottom walls. A Poiseuille flow profile is used for the inlet velocity. Extrapolation outflow boundary conditions are used for the outlet to allow the fluid flow to exit the domain freely. The

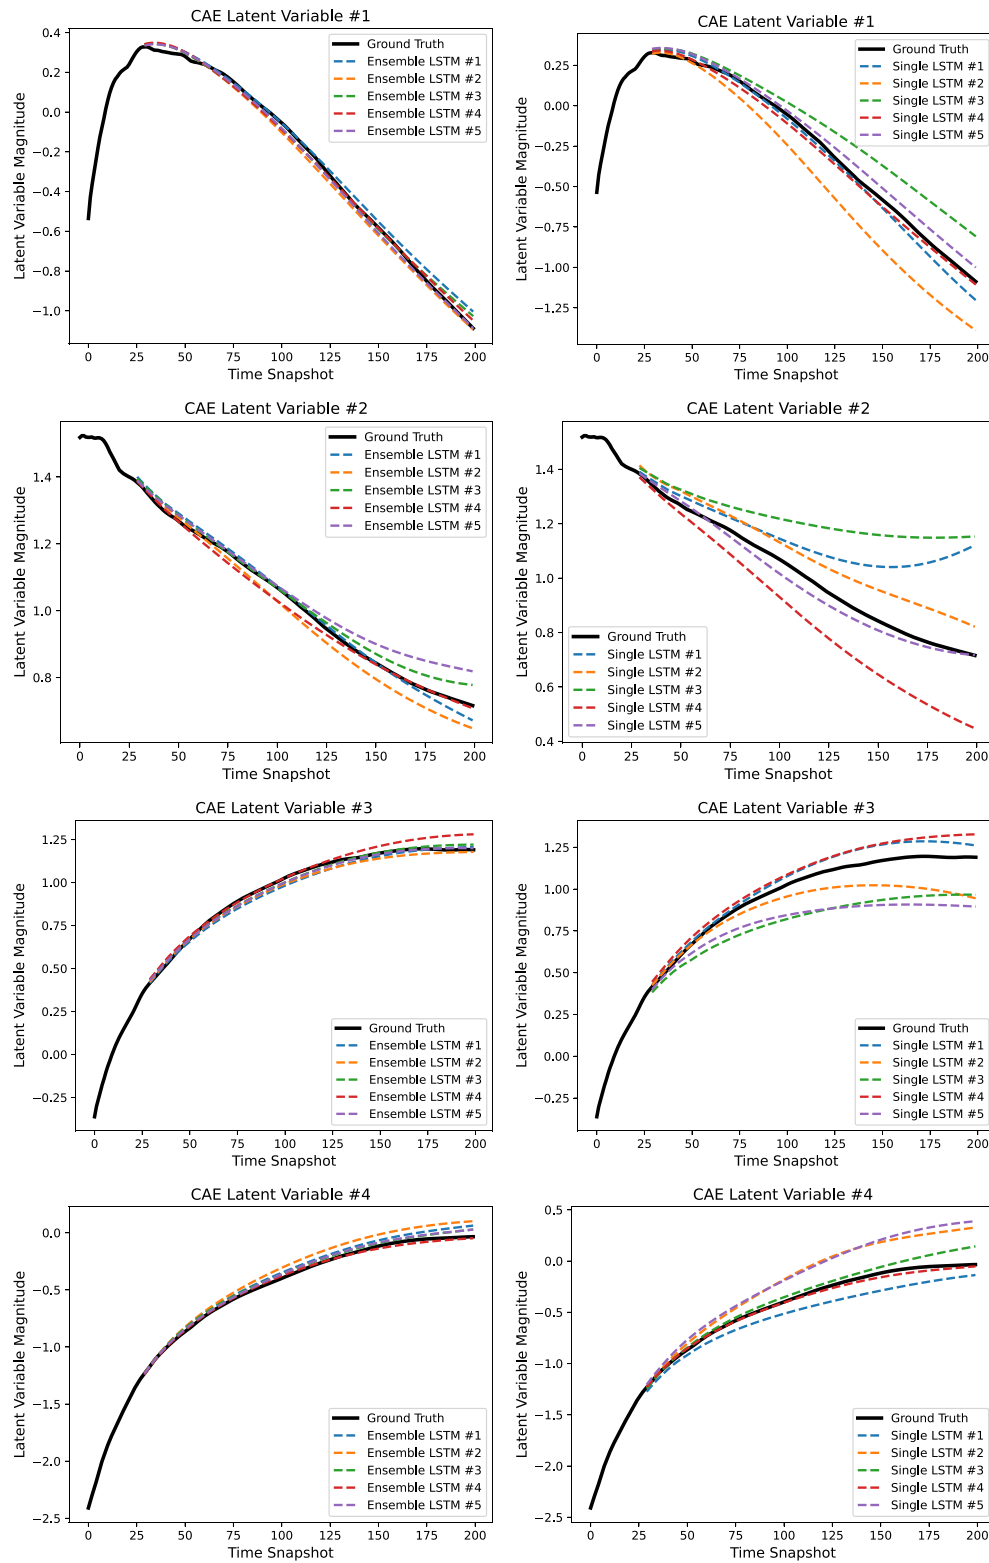


FIG. 7. Comparison of latent variable trajectories for the lid-driven cavity case at  $\mu^* = [1.299, 1.689, -0.0367]$ .

21 July 2024 19:03:28




FIG. 8. Time snapshot of  $u$  (top) and  $v$  (bottom) at  $T=5$  and errors averaged over the last 10% of the simulation at  $\mu^* = [1.299, 1.689, -0.0367]$ .

computational domain measures  $1536 \times 512$  voxels that are uniformly spaced. The cylinder is centered at  $x_c, y_c = [160, 256]$  (zero-based indexing is used). Simulation results are down-sampled onto a grid that measures  $384 \times 128$  before being used for the ROM, resulting in  $N=49152$ , as the original domain's large size leads to a very large training cost as well as memory usage. Two design parameters are

used, the diameter  $d$  of the cylinder and the Reynolds number  $Re$ . The bounds of the design parameters are given as

$$\begin{aligned} \mu_1 &= d \in [48, 68], \\ \mu_2 &= Re \in [120, 240]. \end{aligned}$$

The diameter  $d$  is set to an integer quantity by rounding to the nearest whole number. The wake structure behind the cylinder undergoes instabilities<sup>43,44</sup> at a critical Reynolds number of approximately  $Re_c=180$ , where the vortices transition to becoming turbulent. As a result, the prescribed range given for  $Re$  in the parameter space includes a variety of physical regimes. Additionally, both  $Re$  and  $d$  control the size and periodicity of the shed vortices, making this a difficult prediction problem for ROMs. The freestream velocity in the  $x$ -direction is set to  $u_\infty = 0.001$  in non-dimensional units. The simulation is run for  $T=750000$  timesteps with data being saved every 1000 steps, leading to 750 snapshots for a single simulation. In LBM simulations, time and space are often non-dimensionalized by  $\Delta t$  and  $\Delta x$  to arrive at dimensionless values of  $\Delta x^* = \Delta t^* = 1$ . However, one needs to ensure that  $\Delta t/\Delta x \ll 1$  to avoid compressibility errors. In these runs, we picked  $\Delta t/\Delta x = 0.001$ . The initial condition for a simulation of a given diameter  $d$  is the solution to steady flow at  $Re = 20$ .

In total, 50 sets of design parameters are generated and split into 45 training samples and five test points, shown in Fig. 9. Table V also lists the test point design parameters. The ROM uses  $m=96$  bagged LSTMs and a window size of  $w=30$ , which were again chosen

TABLE II. Average relative errors in  $u$  and  $v$  for the lid-driven cavity case. Boldface denotes the lowest values.

Test case index	$\bar{\epsilon}, u$ (ensemble)	$\bar{\epsilon}, u$ (single)	$\bar{\epsilon}, v$ (ensemble)	$\bar{\epsilon}, v$ (single)
1	<b>0.024 21</b>	0.071 19	<b>0.025 14</b>	0.072 04
2	<b>0.019 58</b>	0.041 99	<b>0.021 74</b>	0.038 70
3	<b>0.026 22</b>	0.052 82	<b>0.028 68</b>	0.062 96
4	<b>0.036 40</b>	0.065 75	<b>0.039 24</b>	0.067 42
5	<b>0.024 42</b>	0.065 89	<b>0.019 96</b>	0.051 31
6	<b>0.041 77</b>	0.049 39	<b>0.039 64</b>	0.047 33
7	<b>0.024 45</b>	0.037 71	<b>0.027 89</b>	0.039 15
8	<b>0.045 59</b>	0.053 04	<b>0.042 51</b>	0.051 94
9	<b>0.033 13</b>	0.046 78	<b>0.038 91</b>	0.049 00
10	<b>0.028 77</b>	0.042 91	<b>0.027 17</b>	0.040 79

**TABLE III.** Standard deviation in relative errors of  $u$  and  $v$  for the lid-driven cavity case. Boldface denotes the lowest values.

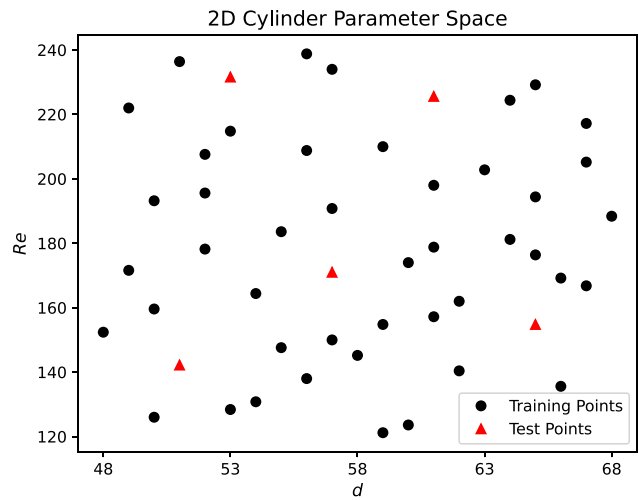
Test case index	$\sigma_s, u$ (ensemble)	$\sigma_s, u$ (single)	$\sigma_s, v$ (ensemble)	$\sigma_s, v$ (single)
1	<b>0.005 225</b>	0.028 19	<b>0.004 470</b>	0.030 51
2	<b>0.001 637</b>	0.009 202	<b>0.001 766</b>	0.012 67
3	<b>0.004 603</b>	0.018 85	<b>0.004 040</b>	0.022 63
4	<b>0.008 210</b>	0.020 07	<b>0.006 915</b>	0.011 72
5	<b>0.003 980</b>	0.011 21	<b>0.004 222</b>	0.007 739
6	<b>0.003 244</b>	0.015 55	<b>0.002 547</b>	0.017 26
7	<b>0.002 077</b>	0.007 614	<b>0.002 466</b>	0.006 886
8	<b>0.002 393</b>	0.020 31	<b>0.002 030</b>	0.020 41
9	<b>0.002 515</b>	0.007 371	<b>0.002 550</b>	0.013 24
10	<b>0.002 467</b>	0.016 27	<b>0.001 759</b>	0.014 93

through a trial-and-error process to maximize accuracy while keeping the number of weak learners to a minimum. The CAE latent dimension is set to  $k = 10$ . Similar to the lid-driven cavity case, values below  $k = 10$  resulted in higher reconstruction errors and increasing the latent dimension further offers no improvement. A bidirectional LSTM architecture<sup>45</sup> is used. Bidirectional recurrent neural networks process sequential data in both forward and backward directions, allowing the model to learn both past and future context. For problems that are characterized by regularly repeating patterns such as a vortex street, bidirectional LSTMs can better learn the cyclic nature of the pattern by leveraging context in both directions. The network consists of three hidden layers with 36 neurons each and a dropout rate of 0.15. The output layer again contains a sigmoid activation function, and the Adam optimizer with the same initial learning rate  $\eta = 5 \times 10^{-4}$  and weight decay of  $\lambda = 1 \times 10^{-6}$  is used for both the LSTM and CAE. Again, the CAE is trained for 200 epochs while an individual LSTM is trained for 250 epochs. The CAE architecture is given in Appendix B. At the test points, the FOM is run for  $T_i = 300\,000$  s (300 snapshots) or 40% of the total simulation time. This value is required to be high as the flow exhibits highly oscillatory behavior initially, leading to very noisy latent variables that cannot be used for model training. As a result, latent variable sequences are generated starting at the 200th snapshot, and  $\mathbf{A}_{\text{train}}$  does not contain time-series data of latent variables before this point.

Figure 10 shows the latent variable trajectories for the first four latent variables at the test point  $\mu^* = [51, 142.2]$ . For each latent variable, the predictions given by single LSTMs are initially similar, but eventually diverge in terms of both the amplitude and frequency of the latent variables, leading to large inaccuracies. Using LSTM

**TABLE IV.** Computational costs associated with the lid-driven cavity problem.

Task	Wall time (s)
CAE	749
Single LSTM	81
CFD (1 CPU)	136
ROM inference	14.5



**FIG. 9.** Training and test design parameters for the 2D cylinder case.

ensembles greatly reduces this effect, and the resultant latent variable trajectories follow the ground truth closely and do not differ greatly in amplitude nor frequency. The regular pattern of the Kármán vortex street is well-predicted given a small amount of initial latent variable history.

Figure 11 shows snapshots of  $u$  and  $v$  at  $T = 750\,000$  as well as ROM errors averaged over the last 10% of the simulation for a single seed. The errors are significantly lower throughout the computational domain when using LSTM ensembles, which are greatest in the wake of the cylinder and immediately downstream of it. Table VI lists the seed-averaged relative errors in  $u$  and  $v$  at the test points. The ensemble ROM again offers better performance in predicting both fields at all of the test points. Table VII lists the standard deviation of these errors, which are again lower at all of the test points. At the fourth test case index ( $\mu^* = [57, 171.0]$ ), the errors and standard deviations are similar. This test point lies in the middle of the parameter space and is well-surrounded by training samples; as a result, the advantage gained using an ensemble method may be marginal. Table VIII lists the computational costs associated with the problem. Using the ROM for inference offers a speed-up of approximately  $27\times$  over CFD (again, the given CFD and ROM inference costs are for the portion of the simulation over the prediction time horizon). Due to the increased domain size and amount of training data, the CAE and LSTM are more expensive to train when compared to the lid-driven cavity case.

**TABLE V.** Test case design parameters for the 2D cylinder case.

Test case index	$d$	$Re$
1	53	231.6
2	51	142.2
3	65	154.8
4	57	171.0
5	61	225.6

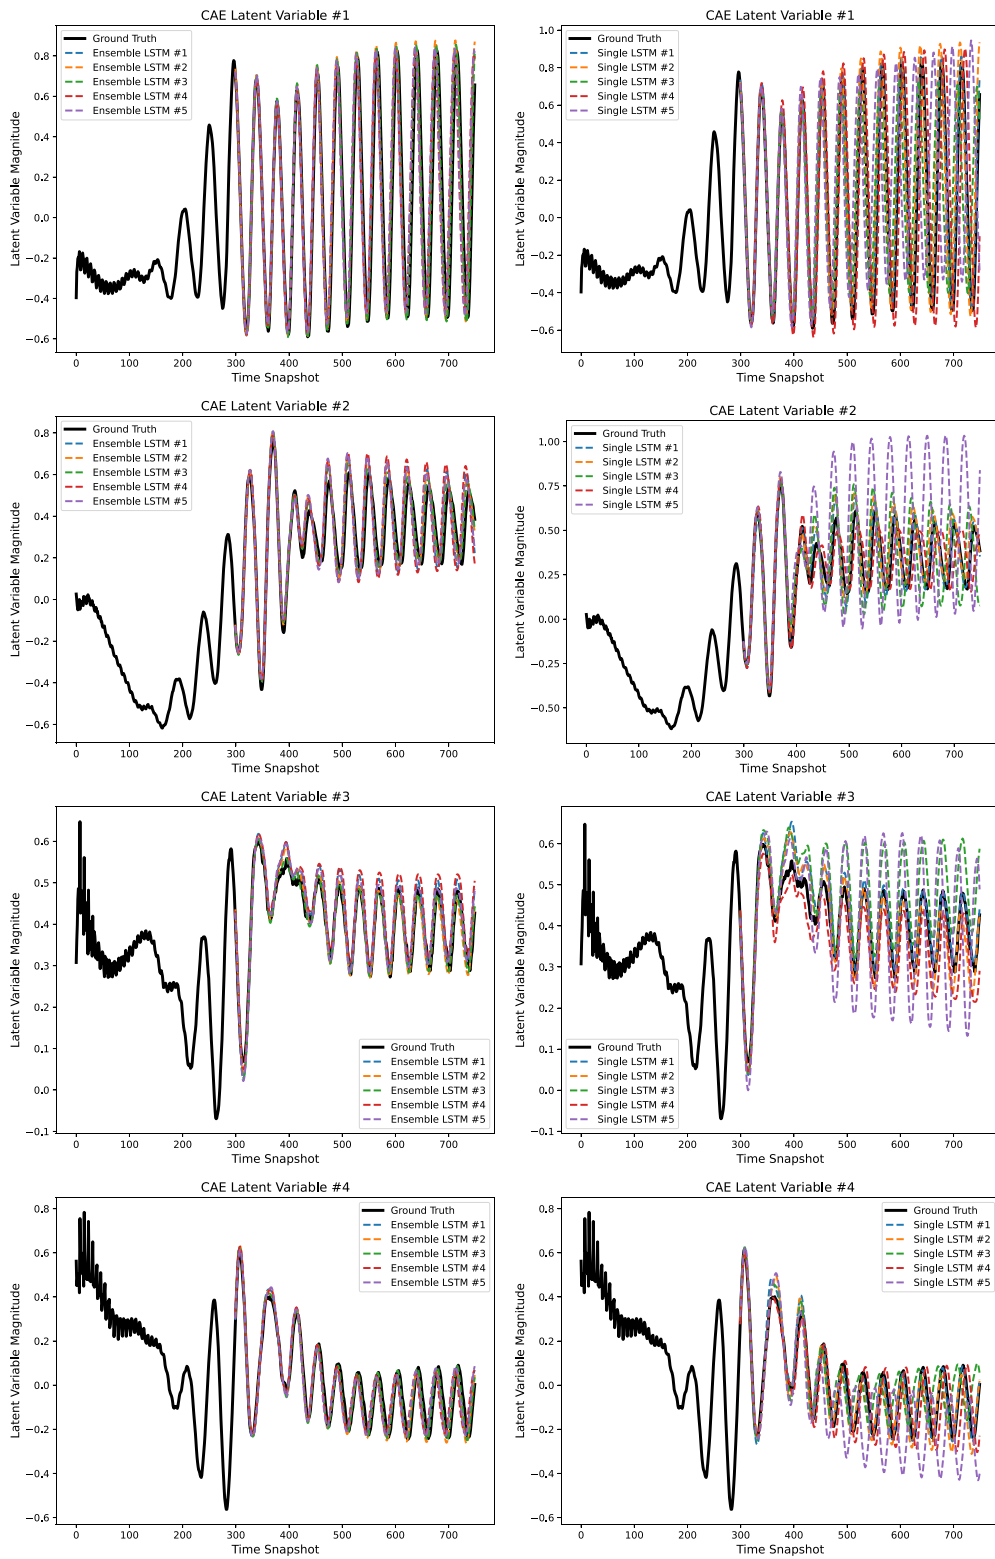


FIG. 10. Comparison of latent variable trajectories for the 2D cylinder case at  $\mu^* = [51, 142.2]$ .

21 July 2024 19:03:28

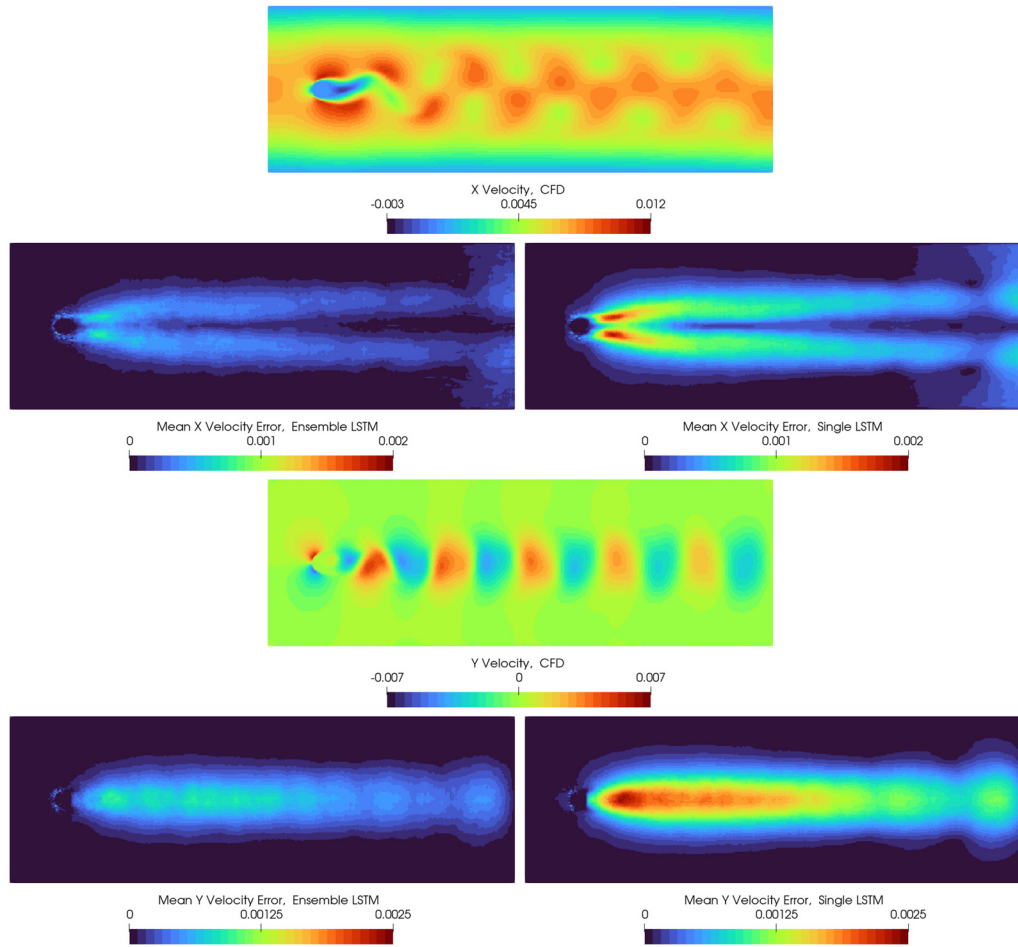


FIG. 11. Time snapshots of  $u$  and  $v$  at  $T = 750\,000$  and errors averaged over the last 10% of the simulation at  $\mu^* = [51, 142.2]$ .

TABLE VI. Average relative errors in  $u$  and  $v$  for the 2D cylinder case. Boldface denotes the lowest values.

Test case index	$\bar{\epsilon}, u$ (ensemble)	$\bar{\epsilon}, u$ (single)	$\bar{\epsilon}, v$ (ensemble)	$\bar{\epsilon}, v$ (single)
1	<b>0.043 24</b>	0.079 75	<b>0.3379</b>	0.6162
2	<b>0.034 63</b>	0.058 38	<b>0.2971</b>	0.4871
3	<b>0.023 80</b>	0.051 95	<b>0.1399</b>	0.3315
4	<b>0.054 46</b>	0.066 36	<b>0.3995</b>	0.4900
5	<b>0.036 55</b>	0.060 29	<b>0.2269</b>	0.4059

TABLE VIII. Computational costs associated with the 2D cylinder case.

Task	Wall time (s)
CAE	2933
Single LSTM	140
CFD (2 GPUs)	1305
ROM inference	48.4

TABLE VII. Standard deviation in relative errors of  $u$  and  $v$  for the 2D cylinder case. Boldface denotes the lowest values.

Test case index	$\sigma_s, u$ (ensemble)	$\sigma_s, u$ (single)	$\sigma_s, v$ (ensemble)	$\sigma_s, v$ (single)
1	<b>0.013 21</b>	0.041 84	<b>0.1138</b>	0.3372
2	<b>0.007 607</b>	0.033 04	<b>0.070 05</b>	0.2720
3	<b>0.004 610</b>	0.015 87	<b>0.034 93</b>	0.1136
4	<b>0.015 21</b>	0.020 26	<b>0.1186</b>	0.1536
5	<b>0.007 528</b>	0.033 98	<b>0.059 04</b>	0.2684

#### IV. CONCLUSION

This work presents a fully data-driven ROM framework for time-dependent physics simulations using convolutional autoencoders and LSTM ensembles. The combined framework, referred to as the CAE-eLSTM ROM, is used to mitigate the issue of error propagation when performing autoregressive time-series prediction at unseen data sets, a common problem in data-driven CFD applications. The ROM obtains




FIG. 12. Reconstruction errors at  $T = 5$  in  $u$  and  $v$  between POD and CAE ( $k = 4$ ) at  $\mu^* = [1.299, 1.689, -0.0367]$ .

21 July 2024 19:03:28

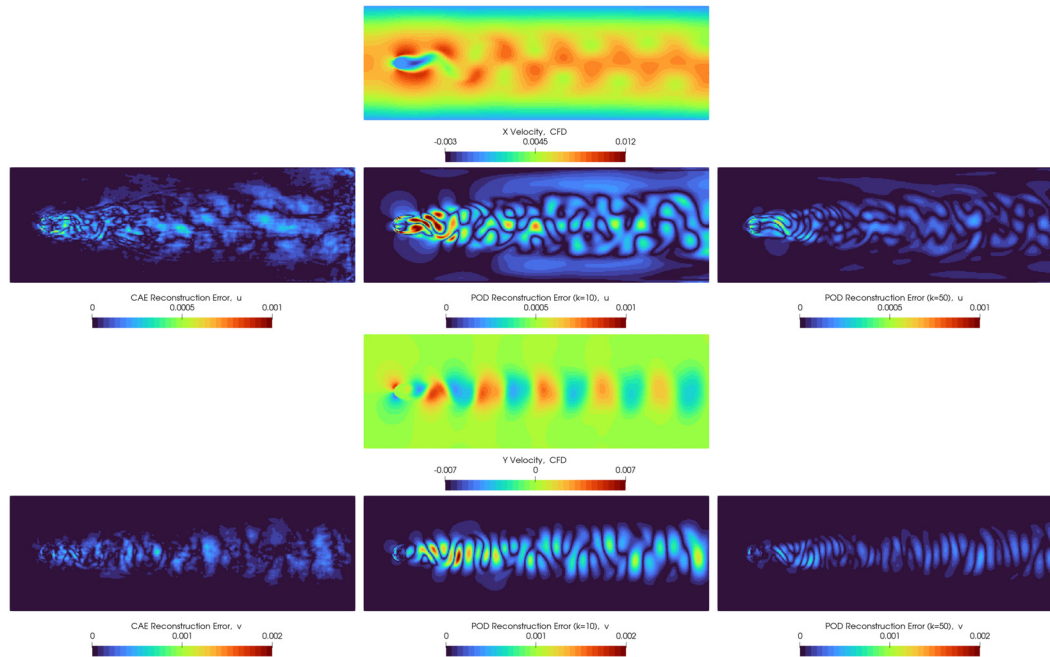


FIG. 13. Reconstruction errors at  $T=5$  in  $u$  and  $v$  between POD and CAE ( $k=4$ ) at  $\mu^* = [51, 142.2]$ .

a low-dimensional solution manifold using convolutional autoencoders, a type of artificial neural network useful for reconstructing spatially distributed data. Using the encoder section of the autoencoder, the low-dimensional latent variables of the training data are used to generate multivariate time-series sequences that are used to train LSTMs, a type of recurrent neural network useful for modeling sequential data. Bagging, a popular ensemble learning technique, is used to train multiple base LSTMs on subsets of the training sequences sampled randomly with replacement. Ensemble learning is chosen as it is a useful tool for improving the stability and accuracy of machine learning methods. At unseen test points, the full-order model is run to obtain an initial sequence of latent variables through the encoder and autoregressive time-series prediction is used to predict the latent variables over a prescribed time horizon. The predicted latent variables are then passed through the decoder to obtain spatial reconstructions of the FOM at different points in time.

When applied to two incompressible, unsteady, laminar fluid dynamics problems, the proposed ROM exhibits high levels of accuracy when compared to a ROM using a single LSTM model. The use of LSTM ensembles significantly diminishes the error propagation issue, with latent variable trajectories from different weight initializations showing low levels of variance. Error metrics show that using the ensemble ROM leads to better predictive performance at all test points. Furthermore, the fully data-driven nature of the ROM allows it to be applied to two different CFD solvers. Although the cost associated with training ensemble neural networks is high, bagging allows for models to be trained in parallel, and multi-GPU architectures can be used to significantly reduce the cost. The addition of an ensemble model for time-series predictions, while simple, greatly improves the performance of the ROM and does not require additional training data. A limitation of the presented framework is its significantly

increased offline training cost. While this can be alleviated to a great degree by training weak learners in parallel, multi-GPU platforms are expensive and may not be readily available for practitioners. Additionally, the use of convolutional autoencoders limits the application of the ROM to structured grids unless a mesh interpolation method is used. Future work will focus on further developing the method to apply it to three-dimensional geometries, problems exhibiting turbulent flow, and more complex geometries involving unstructured meshes. Problems consisting of time-series datasets that exhibit non-smooth variations over long periods of time will be of particular interest, as they are more representative of real-world applications and can reap greater benefits from ensemble learning. Computationally efficient optimal hyperparameter selection methods will also be explored to lower the offline costs of training the ROM.

## AUTHOR DECLARATIONS

### Conflict of Interest

The authors have no conflicts to disclose.

### Author Contributions

**Rakesh Halder:** Conceptualization (equal); Formal analysis (lead); Methodology (lead); Software (equal); Writing – original draft (lead). **Mohammadmehdi Ataei:** Conceptualization (equal); Software (equal); Supervision (equal); Writing – review & editing (equal). **Hesam Salehipour:** Conceptualization (equal); Software (equal); Supervision (equal); Writing – review & editing (equal). **Krzysztof Fidkowski:** Writing – review & editing (equal). **Kevin Maki:** Writing – review & editing (equal).

**DATA AVAILABILITY**

The data that support the findings of this study are available from the corresponding author upon reasonable request.

**APPENDIX A: RECONSTRUCTION ERRORS**

Reconstruction error comparisons between CAE and POD in  $u$  and  $v$  are given for both test cases. The same test points from Sec. III are used. Figure 12 shows the reconstruction errors between CAE ( $k=4$ ) and POD at  $k=4$  and  $k=25$  for the lid-driven cavity case at  $\mu^* = [1.299, 1.689, -0.0367]$ . For the same latent dimension, the CAE reconstruction errors are significantly smaller than those from POD throughout the domain in both  $u$  and  $v$ . For  $k=25$ , reconstruction errors from POD are similar to that of CAE. Figure 13 shows reconstruction errors for the cylinder case at  $\mu^* = [51, 142.2]$ . Again, it is shown that for the same latent dimension, the reconstruction error from a CAE is much smaller than

from POD. At  $k=50$ , the error contours are comparable. Using POD would require the LSTM model to handle a very large multivariate time-series problem, which would degrade the accuracy of predictions.

**APPENDIX B: CONVOLUTIONAL AUTOENCODER ARCHITECTURES**

The convolutional autoencoder architectures for the lid-driven cavity and 2D cylinder cases are given in Tables IX and X, respectively. Both architectures are similar and use convolutional layers in the encoder section to progressively reduce the number of pixels or grow the number of filters in subsequent layers. Convolutional layers are usually followed max-pooling layers that have batch normalization<sup>46</sup> applied to them (referred to as pool-norm layers). Batch normalization normalizes the input of each layer over a mini-batch, reducing internal covariate shift, leading to more efficient

**TABLE IX.** Convolutional autoencoder architecture for the lid-driven cavity case.

Layer	Number of filters	Kernel size	Activation function	Size of output
Input				$128 \times 128 \times 2$
Convolutional	8	$3 \times 3$	Leaky ReLU	$128 \times 128 \times 8$
Pool-norm		$2 \times 2$		$64 \times 64 \times 8$
Convolutional	16	$3 \times 3$	Leaky ReLU	$64 \times 64 \times 16$
Pool-norm		$2 \times 2$		$32 \times 32 \times 16$
Convolutional	32	$3 \times 3$	Leaky ReLU	$32 \times 32 \times 32$
Pool-norm		$2 \times 2$		$16 \times 16 \times 32$
Convolutional	64	$3 \times 3$	Leaky ReLU	$16 \times 16 \times 64$
Pool-norm		$2 \times 2$		$8 \times 8 \times 64$
Reshape				4096
Fully connected			Leaky ReLU	128
Batch norm				128
Fully connected (latent space)				4
Fully connected			Leaky ReLU	128
Batch norm				128
Fully connected			Leaky ReLU	4096
Batch norm				4096
Reshape				$8 \times 8 \times 64$
Convolutional transpose	32	$3 \times 3$	Leaky ReLU	$16 \times 16 \times 32$
Batch norm				$16 \times 16 \times 32$
Convolutional transpose	16	$3 \times 3$	Leaky ReLU	$32 \times 32 \times 16$
Batch norm				$32 \times 32 \times 16$
Convolutional transpose	8	$3 \times 3$	Leaky ReLU	$64 \times 64 \times 8$
Batch norm				$64 \times 64 \times 8$
Convolutional transpose	2	$3 \times 3$	Sigmoid	$128 \times 128 \times 2$

**TABLE X.** Convolutional autoencoder architecture for the 2D cylinder case.

Layer	Number of filters	Kernel size	Activation function	Size of output
Input				$384 \times 128 \times 2$
Convolutional	8	$5 \times 5$	Leaky ReLU	$192 \times 64 \times 8$
Batch norm				$192 \times 64 \times 8$
Convolutional	16	$5 \times 5$	Leaky ReLU	$96 \times 32 \times 16$
Batch norm				$96 \times 32 \times 16$
Convolutional	32	$3 \times 3$	Leaky ReLU	$96 \times 32 \times 32$
Pool-norm		$2 \times 2$		$48 \times 16 \times 32$
Convolutional	64	$3 \times 3$	Leaky ReLU	$48 \times 16 \times 64$
Pool-norm		$2 \times 2$		$24 \times 8 \times 64$
Reshape				12 288
Fully connected			Leaky ReLU	128
Batch norm				128
Fully connected (latent space)				10
Fully connected			Leaky ReLU	128
Batch norm				128
Fully connected			Leaky ReLU	12 288
Batch norm				12 288
Reshape				$24 \times 8 \times 64$
Convolutional transpose	32	$3 \times 3$	Leaky ReLU	$48 \times 16 \times 32$
Batch norm				$48 \times 16 \times 32$
Convolutional transpose	16	$3 \times 3$	Leaky ReLU	$96 \times 32 \times 16$
Batch norm				$96 \times 32 \times 16$
Convolutional transpose	8	$5 \times 5$	Leaky ReLU	$192 \times 64 \times 8$
Batch norm				$192 \times 64 \times 8$
Convolutional transpose	2	$5 \times 5$	Sigmoid	$384 \times 128 \times 2$

21 July 2024, 19:03:28

gradient flow during backpropagation. Fully connected layers are present in both the encoder and decoder. Although their inclusion leads to a larger number of network parameters, reconstruction accuracies are significantly improved. The decoder consists of convolutional transpose layers that are used to progressively increase the number of pixels. The output layer uses a sigmoid activation function to constrain the range of outputs to  $[0,1]$ . The leaky ReLU activation function is used for convolutional and fully connected layers and is given as follows:

$$\phi(x) = \begin{cases} x, & \text{if } x \geq 0, \\ \alpha x, & \text{if } x < 0. \end{cases} \quad (\text{B1})$$

A value of  $\alpha = 0.25$  is used for all leaky ReLU activation functions. The activation function is often used over the standard ReLU activation function ( $\alpha = 0$ ). For inputs less than zero, the ReLU activation function returns a gradient of zero, effectively rendering certain neurons inactive and inhibiting gradient flow, referred to as the dying ReLU problem.

## REFERENCES

- <sup>1</sup>D. J. Lucia, P. S. Beran, and W. A. Silva, "Reduced-order modeling: New approaches for computational physics," *Prog. Aerosp. Sci.* **40**(1–2), 51–117 (2004).
- <sup>2</sup>G. Berkooz, P. J. Holmes, and J. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annu. Rev. Fluid Mech.* **25**, 539–575 (2003).
- <sup>3</sup>J. S. Hesthaven and S. Ubbiali, "Non-intrusive reduced order modeling of non-linear problems using neural networks," *J. Comput. Phys.* **363**, 55–78 (2018).
- <sup>4</sup>Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature* **521**(7553), 436–444 (2015).
- <sup>5</sup>K. Lee and K. Carlberg, "Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders," *J. Comput. Phys.* **404**, 108973 (2020).
- <sup>6</sup>R. Halder, K. J. Fidkowski, and K. J. Maki, "Non-intrusive reduced-order modeling using convolutional autoencoders," *Int. J. Numer. Methods Eng.* **123**(21), 5369–5390 (2022).
- <sup>7</sup>H. Eivazi, S. Le Clainche, S. Hoyas, and R. Vinuesa, "Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows," *Expert Syst. Appl.* **202**, 117038 (2022).
- <sup>8</sup>A. Solera-Rico, C. S. Vila, M. A. Gómez, Y. Wang, A. Almashjary, S. T. M. Dawson, and R. Vinuesa, " $\beta$ -variational autoencoders and transformers for reduced-order modelling of fluid flows," *Nat. Commun.* **15**, 1361 (2023).
- <sup>9</sup>S. Wang, H. Wang, and P. Perdikaris, "Learning the solution operator of parametric partial differential equations with physics-informed DeepONets," *Sci. Adv.* **7**(40), eabi8605 (2021).
- <sup>10</sup>Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Aizzadenesheli, and A. Anandkumar, "Physics-informed neural operator for learning partial differential equations," *ACM/JMS J. Data Sci.* **1**, 9 (2021).
- <sup>11</sup>S. Grimberg, C. Farhat, and N. Youkilis, "On the stability of projection-based model order reduction for convection-dominated laminar and turbulent flows," *J. Comput. Phys.* **419**, 109681 (2020).
- <sup>12</sup>S. M. Rahman, S. Pawar, O. San, A. Rasheed, and T. Iliescu, "Nonintrusive reduced order modeling framework for quasigeostrophic turbulence," *Phys. Rev. E* **100**(5), 053306 (2019).
- <sup>13</sup>W. Xin Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong *et al.*, "A survey of large language models," [arXiv:2303.18223](https://arxiv.org/abs/2303.18223) (2023).
- <sup>14</sup>F. A. Gers, J. Schmidhuber, and F. Cummins, "Learning to forget: Continual prediction with LSTM," *Neural Comput.* **12**(10), 2451–2471 (2000).
- <sup>15</sup>T. Hill, M. O'Connor, and W. Remus, "Neural network models for time series forecasts," *Manage. Sci.* **42**(7), 1082–1092 (1996).
- <sup>16</sup>S. Lee and D. You, "Data-driven prediction of unsteady flow over a circular cylinder using deep learning," *J. Fluid Mech.* **879**, 217–254 (2019).
- <sup>17</sup>R. Vinuesa and S. L. Brunton, "Enhancing computational fluid dynamics with machine learning," *Nat. Comput. Sci.* **2**(6), 358–366 (2022).
- <sup>18</sup>P. Wu, F. Qiu, W. Feng, F. Fang, and C. Pain, "A non-intrusive reduced order model with transformer neural network and its application," *Phys. Fluids* **34**(11), 115130 (2022).
- <sup>19</sup>R. Maulik, B. Lusch, and P. Balaprakash, "Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders," *Phys. Fluids* **33**(3), 037106 (2021).
- <sup>20</sup>K. Hasegawa, K. Fukami, T. Murata, and K. Fukagata, "Machine-learning-based reduced-order modeling for unsteady flows around bluff bodies of various shapes," *Theor. Comput. Fluid Dyn.* **34**, 367–383 (2020).
- <sup>21</sup>J. Jeon, J. Lee, R. Vinuesa, and S. J. Kim, "Residual-based physics-informed transfer learning: A hybrid method for accelerating long-term CFD simulations via deep learning," *Int. J. Heat Mass Transfer* **220**, 124900 (2024).
- <sup>22</sup>O. Sagi and L. Rokach, "Ensemble learning: A survey," *Wiley Interdiscip. Rev.* **8**(4), e1249 (2018).
- <sup>23</sup>M. Mrosek, C. Othmer, and R. Radespiel, "Reduced-order modeling of vehicle aerodynamics via proper orthogonal decomposition," *SAE Int. J. Passeng. Cars—Mech. Syst.* **12**(3), 225–236 (2019).
- <sup>24</sup>Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel, "Handwritten digit recognition with a back-propagation network," in *Advances in Neural Information Processing Systems* (Morgan-Kaufmann, 1989), Vol. 2.
- <sup>25</sup>O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *Int. J. Comput. Vis.* **115**, 211–252 (2015).
- <sup>26</sup>J. Xu and K. Duraisamy, "Multi-level convolutional autoencoder networks for parametric prediction of spatio-temporal dynamics," *Comput. Methods Appl. Mech. Eng.* **372**, 113379 (2020).
- <sup>27</sup>D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature* **323**(6088), 533–536 (1986).
- <sup>28</sup>S. Hochreiter, "The vanishing gradient problem during learning recurrent neural nets and problem solutions," *Int. J. Uncertainty Fuzziness Knowl. Based Syst.* **06**(02), 107–116 (1998).
- <sup>29</sup>S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.* **9**(8), 1735–1780 (1997).
- <sup>30</sup>Y. Freund, R. Schapire, and N. Abe, "A short introduction to boosting," *J. Japan. Soc. Artif. Intell.* **14**(771–780), 1612 (1999).
- <sup>31</sup>L. Breiman, "Bagging predictors," *Mach. Learn.* **24**, 123–140 (1996).
- <sup>32</sup>A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2019), Vol. 32.
- <sup>33</sup>R. Jin, W. Chen, and A. Sudjianto, "An efficient algorithm for constructing optimal design of computer experiments," *J. Stat. Plann. Inference* **134**(1), 268–287 (2005).
- <sup>34</sup>J. Sola and J. Sevilla, "Importance of input data normalization for the application of neural networks to complex industrial problems," *IEEE Trans. Nucl. Sci.* **44**, 1464–1468 (1997).
- <sup>35</sup>H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, "A tensorial approach to computational continuum mechanics using object-oriented techniques," *Comput. Phys.* **12**(6), 620–631 (1998).
- <sup>36</sup>N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014).
- <sup>37</sup>D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).
- <sup>38</sup>D. J. Tritton, "Experiments on the flow past a circular cylinder at low Reynolds numbers," *J. Fluid Mech.* **6**(4), 547–567 (1959).
- <sup>39</sup>E. M. Saiki and S. Biringen, "Numerical simulation of a cylinder in uniform flow: Application of a virtual boundary method," *J. Comput. Phys.* **123**(2), 450–465 (1996).
- <sup>40</sup>M. Ataei and H. Salehipour, "XLB: A differentiable massively parallel lattice Boltzmann library in Python," *Comput. Phys. Commun.* **300**, 109187 (2024).

- <sup>41</sup>S. Chen and G. D. Doolen, "Lattice Boltzmann method for fluid flows," *Annu. Rev. Fluid Mech.* **30**(1), 329–364 (1998).
- <sup>42</sup>J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python+NumPy programs, 2018; see <https://github.com/google/jax>.
- <sup>43</sup>J. H. Gerrard, "Flow around circular cylinders; volume 1. fundamentals. by M. M. Zdravkovich. oxford science publications, 1997," *J. Fluid Mech.* **350**, 375–378 (1997).
- <sup>44</sup>B. N. Rajani, A. Kandasamy, and S. Majumdar, "Numerical simulation of laminar flow past a circular cylinder," *Appl. Math. Modell.* **33**(3), 1228–1247 (2009).
- <sup>45</sup>M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Trans. Signal Process.* **45**(11), 2673–2681 (1997).
- <sup>46</sup>S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning* (PMLR, 2015), pp. 448–456.