

Bridging Autoencoders and Dynamic Mode Decomposition for Reduced-order Modeling and Control of PDEs

Priyabrata Saha and Saibal Mukhopadhyay

Abstract—Modeling and controlling complex spatiotemporal dynamical systems driven by partial differential equations (PDEs) often necessitate dimensionality reduction techniques to construct lower-order models for computational efficiency. This paper explores a deep autoencoding learning method for reduced-order modeling and control of dynamical systems governed by spatiotemporal PDEs. We first analytically show that an optimization objective for learning a linear autoencoding reduced-order model can be formulated to yield a solution closely resembling the result obtained through the *dynamic mode decomposition with control* algorithm. We then extend this linear autoencoding architecture to a deep autoencoding framework, enabling the development of a nonlinear reduced-order model. Furthermore, we leverage the learned reduced-order model to design controllers using stability-constrained deep neural networks. Numerical experiments are presented to validate the efficacy of our approach in both modeling and control using the example of a reaction-diffusion system.

I. INTRODUCTION

Performing high-fidelity simulations of physical systems governed by partial differential equations (PDEs) incurs substantial computational costs, rendering subsequent tasks, such as control, extremely challenging if not infeasible. To address this, reduced-order models (ROMs) are often developed using dimensionality reduction techniques, enabling efficient simulation and control. For controlled dynamical systems, the reduced-order modeling approaches either combine analytical techniques with empirical approximation [1] or are purely data-driven [2], [3], [4]. Among these, the dynamic mode decomposition (DMD) based methods have become widely popular in recent years due to a strong connection between DMD and Koopman operator theory [5]. Another recent research trend involves utilizing deep neural networks (DNNs), particularly autoencoders, for modeling and control of high-dimensional dynamical systems. Most research in this area focuses only on the modeling and prediction of such complex dynamics [6], [7], [8].

A second line of research, though relatively less prevalent than modeling and prediction, leverages deep learning for controlling PDE-driven systems. Deep reinforcement learning (RL) is one such approach utilized to learn control policies for these systems [9], [10]. However, model-free RL methods require running numerical solvers in every iteration to provide feedback to the agents, which is computationally expensive. The same concern arises for the methods involving differentiable simulators as in [11], [12]. The alternative

to model-free methods for control design takes the traditional approach: develop a model first and then use that to design controllers [13], [14], [15]. These model-based methods constrain the latent dynamic models to be linear and work well within a short time-window. However, linear combination of a finite number of dynamic modes may not accurately represent the long-term nonlinear characteristics of complex dynamics [16] and adding nonlinear forcing terms yields better prediction accuracy [7]. It is also necessary to update the linear ROMs with online observations during operation for better prediction accuracy. Accordingly, model-based approaches typically leverage the model-predictive control (MPC) scheme to optimize the control policy online using the updated ROMs. Running online optimization at each step may not be computationally feasible in many cases.

In this paper, we explore the potential of nonlinear ROMs to achieve more accurate long-term predictions, enabling the use of offline control learning methods. Specifically, we develop autoencoder-based ROMs for PDE-driven controlled dynamical systems and leverage these ROMs to learn control policies for the original systems. Our main contributions are:

- Drawing inspiration from the *dynamic mode decomposition with control* (DMDc) algorithm [4], we propose an autoencoding framework that effectively captures the dynamic characteristics of the underlying systems, which is crucial for the controller's performance on the full-order system.
- We formulate an objective function for data-driven model learning in a linear autoencoding configuration and show analytically that this function encourages a linear ROM resembling the model obtained via DMDc.
- We configure the linear autoencoding architecture such that its components can be substituted with DNNs and optimized via gradient descent to obtain a nonlinear ROM. We show that DNN-based nonlinear ROMs yield more accurate long-term predictions and facilitates offline control learning for the underlying system.
- We leverage the learned nonlinear ROM to design controllers using stability-constrained DNNs, ensuring the stability of the control policies.

Section II provides the problem statement and background on DMDc. Section III details the linear autoencoding framework, its connection to DMDc, extends it to nonlinear autoencoders, and outlines the control learning algorithm. Section IV presents numerical experiments to validate the proposed methods, and Section V concludes the paper with a summary of findings and potential future research directions.

The authors are with School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, USA. Emails: priyabratasaha@gatech.edu, saibal.mukhopadhyay@ece.gatech.edu

II. PROBLEM AND PRELIMINARIES

A. Problem statement

Consider a time-invariant controlled dynamical system driven by a PDE

$$\frac{\partial \mathcal{X}}{\partial t} = \mathcal{M}\left(\mathcal{X}, \frac{\partial \mathcal{X}}{\partial \zeta}, \frac{\partial^2 \mathcal{X}}{\partial \zeta^2}, \dots, \mathbf{u}\right), \quad (1)$$

where $\mathcal{X}(\zeta, t) \in \mathbb{R}$ is the system state at location ζ and time t , and $\mathbf{u}(t) \in \mathbb{U} \subset \mathbb{R}^{d_u}$ is the actuation (or control input) at time t . \mathcal{M} is a linear or nonlinear function of the system state, its spatial derivatives of different orders, and the actuation. Space discretization of the state leads to a system of ordinary differential equations (ODEs) that can be written as

$$\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \mathbf{u}). \quad (2)$$

Here $\mathbf{x}(t) \in \mathbb{X} \subset \mathbb{R}^{d_x}$, $d_x \gg 1$ is the space-discretized state vector at time t . Our objective is to learn a reduced-order model for this high-dimensional ($d_x \gg 1$) system of (2) and use that ROM to learn a feedback controller $\mathbf{u} = \pi(\mathbf{x})$ that stabilizes the system at a desired state. We consider a data-driven learning scenario and assume that we have observations from the system consisting of time series data $\mathbf{x}(t_i), i = 0, 1, \dots, n$ subjected to random values of actuations $\mathbf{u}(t_i), i = 0, 1, \dots, (n-1)$. Note, we use v (in place of $v(t)$ for brevity) as notation for any continuous-time variable (e.g., system state, control input), whereas $v(t_i)$ is used to denote their discrete sample at time instance t_i . We further assume that the system we are aiming to stabilize at an equilibrium point is *locally stabilizable*, i.e., there exists a control policy such that the desired state is *asymptotically stable* for the closed-loop system.

B. Dynamic mode decomposition with control

DMD [17] is a data-driven method that reconstructs the underlying dynamics using only a time series of snapshots from the system. DMDc [4] is an extension of DMD for dynamical systems with control. DMDc seeks best-fit linear operators \mathbf{A} and \mathbf{B} between successive observed states and the actuations:

$$\hat{\mathbf{x}}(t_{i+1}) = \mathbf{A}\mathbf{x}(t_i) + \mathbf{B}\mathbf{u}(t_i), \quad i = 0, 1, \dots, n-1, \quad (3)$$

where $\hat{\mathbf{x}}(t)$ denotes an approximation of $\mathbf{x}(t)$, $\mathbf{A} \in \mathbb{R}^{d_x \times d_x}$, and $\mathbf{B} \in \mathbb{R}^{d_x \times d_u}$. Direct analysis of (3) could be computationally prohibitive for $d_x \gg 1$. DMDc leverages dimensionality reduction to compute a ROM

$$\mathbf{x}_{\text{R,DMDc}}(t_i) = \mathbf{E}_{\text{DMDc}}\mathbf{x}(t_i), \quad (4a)$$

$$\mathbf{x}_{\text{R,DMDc}}(t_{i+1}) = \mathbf{A}_{\text{R,DMDc}}\mathbf{x}_{\text{R,DMDc}}(t_i) + \mathbf{B}_{\text{R,DMDc}}\mathbf{u}(t_i), \quad (4b)$$

$$i = 0, 1, \dots, n-1,$$

which retains the dominant dynamic modes of (3). Here, $\mathbf{x}_{\text{R,DMDc}}(t_i) \in \mathbb{R}^{r_x}$ is the reduced state, where $r_x \ll d_x$, and $\mathbf{E}_{\text{DMDc}} \in \mathbb{R}^{r_x \times d_x}$, $\mathbf{A}_{\text{R,DMDc}} \in \mathbb{R}^{r_x \times r_x}$, $\mathbf{B}_{\text{R,DMDc}} \in \mathbb{R}^{r_x \times d_u}$. The full state is reconstructed from the reduced state using the transformation $\hat{\mathbf{x}}(t_i) = \mathbf{D}_{\text{DMDc}}\mathbf{x}_{\text{R,DMDc}}(t_i)$,

where $\mathbf{D}_{\text{DMDc}} \in \mathbb{R}^{d_x \times r_x}$. DMDc computes truncated singular value decomposition (SVD) of the data matrices $\mathbf{Y} = [\mathbf{x}(t_1), \mathbf{x}(t_2), \dots, \mathbf{x}(t_n)] \in \mathbb{R}^{d_x \times n}$ and $\mathbf{\Omega} = [\boldsymbol{\omega}(t_0), \boldsymbol{\omega}(t_1), \dots, \boldsymbol{\omega}(t_{n-1})] \in \mathbb{R}^{(d_x+d_u) \times n}$, $\boldsymbol{\omega}(t_i) = [\mathbf{x}(t_i)^\top, \mathbf{u}(t_i)^\top]^\top \in \mathbb{R}^{d_x+d_u}$ as follows:

$$\mathbf{Y} = \hat{\mathbf{U}}_{\mathbf{Y}} \hat{\boldsymbol{\Sigma}}_{\mathbf{Y}} \hat{\mathbf{V}}_{\mathbf{Y}}^\top, \quad \mathbf{\Omega} = \hat{\mathbf{U}}_{\boldsymbol{\Omega}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\Omega}} \hat{\mathbf{V}}_{\boldsymbol{\Omega}}^\top, \quad (5)$$

where $\hat{\mathbf{U}}_{\mathbf{Y}} \in \mathbb{R}^{d_x \times r_x}$, $\hat{\boldsymbol{\Sigma}}_{\mathbf{Y}} \in \mathbb{R}^{r_x \times r_x}$, $\hat{\mathbf{V}}_{\mathbf{Y}} \in \mathbb{R}^{n \times r_x}$, $\hat{\mathbf{U}}_{\boldsymbol{\Omega}} \in \mathbb{R}^{(d_x+d_u) \times r_{xu}}$, $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\Omega}} \in \mathbb{R}^{r_{xu} \times r_{xu}}$, and $\hat{\mathbf{V}}_{\boldsymbol{\Omega}} \in \mathbb{R}^{n \times r_{xu}}$. Here, $r_x < \min(d_x, n)$ and $r_{xu} < \min(d_x + d_u, n)$ denote the truncation dimensions of SVDs. Utilizing the SVDs of (5), the parameters of the ROM (4) is obtained as

$$\mathbf{E}_{\text{DMDc}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top, \quad \mathbf{D}_{\text{DMDc}} = \hat{\mathbf{U}}_{\mathbf{Y}}, \quad (6a)$$

$$\mathbf{A}_{\text{R,DMDc}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \hat{\mathbf{V}}_{\boldsymbol{\Omega}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\Omega}}^{-1} \hat{\mathbf{U}}_{\boldsymbol{\Omega},1}^\top \hat{\mathbf{U}}_{\mathbf{Y}}, \quad (6b)$$

$$\mathbf{B}_{\text{R,DMDc}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \hat{\mathbf{V}}_{\boldsymbol{\Omega}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\Omega}}^{-1} \hat{\mathbf{U}}_{\boldsymbol{\Omega},2}^\top, \quad (6c)$$

where $\hat{\mathbf{U}}_{\boldsymbol{\Omega},1} \in \mathbb{R}^{d_x \times r_{xu}}$, $\hat{\mathbf{U}}_{\boldsymbol{\Omega},2} \in \mathbb{R}^{d_u \times r_{xu}}$, and $\hat{\mathbf{U}}_{\boldsymbol{\Omega}}^\top = [\hat{\mathbf{U}}_{\boldsymbol{\Omega},1}^\top \quad \hat{\mathbf{U}}_{\boldsymbol{\Omega},2}^\top]$.

III. METHOD

A. Learning a reduced order model

To develop a nonlinear ROM utilizing DNNs that effectively capture the underlying dynamics, we first investigate if we can obtain a linear ROM similar to DMDc, in a gradient descent arrangement. Specifically, we analyze optimization objectives that encourage a DMDc-like solution for a reduced-order modeling problem using linear networks (single layer network without nonlinear activation). Consider the following reduced-order modeling problem

$$\mathbf{x}_{\text{R}}(t_i) = \mathbf{E}_{\mathbf{x}}\mathbf{x}(t_i), \quad (7a)$$

$$\mathbf{x}_{\text{R}}(t_{i+1}) = \mathbf{A}_{\text{R}}\mathbf{x}_{\text{R}}(t_i) + \mathbf{B}_{\text{R}}\mathbf{u}(t_i), \quad (7b)$$

$$\hat{\mathbf{x}}(t_i) = \mathbf{D}_{\mathbf{x}}\mathbf{x}_{\text{R}}(t_i), \quad i = 0, 1, \dots, n-1, \quad (7c)$$

where the linear operators $\mathbf{E}_{\mathbf{x}} \in \mathbb{R}^{r_x \times d_x}$ and $\mathbf{D}_{\mathbf{x}} \in \mathbb{R}^{d_x \times r_x}$ projects and reconstructs back, respectively, the high-dimensional system state to and from a low-dimensional feature $\mathbf{x}_{\text{R}} \in \mathbb{R}^{r_x}$. The linear operators $\mathbf{A}_{\text{R}} \in \mathbb{R}^{r_x \times r_x}$ and $\mathbf{B}_{\text{R}} \in \mathbb{R}^{r_x \times d_u}$ describe the relations between successive reduced states and actuations. We refer to this reduced-order model with linear networks as *linear autoencoding ROM* or LAROM. In the following, we first analyze the solution of the optimization objective of LAROM for a fixed *encoder* $\mathbf{E}_{\mathbf{x}}$. Then we establish a connection between the solution of LAROM and the solution of DMDc, and further discuss the choice of the encoder to promote similarity between the two. Finally, we extend the linear model to a DNN-based model, which we refer to as DeepROM.

1) *Analysis of the linear reduced-order model for a fixed encoder:* The DMDc algorithm essentially solves for $\tilde{\mathbf{G}} \in \mathbb{R}^{r_x \times (d_x+d_u)}$ to minimize $\frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_{\mathbf{x}}\mathbf{x}(t_{i+1}) - \tilde{\mathbf{G}}\boldsymbol{\omega}(t_i)\|^2$ for a fixed projection matrix $\mathbf{E}_{\mathbf{x}} = \mathbf{E}_{\text{DMDc}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top$. Here, $\boldsymbol{\omega}(t_i)$ is the concatenated vector of state and actuation as defined in section II-B. The optimal solution $\tilde{\mathbf{G}}_{\text{opt}}$ is then partitioned as $[\tilde{\mathbf{A}} \quad \tilde{\mathbf{B}}]$ such that $\tilde{\mathbf{A}} \in \mathbb{R}^{r_x \times d_x}$, $\tilde{\mathbf{B}} \in \mathbb{R}^{r_x \times d_u}$.

Finally, $\tilde{\mathbf{A}}$ is post-multiplied with the reconstruction operator $\mathbf{D}_{\text{DMDc}} = \hat{\mathbf{U}}_Y$ to get the ROM components $\mathbf{A}_{\text{R,DMDc}}$ and $\mathbf{B}_{\text{R,DMDc}}$. Note, that the final step of this process (multiplication of the operators) is feasible only for the linear case, not in the case when the projection and reconstruction operators are nonlinear (e.g. DNNs). Therefore, we use an alternative formulation with the following results to design a loss function that encourages a DMDc-like solution for (7) and also offers dimensionality reduction when nonlinear components are used.

Theorem 1: Consider the following objective function

$$L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) = \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x \mathbf{x}(t_{i+1}) - \mathbf{G} \mathbf{E}_{x_u} \boldsymbol{\omega}(t_i)\|^2, \quad (8)$$

where $\mathbf{G} = [\mathbf{A}_R \quad \mathbf{B}_R] \in \mathbb{R}^{r_x \times (r_x + d_u)}$, $\mathbf{E}_{x_u} = \begin{bmatrix} \mathbf{E}_x & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix} \in \mathbb{R}^{(r_x + d_u) \times (d_x + d_u)}$, \mathbf{I}_{d_u} being the identity matrix of order d_u . For any fixed matrix \mathbf{E}_x , the objective function L_{pred} is convex in the coefficients of \mathbf{G} and attains its minimum for any \mathbf{G} satisfying

$$\mathbf{G} \mathbf{E}_{x_u} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{x_u}^\top = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top \mathbf{E}_{x_u}^\top, \quad (9)$$

where \mathbf{Y} and $\boldsymbol{\Omega}$ are the data matrices as defined in section II-B. If \mathbf{E}_x has full rank r_x , and $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top$ is non-singular, then L_{pred} is strictly convex and has a unique minimum for

$$\mathbf{G} = [\mathbf{A}_R \quad \mathbf{B}_R] = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top \mathbf{E}_{x_u}^\top (\mathbf{E}_{x_u} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{x_u}^\top)^{-1}. \quad (10)$$

Proof. See appendix I-A.

Remark 1: For a unique solution, we assume that \mathbf{E}_x has full rank. The other scenario, i.e., \mathbf{E}_x is rank-deficient suggests poor utilization of the hidden units of the model. In that case, the number of hidden units (which represents the dimension of the reduced state) should be decreased. The assumption that the covariance matrix $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top$ is invertible can be ensured when $n \geq d_x + d_u$, by removing any linearly dependent features in system state and actuation. When $n < d_x + d_u$, the covariance matrix $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top$ is not invertible. However, similar results can be obtained by adding ℓ_2 regularization (for the coefficients/entries of \mathbf{G}) to the objective function.

2) *The connection between the solutions of the linear autoencoding model and DMDc:* The connection between the ROM obtained by minimizing L_{pred} (for a fixed \mathbf{E}_x), i.e., (10) and the DMDc ROM of (6) is not readily apparent. To interpret the connection, we formulate an alternative representation of (10) utilizing the SVD and the Moore-Penrose inverse of matrices. This alternative representation leads to the following result.

Corollary 1.1: Consider the (full) SVD of the data matrix $\boldsymbol{\Omega}$ given by $\boldsymbol{\Omega} = \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \mathbf{V}_\Omega^\top$, where $\mathbf{U}_\Omega \in \mathbb{R}^{(d_x + d_u) \times (d_x + d_u)}$, $\boldsymbol{\Sigma}_\Omega \in \mathbb{R}^{(d_x + d_u) \times n}$, and $\mathbf{V}_\Omega \in \mathbb{R}^{n \times n}$. If $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$ and $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top$ is non-singular, then the solution for $\mathbf{G} = [\mathbf{A}_R \quad \mathbf{B}_R]$ corresponding to the unique minimum of L_{pred} can be expressed as

$$\mathbf{A}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}^* \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y, \quad \mathbf{B}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}^* \mathbf{U}_{\Omega,2}^\top, \quad (11)$$

where $[\mathbf{U}_{\Omega,1}^\top \quad \mathbf{U}_{\Omega,2}^\top] = \mathbf{U}_\Omega^\top$ with $\mathbf{U}_{\Omega,1} \in \mathbb{R}^{d_x \times (d_x + d_u)}$, $\mathbf{U}_{\Omega,2} \in \mathbb{R}^{d_u \times (d_x + d_u)}$, and $\boldsymbol{\Sigma}^* = \lim_{\varepsilon \rightarrow 0} (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \boldsymbol{\Sigma}_\Omega + \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \boldsymbol{\Sigma}_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top$.

Proof. See appendix I-C.

Remark 2: It can be verified easily that if we use the truncated SVD (as defined by 5), instead of the full SVD, for $\boldsymbol{\Omega}$ in corollary 1.1, we get an approximation of (11):

$$\hat{\mathbf{A}}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \hat{\mathbf{V}}_\Omega \hat{\boldsymbol{\Sigma}}^* \hat{\mathbf{U}}_{\Omega,1}^\top \hat{\mathbf{U}}_Y, \quad \hat{\mathbf{B}}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \hat{\mathbf{V}}_\Omega \hat{\boldsymbol{\Sigma}}^* \hat{\mathbf{U}}_{\Omega,2}^\top, \quad (12)$$

where $\hat{\boldsymbol{\Sigma}}^* = \lim_{\varepsilon \rightarrow 0} (\hat{\boldsymbol{\Sigma}}_\Omega^\top \hat{\mathbf{U}}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \hat{\mathbf{U}}_{\Omega,1} \hat{\boldsymbol{\Sigma}}_\Omega + \hat{\boldsymbol{\Sigma}}_\Omega^\top \hat{\mathbf{U}}_{\Omega,2}^\top \hat{\mathbf{U}}_{\Omega,2} \hat{\boldsymbol{\Sigma}}_\Omega + \varepsilon^2 \mathbf{I}_{r_{x_u}})^{-1} \hat{\boldsymbol{\Sigma}}_\Omega^\top$. We can see that (12) has the same form as (6b), except $\hat{\boldsymbol{\Sigma}}_\Omega^{-1}$ is replaced with $\hat{\boldsymbol{\Sigma}}^*$.

All the aforementioned results are derived for a fixed \mathbf{E}_x and the relation to the DMDc is specific to the case $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$. Note that the columns of the $\hat{\mathbf{U}}_Y$ are the left singular vectors, corresponding to the leading singular values, of \mathbf{Y} . Equivalently, those are also the eigenvectors, corresponding to the leading eigenvalues, of the covariance matrix $\mathbf{Y} \mathbf{Y}^\top$. L_{pred} alone does not constrain \mathbf{E}_x to take a similar form and we need another loss term to encourage such form for the encoder. To this end, we build on the findings in [18], which explore the similarity between principle component analysis and linear autoencoders optimized with the objective function: $L_{\text{recon}}(\mathbf{E}_x, \mathbf{D}_x) = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}(t_i) - \mathbf{D}_x \mathbf{E}_x \mathbf{x}(t_i)\|^2$. It is shown in [18] that all the critical points of L_{recon} correspond to projections onto subspaces associated with subsets of eigenvectors of the covariance matrix $\mathbf{Y} \mathbf{Y}^\top$. Moreover, L_{recon} has a unique global minimum corresponding to the first r_x (i.e., the desired dimension of the reduced state) number of eigenvectors of $\mathbf{Y} \mathbf{Y}^\top$, associated with the leading r_x eigenvalues. In other words, for any invertible matrix $\mathbf{C} \in \mathbb{R}^{r_x \times r_x}$, $\mathbf{D}_x = \mathbf{U}_{r_x} \mathbf{C}$ and $\mathbf{E}_x = \mathbf{C}^{-1} \mathbf{U}_{r_x}^\top$ globally minimizes L_{recon} , where \mathbf{U}_{r_x} denotes the matrix containing leading r_x eigenvectors of $\mathbf{Y} \mathbf{Y}^\top$. Since the left singular vectors of \mathbf{Y} are the eigenvectors of $\mathbf{Y} \mathbf{Y}^\top$, we have $\mathbf{U}_{r_x} = \hat{\mathbf{U}}_Y$. Hence, we consider to utilize L_{recon} to promote learning an encoder \mathbf{E}_x in the form of $\mathbf{C}^{-1} \hat{\mathbf{U}}_Y^\top$. Accordingly, we propose to minimize the following objective function to encourage a DMDc-like solution for LAROM:

$$L(\mathbf{E}_x, \mathbf{D}_x, \mathbf{G}) = L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) + \beta_1 L_{\text{recon}}(\mathbf{E}_x, \mathbf{D}_x), \quad (13)$$

where $\beta_1 > 0$ is a tunable hyperparameter.

It is important to note that L_{recon} is minimized for any invertible matrix \mathbf{C} , $\mathbf{D}_x = \hat{\mathbf{U}}_Y \mathbf{C}$, and $\mathbf{E}_x = \mathbf{C}^{-1} \hat{\mathbf{U}}_Y^\top$. When optimized using gradient descent, it is highly unlikely to get \mathbf{C} as the identity matrix like DMDc. Rather, we expect a random \mathbf{C} . Therefore, we need additional constraints to promote similarity with DMDc. For this purpose, we tie the matrices \mathbf{E}_x and \mathbf{D}_x to be the transpose of each other and add a semi-orthogonality constraint $\beta_4 \|\mathbf{E}_x \mathbf{E}_x^\top - \mathbf{I}_{r_x}\|$, $\beta_4 > 0$ to the optimization objective of (13).

3) *Extending the linear model to a deep model:* Here, we discuss the process of extending LAROM to a non-

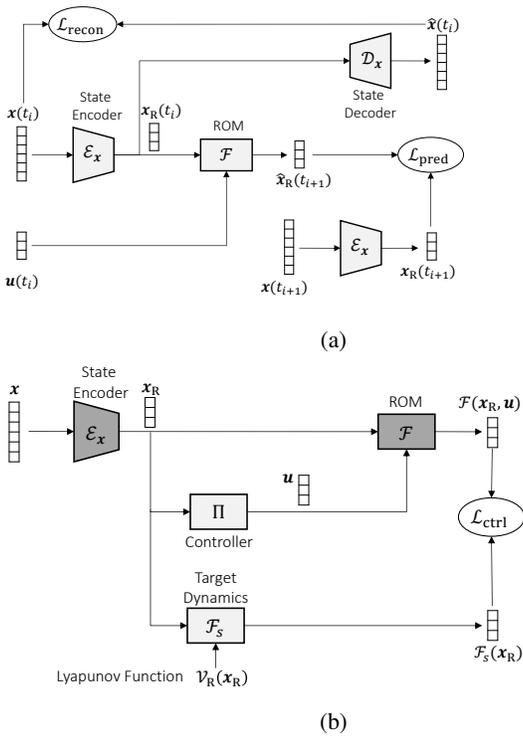


Fig. 1: (a): Autoencoding architecture for reduced-order modeling. The state encoder \mathcal{E}_x reduces the dimension of the state. The ROM \mathcal{F} takes the current reduced state and actuation to predict the next reduced state, which is then uplifted to the full state by the state decoder \mathcal{D}_x . All modules are trained together using a combined loss involving $\mathcal{L}_{\text{pred}}$ and $\mathcal{L}_{\text{recon}}$. (b): The control learning process. Given a reduced state, \mathcal{F}_s predicts a target dynamics for the closed-loop system, and the controller Π predicts an actuation to achieve that target. Both the modules are trained jointly using the loss function $\mathcal{L}_{\text{ctrl}}$. Parameters of the dark-shaded modules are kept fixed during this process.

linear reduced-order modeling framework. We replace all the trainable components of LAROM, i.e., \mathbf{E}_x , \mathbf{D}_x , and \mathbf{G} , with DNNs. Specifically, we use an encoding function or *encoder* $\mathcal{E}_x : \mathbb{X} \rightarrow \mathbb{R}^{r_x}$ and a decoding function or *decoder* $\mathcal{D}_x : \mathbb{R}^{r_x} \rightarrow \mathbb{X}$ to transform the high-dimensional system state to low-dimensional features and reconstruct it back, respectively, i.e., $\mathbf{x}_R = \mathcal{E}_x(\mathbf{x})$, $\hat{\mathbf{x}} = \mathcal{D}_x(\mathbf{x}_R)$, where $\mathbf{x}_R \in \mathbb{R}^{r_x}$ denotes the reduced state, and $\hat{\mathbf{x}}$ is the reconstruction of \mathbf{x} . The encoded state and control are then fed to another DNN that represents the reduced order dynamics $\frac{d\mathbf{x}_R}{dt} = \mathcal{F}(\mathbf{x}_R, \mathbf{u})$, where $\mathcal{F} : \mathbb{R}^{r_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{r_x}$. Given the current reduced state $\mathbf{x}_R(t_i)$ and control input $\mathbf{u}(t_i)$, the next reduced state $\mathbf{x}_R(t_{i+1})$ can be computed by integrating \mathcal{F} using a numerical integrator: $\mathbf{x}_R(t_{i+1}) = \mathbf{x}_R(t_i) + \int_{t_i}^{t_{i+1}} \mathcal{F}(\mathbf{x}_R(t), \mathbf{u}(t)) dt \triangleq \mathcal{G}(\mathbf{x}_R(t_i), \mathbf{u}(t_i))$. We can say that \mathcal{G} is the nonlinear counterpart of \mathbf{G} .

Note, here the ROM is represented as a continuous-time dynamics, unlike the linear case where we used a discrete-time model. We use a discrete-time formulation for LAROM to establish its similarity with DMDc, which is

formulated in discrete time. DeepROM can be formulated in a similar fashion as well. However, the specific control learning algorithm we used, which will be discussed in the next subsection, requires vector fields of the learned ROM for training. Therefore, we formulate the ROM in continuous time so that it provides the vector field $\mathcal{F}(\mathbf{x}_R, \mathbf{u})$ of the dynamics. In cases where only the prediction model is of interest and control learning is not required, a discrete-time formulation should be used for faster training of the ROM.

We train \mathcal{E}_x , \mathcal{D}_x , and \mathcal{F} by minimizing the following loss function, analogous to (13),

$$\mathcal{L}(\mathcal{E}_x, \mathcal{D}_x, \mathcal{F}) = \mathcal{L}_{\text{pred}}(\mathcal{E}_x, \mathcal{F}) + \beta_2 \mathcal{L}_{\text{recon}}(\mathcal{E}_x, \mathcal{D}_x), \quad (14)$$

where $\beta_2 > 0$ is a tunable hyperparameter. $\mathcal{L}_{\text{pred}}$ and $\mathcal{L}_{\text{recon}}$ are defined as $\mathcal{L}_{\text{pred}}(\mathcal{E}_x, \mathcal{F}) = \frac{1}{n} \sum_{i=0}^{n-1} \left\| \mathcal{E}_x(\mathbf{x}(t_{i+1})) - \mathcal{G}(\mathcal{E}_x(\mathbf{x}(t_i)), \mathbf{u}(t_i)) \right\|^2$ and $\mathcal{L}_{\text{recon}}(\mathcal{E}_x, \mathcal{D}_x) = \frac{1}{n} \sum_{i=1}^n \left\| \mathbf{x}(t_i) - \mathcal{D}_x \circ \mathcal{E}_x(\mathbf{x}(t_i)) \right\|^2$. Here, the operator \circ denotes the composition of two functions. In experiments, $\mathcal{L}_{\text{recon}}$ also includes the reconstruction loss of the desired state where we want to stabilize the system. Figure 1a shows the overall framework for training DeepROM.

B. Learning control

Once we get a trained ROM of the form $\frac{d\mathbf{x}_R}{dt} = \mathcal{F}(\mathbf{x}_R, \mathbf{u})$ using the method proposed in section III-A, the next goal is to design a controller for the system utilizing that ROM. Since our ROM is represented by DNNs, we need a data-driven method to develop the controller. We adopt the approach presented in [19] for learning control laws for nonlinear systems, represented by DNNs. The core idea of the method is to hypothesize a target dynamics that ensures exponential stability of the desired state, and simultaneously learn a control policy to realize that target dynamics in the closed loop. A DNN is used to represent the vector field $\mathcal{F}_s : \mathbb{R}^{r_x} \rightarrow \mathbb{R}^{r_x}$ of the target dynamics $\frac{d\mathbf{x}_R}{dt} = \mathcal{F}_s(\mathbf{x}_R)$. We use another DNN to represent a controller $\Pi : \mathbb{R}^{r_x} \rightarrow \mathbb{R}^{d_u}$ that provides the necessary actuation for a given reduced state \mathbf{x}_R : $\mathbf{u} = \Pi(\mathbf{x}_R)$. Finally, the reduced state \mathbf{x}_R and actuation \mathbf{u} are fed to the (trained) ROM of $\frac{d\mathbf{x}_R}{dt} = \mathcal{F}(\mathbf{x}_R, \mathbf{u})$ to get $\mathcal{F}(\mathbf{x}_R, \mathbf{u})$. The overall framework for learning control is referred to as *deep reduced-order control* (DeepROC) and is shown in Figure 1b.

Our training objective is to minimize the difference between $\mathcal{F}(\mathbf{x}_R, \mathbf{u})$ and $\mathcal{F}_s(\mathbf{x}_R)$, i.e.,

$$\mathcal{L}_{\text{ctrl}}(\mathcal{F}_s, \Pi) = \frac{1}{n} \sum_{i=1}^n \left\| \mathcal{F}(\mathcal{E}_x(\mathbf{x}(t_i)), \Pi \circ \mathcal{E}_x(\mathbf{x}(t_i))) - \mathcal{F}_s \circ \mathcal{E}_x(\mathbf{x}(t_i)) \right\|^2. \quad (15)$$

To minimize the control effort, we add a regularization loss with (15), and the overall training objective for learning control is given by

$$\mathcal{L}_{\text{ctrl,reg}}(\mathcal{F}_s, \Pi) = \mathcal{L}_{\text{ctrl}}(\mathcal{F}_s, \Pi) + \beta_3 \frac{1}{n} \sum_{i=1}^n \left\| \Pi(\mathbf{x}_R(t_i)) \right\|^2, \quad (16)$$

where $\beta_3 > 0$ is a tunable hyperparameter. Here we jointly train the DNNs representing Π and \mathcal{F}_s only, whereas the previously-trained DNNs for \mathcal{E}_x and \mathcal{F} are kept frozen. Once all the DNNs are trained, we only need \mathcal{E}_x and Π during evaluation to generate actuation for the actual system, given a full-state observation: $\mathbf{u} = \Pi \circ \mathcal{E}_x(\mathbf{x}) = \pi(\mathbf{x})$. As we mentioned earlier, we require that the target dynamics, hypothesized by a DNN, ensures exponential stability of the desired state. Without loss of generality, we consider stability at $\mathbf{x}_R = \mathbf{0}$. The system can be stabilized at any desired state by adding a feedforward component to the control. Dynamics represented by a standard neural network does not generally guarantee stability of any equilibrium point. However, it is possible to design a DNN, by means of Lyapunov functions, to represent a dynamics that ensures exponential stability of an equilibrium point, as demonstrated in [20]. Accordingly, we represent our target dynamics as follows:

$$\frac{d\mathbf{x}_R}{dt} = \mathcal{F}_s(\mathbf{x}_R) = \mathcal{P}(\mathbf{x}_R) - \frac{\text{ReLU}(\nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R) + \alpha \mathcal{V}_R(\mathbf{x}_R))}{\|\nabla \mathcal{V}_R(\mathbf{x}_R)\|^2} \nabla \mathcal{V}_R(\mathbf{x}_R), \quad (17)$$

where α is a positive constant, $\text{ReLU}(z) = \max(0, z)$, $z \in \mathbb{R}$, $\mathcal{P} : \mathbb{R}^{r_x} \rightarrow \mathbb{R}^{r_x}$ is an arbitrary function represented by a DNN, and $\mathcal{V}_R : \mathbb{R}^{r_x} \rightarrow \mathbb{R}$ is a candidate Lyapunov function. We use

$$\mathcal{V}_R(\mathbf{x}_R) = \mathbf{x}_R^\top \mathbf{K} \mathbf{x}_R, \quad (18)$$

where $\mathbf{K} \in \mathbb{R}^{r_x \times r_x}$ is a positive definite matrix. It is shown in [20] that the origin is exponentially stable for the target dynamics of (17) for any arbitrary DNN \mathcal{P} .

IV. NUMERICAL EXPERIMENTS

We consider the Newell-Whitehead-Segel reaction-diffusion equation

$$\begin{aligned} \frac{\partial q}{\partial t} &= \sigma \nabla^2 q + q(1 - q^2) + \mathbf{1}_{\mathbb{W}} w \quad \text{in } \mathbb{I} \times \mathbb{R}^+, \\ \nabla q(\zeta_l, t) &= \nabla q(\zeta_r, t) = 0, \quad t \in \mathbb{R}^+, \end{aligned} \quad (19)$$

which is used to describe various nonlinear physical systems including Rayleigh-Bénard convection. In (19), $q(\zeta, t) \in \mathbb{R}$ denotes the measurement variable such as concentration or temperature at location $\zeta \in \mathbb{I} \subset \mathbb{R}$ and time t ; σ denotes the diffusion coefficient; $w(t) \in \mathbb{R}$ is the actuation at time t and $\mathbf{1}_{\mathbb{W}}(\zeta)$ is the indicator function with $\mathbb{W} \subset \mathbb{I}$; ζ_l and ζ_r denote the boundary points of \mathbb{I} . The considered system is a bistable system with ± 1 as stable and 0 as unstable equilibria. For the control task, we consider feedback stabilization of this system at the unstable equilibrium 0, as studied in [21]. We use $\mathbb{I} = (-1, 1)$, $\mathbb{W} = (-0.2, 0.2)$, and $\sigma = 0.2$. Details on dataset generation, neural network architectures, and training settings are given in appendix II.

The prediction performance of DeepROM is compared against DMDC and the Deep Koopman model [13]. The Deep Koopman model shares a similar DNN-based autoencoding structure as ours, with the distinction that its (reduced-order) dynamic model is linear. The method proposed in [13]

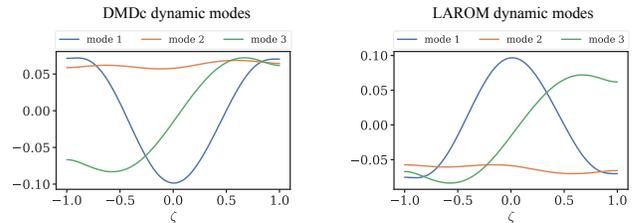


Fig. 2: The first three dynamic modes of the reaction-diffusion system, obtained using DMDC and LAROM.

considers a model predictive scenario, where the state/system matrix of the linear reduced-order model is updated with online observations during operation while the input/control matrix is kept fixed. However, in contrast to the original method, we keep both matrices fixed during the control operation as we consider offline control design in this paper. For the same reason, we apply linear quadratic regulator (LQR) on the ROM obtained from the Deep Koopman method, instead of model predictive control, to compare the control performance with our method: DeepROC. The control performance is also compared against the reduced-order controller obtained by applying LQR on the ROM derived from DMDC.

Similarity with DMDC. To investigate the similarity with DMDC, we first train LAROM using gradient descent to minimize the objective (13) with the semi-orthogonality regularization and enforcing $\mathbf{D}_x = \mathbf{E}_x^\top$, as discussed in III-A.2. The dynamic modes for LAROM are computed as $\varphi_i = \mathbf{D}_x \mathbf{z}_i$, where \mathbf{z}_i is the i^{th} eigenvector of \mathbf{A}_R . Similarly, the dynamic modes for DMDC are computed as $\varphi_{i,\text{DMDC}} = \mathbf{D}_{\text{DMDC}} \mathbf{z}_{i,\text{DMDC}}$, where $\mathbf{z}_{i,\text{DMDC}}$ is the i^{th} eigenvector of $\mathbf{A}_{R,\text{DMDC}}$. Note, these dynamic modes are similar to the ones used in the original DMD algorithm [17], not the exact modes obtained in [4]. Exact modes cannot be computed for LAROM since it does not involve SVD. Modes defined by $\varphi_{i,\text{DMDC}} = \mathbf{D}_{\text{DMDC}} \mathbf{z}_{i,\text{DMDC}} = \hat{\mathbf{U}}_Y \mathbf{z}_{i,\text{DMDC}}$ are the orthogonal projection of the exact modes onto the range of \mathbf{Y} (Theorem 3, [22]). Fig. 2 compares the dynamic modes obtained using DMDC and LAROM for the case when the dimension of the ROMs is 3. It is important to note that the numbering of the modes is arbitrary as the optimal ranking of DMDC modes is not trivial. The correspondence between the DMDC modes and LAROM modes are determined by comparing the eigenvalues of $\mathbf{A}_{R,\text{DMDC}}$ and \mathbf{A}_R . Dynamic modes of both methods are similar except for the different signs of the first two modes.

Prediction performance. We compare the performance of DeepROM, Deep Koopman model, and DMDC in the prediction task. Fig. 3(a) shows the quantitative comparison of the recursive multi-step predictions obtained using DMDC, Deep Koopman model, and DeepROM. The prediction error is computed as *normalized mean squared error* (NMSE) with respect to the solution obtained using the PDE solver. Prediction error increases more quickly for DMDC and Deep Koopman than DeepROM as the linear ROMs become less accurate in the long term. A qualitative comparison of

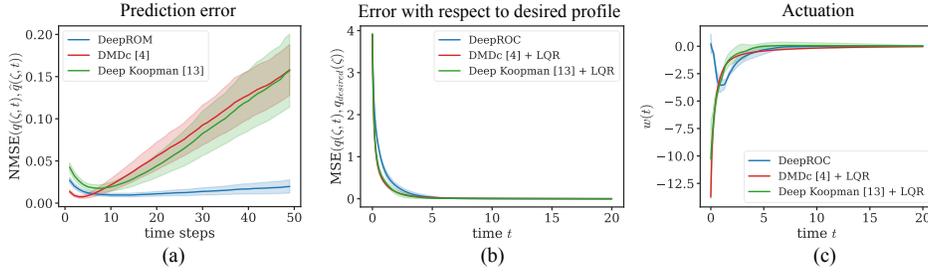


Fig. 3: (a): Prediction performance of different methods in the reaction–diffusion example. The shaded interval shows the 95% confidence interval around the mean from 100 test sequences and 3 different training instances. (b,c): Control performance of different methods in the reaction–diffusion example. The shaded interval shows the 1-standard deviation range around the mean from 3 different training instances.

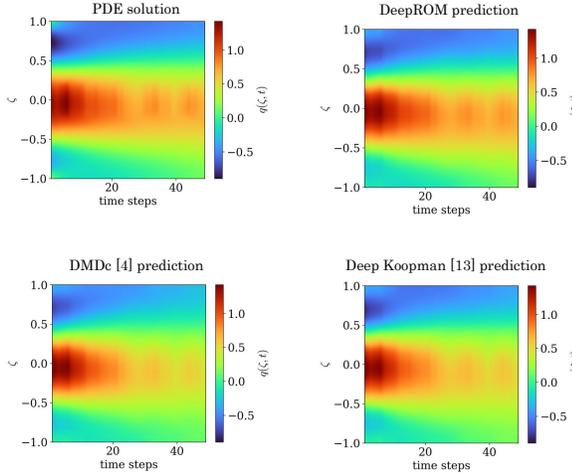


Fig. 4: Qualitative comparison of prediction performance of DMDc, Deep Koopman, and DeepROM for the reaction–diffusion system using one example sequence.

the prediction performance of the methods for an example sequence is shown in Fig. 4.

Control performance. Fig. 3(b,c) show the control performance of DeepROC, Deep Koopman + LQR, and DMDc + LQR in the task of stabilizing the system at the unstable equilibrium 0 from an initial state $2 + \cos(2\pi\zeta) \cos(\pi\zeta)$. We use the following two metrics for comparison:

- (i) mean squared error over time between the controlled solutions and the desired profile
- (ii) the amount of actuation applied

All methods show similar closed-loop error profiles. However, DeepROC requires significantly less amount of actuation in comparison with DMDc + LQR and Deep Koopman + LQR to reach a similar steady-state error. DeepROC can account for the decaying nonlinear term $-q^3$ present in the system (19) and therefore learns to apply less actuation. A qualitative comparison of the uncontrolled solution and the controlled solutions obtained using the three methods is shown in Fig. 5.

V. CONCLUSION

We presented a framework for autoencoder-based modeling and control learning for PDE-driven dynamical systems. The proposed reduced-order modeling framework is

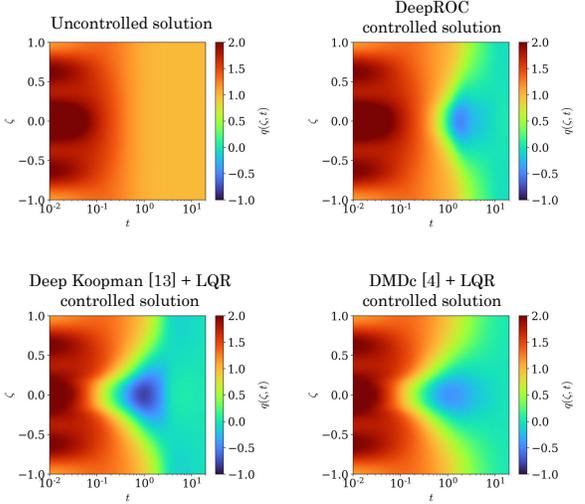


Fig. 5: Visual comparison of the uncontrolled solution and the controlled solutions of the reaction–diffusion system using DeepROC, Deep Koopman + LQR, and DMDc + LQR.

grounded on the connection between dynamic mode decomposition for controlled systems and a linear autoencoding architecture that can be trained using gradient descent. As demonstrated in our numerical experiments, DeepROM offers better prediction accuracy than a linear ROM over a relatively longer prediction horizon when applied to a non-linear system. Additionally, DeepROC requires less amount of actuation to reach the desired state for such a system.

APPENDIX I PROOFS

A. Proof of theorem 1

We can write $L_{\text{pred}}(\mathbf{E}_x, \mathbf{G})$ as follows,

$$\begin{aligned}
 L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) &= \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x \mathbf{x}(t_{i+1}) - \mathbf{G} \mathbf{E}_{x_u} \boldsymbol{\omega}(t_i)\|^2 \\
 &= \frac{1}{n} \|\text{vec}(\mathbf{E}_x \mathbf{Y}) - \text{vec}(\mathbf{G} \mathbf{E}_{x_u} \boldsymbol{\Omega})\|^2 \\
 &= \frac{1}{n} \|\text{vec}(\mathbf{E}_x \mathbf{Y}) - (\boldsymbol{\Omega}^\top \mathbf{E}_{x_u}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{G})\|^2, \tag{20}
 \end{aligned}$$

where \otimes and $\text{vec}(\cdot)$ denotes the Kronecker product and vectorization of matrices, respectively. The third equality is obtained using properties of the Kronecker product. Applying the principles of linear least squares optimization, we obtain the following outcome. For fixed \mathbf{E}_x , (20) is convex in coefficient of \mathbf{G} , and \mathbf{G} corresponds to a global minimum of L_{pred} if and only if

$$\begin{aligned} & (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})^\top (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{G}) \\ &= (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})^\top \text{vec}(\mathbf{E}_x \mathbf{Y}). \end{aligned} \quad (21)$$

Using the properties of the Kronecker product, we can write (21) as

$$(\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{G}) = (\mathbf{E}_{xu} \boldsymbol{\Omega} \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{E}_x \mathbf{Y}), \quad (22)$$

which yields $\mathbf{G} \mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top$, i.e., (9).

If \mathbf{E}_x has full rank r_x , then $\mathbf{E}_{xu} = \begin{bmatrix} \mathbf{E}_x & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix} \in \mathbb{R}^{(r_x+d_u) \times (d_x+d_u)}$ has full rank $(r_x + d_u)$. If $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top \in \mathbb{R}^{(d_x+d_u) \times (d_x+d_u)}$ is non-singular, then $\boldsymbol{\Omega}$ has full row-rank $(d_x + d_u)$. Consequently, applying the properties of matrix rank, we have

$$\begin{aligned} \text{rank}(\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top) &= \text{rank}(\mathbf{E}_{xu} \boldsymbol{\Omega}) \\ &= \text{rank}(\mathbf{E}_{xu}) = r_x + d_u. \end{aligned} \quad (23)$$

Hence, the symmetric positive semidefinite matrix $\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top$ has full rank and is therefore positive definite. Using the properties of the Kronecker product of positive definite matrices, we can see that $(\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})^\top (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) = (\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})$ is positive definite as well. Therefore, (20) is strictly convex in the coefficients of \mathbf{G} and has a unique minimum. Since $\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \succ 0$, it is invertible. Hence, from (9), we can say that the unique minimum of (20) is reached at $\mathbf{G} = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top (\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top)^{-1}$, i.e., (10). ■

B. An alternative representation of (10)

Here we provide an alternative representation of (10) required to prove corollary 1.1.

Lemma 2: Consider the (full) SVD of the data matrix $\boldsymbol{\Omega}$ given by $\boldsymbol{\Omega} = \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \mathbf{V}_\Omega^\top$, where $\mathbf{U}_\Omega \in \mathbb{R}^{(d_x+d_u) \times (d_x+d_u)}$, $\boldsymbol{\Sigma}_\Omega \in \mathbb{R}^{(d_x+d_u) \times n}$, and $\mathbf{V}_\Omega \in \mathbb{R}^{n \times n}$. (10) can be expressed as

$$\begin{aligned} \mathbf{G} &= \lim_{\varepsilon \rightarrow 0} \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \\ &\quad + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top. \end{aligned} \quad (24)$$

Proof. Replacing $\boldsymbol{\Omega}$ with its SVD in (10) we get,

$$\begin{aligned} \mathbf{G} &= \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top (\mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \mathbf{V}_\Omega^\top \mathbf{V}_\Omega \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top)^{-1} \\ &= \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top (\mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top)^{-1} \\ &= \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega (\mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega)^\dagger, \end{aligned} \quad (25)$$

where $(\cdot)^\dagger$ denotes the Moore-Penrose inverse of a matrix. The second equality is due to the orthogonality of \mathbf{V}_Ω .

Substituting $(\mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega)^\dagger$ with the *limit definition* (see [23]) of the Moore-Penrose inverse, we get

$$\begin{aligned} \mathbf{G} &= \lim_{\varepsilon \rightarrow 0} \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \\ &\quad + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top. \end{aligned} \quad (26)$$

C. Proof of Corollary 1.1

By the definition of truncated SVD, the columns of $\widehat{\mathbf{U}}_Y$ are orthonormal. Therefore, $\widehat{\mathbf{U}}_Y^\top$ has full row-rank r_x . Hence, by theorem 1 and lemma 2, if $\mathbf{E}_x = \widehat{\mathbf{U}}_Y^\top$, and $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top$ is non-singular, then the unique minimum of L_{pred} , is reached when

$$\begin{aligned} \mathbf{G} &= \widehat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega (\mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega)^\dagger \\ &= \lim_{\varepsilon \rightarrow 0} \widehat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \\ &\quad + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top. \end{aligned} \quad (27)$$

Now, substituting $\mathbf{E}_x = \widehat{\mathbf{U}}_Y^\top$ in \mathbf{E}_{xu} , and using the partition $\mathbf{U}_\Omega^\top = [\mathbf{U}_{\Omega,1}^\top \quad \mathbf{U}_{\Omega,2}^\top]$, where $\mathbf{U}_{\Omega,1} \in \mathbb{R}^{d_x \times (d_x+d_u)}$, $\mathbf{U}_{\Omega,2} \in \mathbb{R}^{d_u \times (d_x+d_u)}$, we get

$$\mathbf{E}_{xu} \mathbf{U}_\Omega = \begin{bmatrix} \widehat{\mathbf{U}}_Y^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{\Omega,1} \\ \mathbf{U}_{\Omega,2} \end{bmatrix} = \begin{bmatrix} \widehat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \\ \mathbf{U}_{\Omega,2} \end{bmatrix}, \quad (28)$$

and

$$\begin{aligned} \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega &= \begin{bmatrix} \mathbf{U}_{\Omega,1}^\top \widehat{\mathbf{U}}_Y & \mathbf{U}_{\Omega,2}^\top \end{bmatrix} \begin{bmatrix} \widehat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \\ \mathbf{U}_{\Omega,2} \end{bmatrix} \\ &= \mathbf{U}_{\Omega,1}^\top \widehat{\mathbf{U}}_Y \widehat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} + \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2}. \end{aligned} \quad (29)$$

Plugging (28) and (29) into (27) leads to

$$\begin{aligned} \mathbf{G} &= \lim_{\varepsilon \rightarrow 0} \widehat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,1}^\top \widehat{\mathbf{U}}_Y \widehat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \boldsymbol{\Sigma}_\Omega \\ &\quad + \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \boldsymbol{\Sigma}_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top \begin{bmatrix} \widehat{\mathbf{U}}_Y^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix}. \end{aligned} \quad (30)$$

Defining $\boldsymbol{\Sigma}^* \triangleq \lim_{\varepsilon \rightarrow 0} (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,1}^\top \widehat{\mathbf{U}}_Y \widehat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \boldsymbol{\Sigma}_\Omega + \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \boldsymbol{\Sigma}_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top$, we can split (30) into

$\mathbf{A}_R = \widehat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}^* \mathbf{U}_{\Omega,1}^\top \widehat{\mathbf{U}}_Y$, and $\mathbf{B}_R = \widehat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}^* \mathbf{U}_{\Omega,2}^\top$, which is (11). ■

APPENDIX II

DETAILS ON NUMERICAL EXPERIMENT

A. Dataset

We generate 100 training sequences of length 50 (i.e., $n = 50$) with time step size 0.01 and 256 nodes in \mathbb{I} using the finite element method. The initial conditions and actuations of these sequences are given by

$$q(\zeta, 0) = |a| \sum_{k=0}^4 b_k T_k(\zeta), \quad \zeta \in \mathbb{I}, \quad (31)$$

and

$$w(t_i) = 10g_i \max_{\zeta} |q(\zeta, t_{i-1})|, \quad i = 1, 2, \dots, 49, \quad (32)$$

where T_k denotes the k^{th} Chebyshev polynomial of the first kind, and $a \sim \mathcal{N}(0, 1)$, $b_k, g_i \sim \mathcal{U}(-1, 1)$ are chosen randomly. Similarly, 100 sequences are generated for the test set to evaluate the prediction performance.

B. DNN architectures

The state encoder comprises two 1D convolutional layers, followed by two fully connected (FC) layers. The first convolutional layer has 32 filters of size 3, utilizing a ReLU activation function and a stride of 2. The second convolutional layer maintains the same configuration but lacks the ReLU activation. The subsequent FC layer contains 64 neurons with ReLU activation. The final or output FC layer is composed of r_x neurons and does not use any bias. The arrangement of layers in the state decoder is inverted compared to that of the encoder, with transposed convolutional layers replacing the convolutional layers. The ROM is designed by breaking the function \mathcal{F} into two components: $\mathcal{F}(\mathbf{x}_R, \mathbf{u}) = \mathcal{F}_{\text{auto}}(\mathbf{x}_R) + \mathcal{F}_{\text{forced}}(\mathbf{x}_R, \mathbf{u}) - \mathcal{F}_{\text{forced}}(\mathbf{x}_R, \mathbf{0})$. $\mathcal{F}_{\text{auto}}$ represents the autonomous dynamics that does not depend on the actuation, whereas $\mathcal{F}_{\text{forced}}$ is responsible for the impact of actuation on dynamics. We observe this structure results in better performance in experiments than a single neural network representing $\mathcal{F}(\mathbf{x}_R, \mathbf{u})$. Two multilayer perceptrons (MLPs) are employed to realize $\mathcal{F}_{\text{auto}}$ and $\mathcal{F}_{\text{forced}}$, each consisting of three layers. The first and second layers each comprise 100 neurons with ReLU activation, while the output layer consists of r_x neurons. The output of the ROM is integrated using a numerical integrator to get the next state. The target dynamics is represented by an MLP akin to the one employed for $\mathcal{F}_{\text{auto}}$, followed by a stability criterion in the form of (17). The controller is implemented using another MLP with a similar three-layer configuration, while the output layer is adjusted to accommodate the actuation dimension.

For the encoder and decoder of the Deep Koopman model, we use the identical architectures as those employed in our DeepROM model.

C. Training settings

We use $r_x = 5$ in the prediction task and $r_x = 2$ in the control task for all the methods. All modules are implemented in PyTorch. In both of the learning phases, learning ROM and learning controller, we use the Adam optimizer with an initial learning rate of 0.001 and apply an exponential scheduler with a decay of 0.99. Modules are trained for 100 epochs in mini-batches of size 32. 10% of the training data is used for validation to choose the best set of models. For DeepROM training, we use $\beta_2 = 1$ in (14). For learning control, we use $\beta_3 = 0.2$ in (16), $\alpha = 0.2$ in (17), and $\mathbf{K} = 0.5\mathbf{I}_{r_x}$ in (18). Since the learned ROMs from one training instance to another can vary, the hyperparameter pair (α, β_3) may require re-tuning accordingly.

For the Deep Koopman model, the input matrix is optimized by gradient descent during training along with the encoder-decoder parameters, whereas the system matrix is obtained using linear least-squares regression. The datasets are divided into staggered 32-step sequences for training, and the model is trained by generating recursive predictions over 32 steps following [13]. We employ the same learning hyperparameters as our model, adjusting the mini-batch size to 8 to accommodate multistep recursive prediction.

REFERENCES

- [1] K. Willcox and J. Peraire, "Balanced model reduction via the proper orthogonal decomposition," *AIAA journal*, vol. 40, no. 11, pp. 2323–2330, 2002.
- [2] J.-N. Juang and R. S. Pappa, "An eigensystem realization algorithm for modal parameter identification and model reduction," *Journal of guidance, control, and dynamics*, vol. 8, no. 5, pp. 620–627, 1985.
- [3] J.-N. Juang, M. Phan, L. G. Horta, and R. W. Longman, "Identification of observer/kalman filter markov parameters-theory and experiments," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 2, pp. 320–329, 1993.
- [4] J. L. Proctor, S. L. Brunton, and J. N. Kutz, "Dynamic mode decomposition with control," *SIAM Journal on Applied Dynamical Systems*, vol. 15, no. 1, pp. 142–161, 2016.
- [5] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," *Journal of fluid mechanics*, vol. 641, pp. 115–127, 2009.
- [6] B. Lusch, J. N. Kutz, and S. L. Brunton, "Deep learning for universal linear embeddings of nonlinear dynamics," *Nature communications*, vol. 9, no. 1, p. 4950, 2018.
- [7] H. Eivazi, H. Veisi, M. H. Naderi, and V. Esfahanian, "Deep neural networks for nonlinear model order reduction of unsteady flows," *Physics of Fluids*, vol. 32, no. 10, p. 105104, 2020.
- [8] P. R. Vlachas, G. Arampatzis, C. Uhler, and P. Koumoutsakos, "Multiscale simulations of complex systems by learning their effective dynamics," *Nature Machine Intelligence*, vol. 4, no. 4, pp. 359–366, 2022.
- [9] C. Vignon, J. Rabault, and R. Vinuesa, "Recent advances in applying deep reinforcement learning for flow control: Perspectives and future directions," *Physics of fluids*, vol. 35, no. 3, 2023.
- [10] Y.-Z. Wang, J.-Z. Peng, N. Aubry, Y.-B. Li, Z.-H. Chen, and W.-T. Wu, "Control policy transfer of deep reinforcement learning based intelligent forced heat convection control," *International Journal of Thermal Sciences*, vol. 195, p. 108618, 2024.
- [11] P. Holl, N. Thuerey, and V. Koltun, "Learning to control pdes with differentiable physics," in *International Conference on Learning Representations*, 2020.
- [12] T. Takahashi, J. Liang, Y.-L. Qiao, and M. C. Lin, "Differentiable fluids with solid coupling for learning and control," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35(7), 2021, pp. 6138–6146.
- [13] J. Morton, A. Jameson, M. J. Kochenderfer, and F. Witherden, "Deep dynamical modeling and control of unsteady fluid flows," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [14] O. Bounou, J. Ponce, and J. Carpentier, "Online learning and control of dynamical systems from sensory input," in *NeurIPS 2021-Thirty-fifth Conference on Neural Information Processing Systems Year*, 2021.
- [15] K. Chen, J. Lin, Y. Qiu, F. Liu, and Y. Song, "Deep learning-aided model predictive control of wind farms for age considering the dynamic wake effect," *Control Engineering Practice*, vol. 116, p. 104925, 2021.
- [16] M. A. Khodkar, P. Hassanzadeh, and A. Antoulas, "A koopman-based framework for forecasting the spatiotemporal evolution of chaotic dynamics with nonlinearities modeled as exogenous forcings," *arXiv preprint arXiv:1909.00076*, 2019.
- [17] P. J. Schmid, "Dynamic mode decomposition of numerical and experimental data," *Journal of fluid mechanics*, vol. 656, pp. 5–28, 2010.
- [18] P. Baldi and K. Hornik, "Neural networks and principal component analysis: Learning from examples without local minima," *Neural networks*, vol. 2, no. 1, pp. 53–58, 1989.
- [19] P. Saha, M. Egerstedt, and S. Mukhopadhyay, "Neural identification for control," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4648–4655, 2021.
- [20] J. Z. Kolter and G. Manek, "Learning stable deep dynamics models," *Advances in neural information processing systems*, vol. 32, 2019.
- [21] D. Kalise and K. Kunisch, "Polynomial approximation of high-dimensional hamilton-jacobi-bellman equations and applications to feedback control of semilinear parabolic pdes," *SIAM Journal on Scientific Computing*, vol. 40, no. 2, pp. A629–A652, 2018.
- [22] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. K. and, "On dynamic mode decomposition: Theory and applications," *Journal of Computational Dynamics*, vol. 1, no. 2, pp. 391–421, 2014.
- [23] A. Albert, *Regression and the Moore-Penrose Pseudoinverse*. Academic Press, 1972.