

Autoencoding Reduced Order Models for Control through the Lens of Dynamic Mode Decomposition

Anonymous authors

Paper under double-blind review

Abstract

Modeling and control of high-dimensional dynamical systems often involve some dimensionality reduction techniques to construct a lower-order model that makes the associated task computationally feasible or less demanding. In recent years, two techniques have become widely popular for analysis and model order reduction of high-dimensional dynamical systems: (1) dynamic mode decomposition and (2) deep autoencoding learning. This paper establishes a connection between dynamic mode decomposition and autoencoding learning for controlled dynamical systems. Specifically, we first show that an optimization objective for learning a linear autoencoding reduced order model can be formulated such that its solution closely resembles the solution obtained by the *dynamic mode decomposition with control* algorithm. The linear autoencoding architecture is then extended to a deep autoencoding architecture to learn a nonlinear reduced order model. Finally, the learned reduced order model is used to design a controller utilizing stability-constrained deep neural networks. The studied framework does not require knowledge of the governing equations of the underlying system and learns the model and controller solely from time series data of observations and actuations. We provide empirical analysis on modeling and control of spatiotemporal high-dimensional systems, including fluid flow control.

1 Introduction

Designing controllers for high-dimensional dynamical systems remains a challenge as many control algorithms become computationally prohibitive in high dimensions. Typically, a *reduce-then-design* approach (Atwell et al. (2001)) is used in practice, which involves two steps: (1) develop a reduced order model (ROM) using dimensionality reduction techniques and (2) design a controller for that ROM (Figure 1a). For controlled dynamical systems, the model order reduction (MOR) approaches either combine analytical techniques with empirical approximation (Willcox & Peraire (2002)) or are purely data-driven (Juang & Pappa (1985); Juang et al. (1993); Proctor et al. (2016)). 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 (Rowley et al. (2009)). Another recent research trend involves utilizing deep neural networks (DNNs), particularly autoencoders, for modeling and control of high-dimensional dynamical systems (Lusch et al. (2018); Erichson et al. (2019); Eivazi et al. (2020); Morton et al. (2018); Bounou et al. (2021); Chen et al. (2021); Ping et al. (2021)).

In this paper, we provide a perspective connecting DMD and autoencoding ROMs for controlled dynamical systems and present a framework to learn control policies for such systems by means of the DNN-based ROMs. We first formulate an objective function for data-driven learning of controlled dynamical systems in a linear autoencoding configuration. We analytically show that the associated objective function encourages a linear ROM that closely resembles the lower-order model obtained using the *DMD with control* (DMDc) algorithm (Proctor et al. (2016)). The linear autoencoding architecture is designed in such a way that its components can be replaced with DNNs and the corresponding objective function can be optimized by gradient descent to obtain a nonlinear ROM. The architecture with DNN components, DeepROM, closely resembles the aforementioned deep autoencoding architectures used in recent literature for the prediction and control of dynamical systems. The learned DNN-based ROM is then used in a control learning framework,




Figure 1: (a) Reduce-then-design paradigm for designing control for high-dimensional systems, (b): our work in the context.

deep reduced order control (DeepROC), to design a controller for the underlying system. The control policy is learned by jointly training two DNNs: one stability-constrained DNN predicts a target closed-loop dynamics for the learned ROM while the other DNN serves as a controller to achieve that target dynamics. We analytically show that keeping the joint learning objective within a sufficiently small value implies stability for the closed-loop ROM in terms of *ultimate boundedness*, i.e., trajectories starting close to the desired state stay close to the desired state. The overall workflow of this paper is shown in Figure 1(b). We provide empirical analysis using examples of spatiotemporal physical processes such as reaction–diffusion and fluid flow systems. In summary, our contributions are as follows:

- For controlled dynamical systems, we show that an objective function can be formulated in a linear autoencoding configuration and optimized by gradient descent such that the corresponding linear ROM closely resembles the ROM obtained using the DMDc algorithm.
- We extend the linear autoencoding configuration to a deep autoencoding configuration to learn a DNN-based nonlinear ROM.
- We analytically show that a DNN controller can be trained such that the closed-loop trajectories of the learned ROM remain ultimately bounded.
- We empirically show the similarity of the linear autoencoding ROM (LAROM) with DMDc and evaluate the prediction performance of DeepROM and control performance of DeepROC in experiments with high-dimensional systems, including suppressing vortex shedding in fluid flow.

2 Related Work

In recent years, deep learning (DL) has seen widespread application in scientific and engineering problems, including understanding complex dynamics of large-scale or high-dimensional systems and solving associated computational tasks. The majority of the research in this area focuses on the modeling and prediction of such complex dynamics using deep neural networks (DNNs) (Xingjian et al. (2015); Long et al. (2018); Raissi (2018); Seo et al. (2019); Ayed et al. (2019); Donà et al. (2020)) and has found application in several fields including fluid flow (Erichson et al. (2019); Eivazi et al. (2020); Srinivasan et al. (2019)), biochemical and electric power systems (Yeung et al. (2019)), climate and ocean systems (Scher (2018); Ren et al. (2021); Yang et al. (2017); De Bézenac et al. (2019)), and structural analysis Zhang et al. (2020), just to name a few.

A second line of research, though relatively less prevalent than modeling and prediction, is utilizing DL for controlling high-dimensional systems, particularly in fluid flow control tasks. Rabault et al. (2019);

Tang et al. (2020) applied deep reinforcement learning (DRL) in active flow control for vortex shedding suppression and used a system-specific reward function involving lift and drag. Ma et al. (2018) used an autoencoder for encoding high-dimensional fluid state to low-dimensional features and trained DRL agents with those features to control rigid bodies in a coupled 2D system involving both fluid and rigid bodies. Garnier et al. (2021) also used an autoencoder for feature extraction to train a DRL agent for controlling the position of a small cylinder in a two-cylinder fluid system to reduce the total drag. Beintema et al. (2020) used DRL with system-specific rewards to control a Rayleigh–Bénard system with the aim of reducing convective effects. Model-free DRL methods have high sample complexity necessitating a large number of interactions with the environment and often require system-specific reward construction. Furthermore, running numerical solvers in every iteration to provide feedback to the agents is computationally expensive. The same concern arises for the methods involving differentiable simulators. For example, Holl et al. (2020) used a differentiable partial differential equation (PDE) solver to generate gradient-based feedback for a predictor-corrector scheme to control PDE-driven high-dimensional systems. Takahashi et al. (2021) too integrated a differentiable simulator with DNNs to learn control in coupled solid-fluid systems.

The alternative to model-free methods takes the traditional approach: develop a model first and then use that to design controllers. Deep learning is now being used in this process by developing frameworks like DeepMPC (Lenz et al. (2015)) which incorporates DNN features in model predictive control (MPC). Bieker et al. (2020) and Morton et al. (2018) utilized the DeepMPC framework for fluid flow control. Bieker et al. (2020) used a recurrent neural network (RNN) to model the dynamics of only control-relevant quantities (i.e. lift and drag) of the system, which is then employed in an MPC framework for the flow control tasks. Morton et al. (2018) followed the method proposed by Takeishi et al. (2017) and used DNN-based embedding to first learn a linear reduced order model (ROM) in Koopman invariant subspace and then incorporate it in the MPC framework. Similar approaches have been applied to other high-dimensional control tasks like control from video input (Bounou et al. (2021)), automatic generation control in wind farms in presence of dynamic wake effect (Chen et al. (2021)), and transient stabilization in power grids (Ping et al. (2021)). However, Khodkar et al. (2019) showed that the linear combination of a finite number of dynamic modes may not accurately represent the long-term nonlinear characteristics of complex dynamics and adding nonlinear forcing terms yields better prediction accuracy (Eivazi et al. (2020)). Furthermore, optimization at every step in the MPC framework is computationally expensive.

3 Problem and Preliminaries

3.1 Problem statement

Consider a time-invariant controlled dynamical system

$$\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \mathbf{u}), \quad (1)$$

where $\mathbf{x}(t) \in \mathbb{X} \subset \mathbb{R}^{d_{\mathbf{x}}}$, $d_{\mathbf{x}} \gg 1$ and $\mathbf{u}(t) \in \mathbb{U} \subset \mathbb{R}^{d_{\mathbf{u}}}$ are the system state and the actuation (or control input), respectively, at time t . Our objective is to learn a feedback controller $\mathbf{u} = \pi(\mathbf{x})$ for this high-dimensional ($d_{\mathbf{x}} \gg 1$) system of (1) to stabilize it at a desired state in a data-driven reduce-then-design approach when the nonlinear function $f : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}^{d_{\mathbf{x}}}$ is unknown. We 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 .

3.2 Stabilization of controlled systems

Suppose the function f in (1) is locally Lipschitz and $(\mathbf{x} = \mathbf{0}, \mathbf{u} = \mathbf{0})$ is an equilibrium pair of the system, i.e., $f(\mathbf{0}, \mathbf{0}) = \mathbf{0}$. The system (1) is said to be (locally) *stabilizable* with respect to the equilibrium pair if there exists a locally Lipschitz function

$$\pi : \mathbb{X}_0 \rightarrow \mathbb{U}, \quad \pi(\mathbf{0}) = \mathbf{0},$$

defined on some neighborhood $\mathbb{X}_0 \subset \mathbb{X}$ of the origin $\mathbf{x} = \mathbf{0}$ for which the closed-loop system

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

is locally (in state space \mathbb{X}_0) asymptotically stable (i.e. $\|\mathbf{x}(t_0)\| < \delta$ implies $\lim_{t \rightarrow \infty} \mathbf{x}(t) = \mathbf{0}$). If $\mathbb{X}_0 = \mathbb{X}$ and (2) is globally asymptotically stable, then the system (1) is said to be *globally* stabilizable (Sontag (2013)).

Stability of the closed-loop system $\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \pi(\mathbf{x})) = h(\mathbf{x})$ (or any uncontrolled/autonomous systems in general) at equilibrium points can be analyzed using the method of Lyapunov. Let $\mathcal{V} : \mathbb{X} \rightarrow \mathbb{R}$ be a continuously differentiable function such that

$$\mathcal{V}(\mathbf{0}) = 0, \quad \text{and} \quad \mathcal{V}(\mathbf{x}) > 0 \quad \forall \mathbf{x} \in \mathbb{X} \setminus \{\mathbf{0}\}, \quad (3)$$

and the time derivative of \mathcal{V} along the trajectories

$$\frac{d\mathcal{V}}{dt} = \nabla \mathcal{V}(\mathbf{x})^\top \frac{d\mathbf{x}}{dt} = \nabla \mathcal{V}(\mathbf{x})^\top h(\mathbf{x}) \leq 0 \quad \forall \mathbf{x} \in \mathbb{X}. \quad (4)$$

Then, the equilibrium point $\mathbf{x} = \mathbf{0}$ is stable (i.e., for each $\epsilon > 0$, there exists a $\delta = \delta(\epsilon) > 0$ such that $\|\mathbf{x}(t_0)\| < \delta$ implies $\|\mathbf{x}(t)\| < \epsilon$, $\forall t > t_0$) and \mathcal{V} is called a Lyapunov function. If $\frac{d\mathcal{V}}{dt} < 0$ in $\mathbb{X} \setminus \{\mathbf{0}\}$, then $\mathbf{x} = \mathbf{0}$ is asymptotically stable. Moreover, if there exist positive constants c_1, c_2, c_3 and c_4 such that

$$c_1 \|\mathbf{x}\|^2 \leq \mathcal{V}(\mathbf{x}) \leq c_2 \|\mathbf{x}\|^2, \quad (5)$$

and

$$\nabla \mathcal{V}(\mathbf{x})^\top h(\mathbf{x}) \leq -c_3 \|\mathbf{x}\|^2, \quad \forall \mathbf{x} \in \mathbb{X}, \quad (6)$$

then $\mathbf{x} = \mathbf{0}$ is *exponentially* stable (i.e., there exist positive constants δ, λ and γ such that $\|\mathbf{x}(t)\| \leq \lambda \|\mathbf{x}(t_0)\| e^{-\gamma(t-t_0)}$, $\forall \|\mathbf{x}(t_0)\| < \delta$) (Khalil (2002)). For global stability, we need another condition

$$\mathcal{V}(\mathbf{x}) \rightarrow \infty \text{ as } \|\mathbf{x}\| \rightarrow \infty. \quad (7)$$

The converse statements regarding the existence of a Lyapunov function with aforementioned properties for (locally/globally, asymptotically/exponentially) stable systems hold true as well. The notation $(\cdot)^\top$ in the above discussion denotes the transpose operation.

In this paper, we assume that the system we are aiming to stabilize (at an equilibrium) is stabilizable in the sense of the aforementioned definition and criteria, i.e., there exists a continuously differentiable function \mathcal{V} and a Lipschitz continuous control law π such that criteria (3) and (4) (and possibly (5) and (7) as well) are conformed.

Though the above formulation is for stabilization at an equilibrium point $\mathbf{x} = \mathbf{0}$, the same can be used to stabilize the system at any arbitrary point \mathbf{x}_{ss} . In that case, a steady-state control input \mathbf{u}_{ss} is required that can maintain the equilibrium at \mathbf{x}_{ss} , i.e., $f(\mathbf{x}_{ss}, \mathbf{u}_{ss}) = \mathbf{0}$. The change of variables $\mathbf{x}_e = \mathbf{x} - \mathbf{x}_{ss}$, $\mathbf{u}_e = \mathbf{u} - \mathbf{u}_{ss}$ leads to a transformed system where we can apply the aforementioned formulation of stabilization. The overall control, in this case, $\mathbf{u} = \mathbf{u}_e + \mathbf{u}_{ss}$ comprises a feedback component \mathbf{u}_e and a feedforward component \mathbf{u}_{ss} (Khalil (2002)).

3.3 Dynamic mode decomposition with control

DMD (Schmid (2010)) is a data-driven method that reconstructs the underlying dynamics using only a time series of snapshots from the system. DMD computes a modal decomposition where each mode is associated with an oscillation frequency and decay/growth rate. DMD has become a widely used technique for spectral analysis of dynamical systems. DMDc (Proctor et al. (2016)) 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:

$$\hat{\mathbf{x}}(t_{i+1}) = \mathbf{A}\mathbf{x}(t_i) + \mathbf{B}\mathbf{u}(t_i), \quad i = 0, 1, \dots, n-1, \quad (8)$$

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 (8) could be computationally prohibitive for $d_x \gg 1$. DMDc leverages dimensionality reduction to compute a ROM

$$\mathbf{x}_{R,DMDc}(t_i) = \mathbf{E}_{DMDc} \mathbf{x}(t_i), \quad (9a)$$

$$\mathbf{x}_{R,DMDc}(t_{i+1}) = \mathbf{A}_{R,DMDc} \mathbf{x}_{R,DMDc}(t_i) + \mathbf{B}_{R,DMDc} \mathbf{u}(t_i), \quad i = 0, 1, \dots, n-1, \quad (9b)$$

which retains the dominant dynamic modes of (8). Here, $\mathbf{x}_{R,DMDc}(t_i) \in \mathbb{R}^{r_x}$ is the reduced state, where $r_x \ll d_x$, and $\mathbf{E}_{DMDc} \in \mathbb{R}^{r_x \times d_x}$, $\mathbf{A}_{R,DMDc} \in \mathbb{R}^{r_x \times r_x}$, $\mathbf{B}_{R,DMDc} \in \mathbb{R}^{r_x \times d_u}$. The full state is reconstructed from the reduced state using the transformation $\hat{\mathbf{x}}(t_i) = \mathbf{D}_{DMDc} \mathbf{x}_{R,DMDc}(t_i)$, where $\mathbf{D}_{DMDc} \in \mathbb{R}^{d_x \times r_x}$. 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:

$$\mathbf{Y} = \hat{\mathbf{U}}_Y \hat{\boldsymbol{\Sigma}}_Y \hat{\mathbf{V}}_Y^\top, \quad \mathbf{\Omega} = \hat{\mathbf{U}}_\Omega \hat{\boldsymbol{\Sigma}}_\Omega \hat{\mathbf{V}}_\Omega^\top, \quad (10)$$

where $\hat{\mathbf{U}}_Y \in \mathbb{R}^{d_x \times r_x}$, $\hat{\boldsymbol{\Sigma}}_Y \in \mathbb{R}^{r_x \times r_x}$, $\hat{\mathbf{V}}_Y \in \mathbb{R}^{n \times r_x}$, $\hat{\mathbf{U}}_\Omega \in \mathbb{R}^{(d_x+d_u) \times r_{xu}}$, $\hat{\boldsymbol{\Sigma}}_\Omega \in \mathbb{R}^{r_{xu} \times r_{xu}}$, and $\hat{\mathbf{V}}_\Omega \in \mathbb{R}^{n \times r_{xu}}$. $r_x < \min(d_x, n)$ and $r_{xu} < \min(d_x + d_u, n)$ denote the truncation dimensions of SVDs. Utilizing the SVDs of (10) the parameters of the ROM (9) is obtained as

$$\mathbf{E}_{DMDc} = \hat{\mathbf{U}}_Y^\top, \quad \mathbf{D}_{DMDc} = \hat{\mathbf{U}}_Y, \quad (11a)$$

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

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

4 Method

As mentioned earlier, in the reduce-then-design approach, we first need to develop a ROM and then design a controller using that ROM. A controller designed for the ROM is expected to perform well in the full system only if the ROM effectively captures the dynamic characteristics of the underlying system. In this section, we first describe how to design a ROM that effectively captures the relation between successive observations and actuation. Next, we delineate the process for learning controllers utilizing the learned ROM.

4.1 Learning a reduced order model

DMDc can extract the dominant modes of underlying dynamics in a ROM Proctor et al. (2016). In order 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 MOR problem using linear networks (single layer without nonlinear activation). Consider the following MOR problem

$$\mathbf{x}_R(t_i) = \mathbf{E}_x \mathbf{x}(t_i), \quad \mathbf{x}_R(t_{i+1}) = \mathbf{A}_R \mathbf{x}_R(t_i) + \mathbf{B}_R \mathbf{u}(t_i), \quad \hat{\mathbf{x}}(t_i) = \mathbf{D}_x \mathbf{x}_R(t_i), \quad i = 0, 1, \dots, n-1, \quad (12)$$

where the linear operators $\mathbf{E}_x \in \mathbb{R}^{r_x \times d_x}$ and $\mathbf{D}_x \in \mathbb{R}^{d_x \times r_x}$ projects and reconstructs back, respectively, the high-dimensional system state to and from a low-dimensional feature $\mathbf{x}_R \in \mathbb{R}^{r_x}$. The linear operators $\mathbf{A}_R \in \mathbb{R}^{r_x \times r_x}$ and $\mathbf{B}_R \in \mathbb{R}^{r_x \times d_u}$ describe the relations between successive reduced states and actuations.

The DMDc algorithm essentially solves for $\tilde{\mathbf{G}} \in \mathbb{R}^{r_x \times (d_x+d_u)}$ to minimize $\frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x \mathbf{x}(t_{i+1}) - \tilde{\mathbf{G}} \boldsymbol{\omega}(t_i)\|^2$ for a fixed projection matrix $\mathbf{E}_x = \mathbf{E}_{DMDc} = \hat{\mathbf{U}}_Y^\top$. Here, $\boldsymbol{\omega}(t_i)$ is the concatenated vector of state and actuation as defined in section 3.3. 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 $\mathbf{D}_{DMDc} = \hat{\mathbf{U}}_Y$ to get the ROM components $\mathbf{A}_{R,DMDc}$ and $\mathbf{B}_{R,DMDc}$. Details of this process along with the proofs are given in appendix A.5. Note, the final step of this process offers dimensionality reduction 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 (12) and also offers dimensionality reduction when nonlinear components are used.

Theorem 4.1.1. *Consider the following objective function*

$$L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) = \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x \mathbf{x}(t_{i+1}) - \mathbf{G} \mathbf{E}_{xu} \boldsymbol{\omega}(t_i)\|^2, \quad (13)$$

where $\mathbf{G} = [\mathbf{A}_R \ \mathbf{B}_R] \in \mathbb{R}^{r_x \times (r_x + d_u)}$, $\mathbf{E}_{xu} = \begin{bmatrix} \mathbf{E}_x & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix} \in \mathbb{R}^{(r_x + d_u) \times (d_x + d_u)}$, \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

$$\mathbf{G} \mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top, \quad (14)$$

where \mathbf{Y} and $\boldsymbol{\Omega}$ are the data matrices as defined in section (3.3). 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

$$\mathbf{G} = [\mathbf{A}_R \ \mathbf{B}_R] = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top (\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top)^{-1}. \quad (15)$$

Proof sketch. This can be proved by a method similar to the one used for deriving the solution of linear autoencoder in (Baldi & Hornik (1989)). For any fixed \mathbf{E}_x , the objective function of (13) can be written as $L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) = \|\text{vec}(\mathbf{E}_x \mathbf{Y}) - (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{G})\|^2$, where \otimes denotes the Kronecker product and $\text{vec}(\cdot)$ denotes vectorization of a matrix. Optimizing this linear least-square problem, we get (14) and (15), given the stated conditions are satisfied. The complete proof is given in appendix A.1.

Remark. 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. Proof of this is given in appendix A.4.

The connection between the ROM obtained by minimizing L_{pred} (for a fixed \mathbf{E}_x), i.e., (15) and the DMDC ROM of (11b) is not readily apparent. To interpret the connection, we formulate an alternative representation of (15) utilizing the SVD and the Moore-Penrose inverse of matrices. This alternative representation leads to the following result.

Corollary 4.1.1.1. *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 $\mathbf{G} = [\mathbf{A}_R \ \mathbf{B}_R]$ corresponding to the unique minimum of L_{pred} can be expressed as*

$$\mathbf{A}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}_\Omega^* \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y, \quad \text{and} \quad \mathbf{B}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}_\Omega^* \mathbf{U}_{\Omega,2}^\top, \quad (16)$$

where $[\mathbf{U}_{\Omega,1}^\top \ \mathbf{U}_{\Omega,2}^\top] = \mathbf{U}_\Omega^\top$ 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 $\boldsymbol{\Sigma}_\Omega^* = \lim_{\varepsilon \rightarrow 0} (\boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \boldsymbol{\Sigma}_\Omega + \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \boldsymbol{\Sigma}_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \boldsymbol{\Sigma}_\Omega^\top$.

Proof sketch. This can be derived by plugging $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$ into (15), and using the SVD definition and the limit definition (Albert (1972)) of the Moore-Penrose inverse. The complete proof is given in appendix A.3 that uses some preliminary results presented in appendix A.2.

Remark. It can be verified easily that if we use the truncated SVD (as defined by 10), instead of the full SVD, for $\boldsymbol{\Omega}$ in corollary 4.1.1.1, we get an approximation of (16):

$$\hat{\mathbf{A}}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \hat{\mathbf{V}}_\Omega \hat{\boldsymbol{\Sigma}}_\Omega^* \hat{\mathbf{U}}_{\Omega,1}^\top \hat{\mathbf{U}}_Y, \quad \text{and} \quad \hat{\mathbf{B}}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \hat{\mathbf{V}}_\Omega \hat{\boldsymbol{\Sigma}}_\Omega^* \hat{\mathbf{U}}_{\Omega,2}^\top, \quad (17)$$

where $\hat{\boldsymbol{\Sigma}}_\Omega^* = \lim_{\varepsilon \rightarrow 0} (\hat{\boldsymbol{\Sigma}}_\Omega^\top \hat{\mathbf{U}}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \hat{\mathbf{U}}_{\Omega,1} \hat{\boldsymbol{\Sigma}}_\Omega + \hat{\boldsymbol{\Sigma}}_\Omega^\top \hat{\mathbf{U}}_{\Omega,2}^\top \hat{\mathbf{U}}_{\Omega,2} \hat{\boldsymbol{\Sigma}}_\Omega + \varepsilon^2 \mathbf{I}_{r_{xu}})^{-1} \hat{\boldsymbol{\Sigma}}_\Omega^\top$. We can see that (17) has the same form as (11b), except $\hat{\boldsymbol{\Sigma}}_\Omega^{-1}$ is replaced with $\hat{\boldsymbol{\Sigma}}_\Omega^*$.




Figure 2: Autoencoding architecture for model order reduction. The state encoder \mathcal{E}_x and control encoder \mathcal{E}_u reduce the dimension of the state and actuation, respectively. 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 $\mathcal{L}_{\text{recon}}$. The dashed arrow indicates that the \mathcal{E}_u is used only when $d_u \gg 1$; otherwise, the actuation is used as a direct input to ROM.

All the aforementioned results are derived for a fixed \mathbf{E}_x and the relation to the DMDc is specific to the case $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$. To encourage such a choice of \mathbf{E}_x , we include the standard autoencoder reconstruction loss, given by

$$L_{\text{recon}}(\mathbf{E}_x, \mathbf{D}_x) = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}(t_i) - \mathbf{D}_x \mathbf{E}_x \mathbf{x}(t_i)\|^2, \quad (18)$$

to the optimization objective.

Baldi & Hornik (1989) showed 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 $\mathbf{U}_{r_x} = \hat{\mathbf{U}}_Y$. Accordingly, we propose to minimize the following objective function to encourage a DMDc-like solution for LAROM:

$$L(\mathbf{E}_x, \mathbf{D}_x, \mathbf{G}) = L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) + \beta_1 L_{\text{recon}}(\mathbf{E}_x, \mathbf{D}_x), \quad (19)$$

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

Now, we can replace all the trainable components, i.e., \mathbf{E}_x , \mathbf{D}_x , and \mathbf{G} , with DNNs to establish a nonlinear model order reduction framework utilizing gradient descent. 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.,

$$\mathbf{x}_R = \mathcal{E}_x(\mathbf{x}), \quad \hat{\mathbf{x}} = \mathcal{D}_x(\mathbf{x}_R), \quad (20)$$

where $\mathbf{x}_R \in \mathbb{R}^{r_x}$ denotes the reduced state, and $\hat{\mathbf{x}}$ is the reconstruction of \mathbf{x} . Unlike the linear case, we use an encoder $\mathcal{E}_u : \mathbb{U} \rightarrow \mathbb{R}^{r_u}$, $r_u \ll d_u$ for the actuation as well, in cases where the control space is also

high-dimensional (for example, distributed control of spatiotemporal PDEs). The control encoder \mathcal{E}_u maps the high-dimensional actuation to a low-dimensional representation: $\mathbf{u}_R = \mathcal{E}_u(\mathbf{u})$, where $\mathbf{u}_R \in \mathbb{R}^{r_u}$ denotes the encoded actuation. The encoded state and control are then fed to another DNN that represents the reduced order dynamics

$$\frac{d\mathbf{x}_R}{dt} = \mathcal{F}(\mathbf{x}_R, \mathbf{u}_R), \quad (21)$$

where $\mathcal{F} : \mathbb{R}^{r_x} \times \mathbb{R}^{r_u} \rightarrow \mathbb{R}^{r_x}$. Note, here the ROM is represented as a continuous-time dynamics, unlike the linear case where we used a discrete-time model. The reason behind this choice will become apparent in the following subsection.

Given the current reduced state $\mathbf{x}_R(t_i)$ and control input $\mathbf{u}_R(t_i)$, the next reduced state $\mathbf{x}_R(t_{i+1})$ can be computed by integrating \mathcal{F} using standard numerical integrator or neural ODE (Chen et al. (2018)):

$$\mathbf{x}_R(t_{i+1}) = \mathbf{x}_R(t_i) + \int_{t_i}^{t_{i+1}} \mathcal{F}(\mathbf{x}_R(t_i), \mathbf{u}_R(t_i)) dt \triangleq \mathcal{G}(\mathbf{x}_R(t_i), \mathbf{u}_R(t_i)). \quad (22)$$

We can say that \mathcal{G} is the nonlinear counterpart of \mathbf{G} . Figure 2 shows the overall framework for training DeepROM.

We train $\mathcal{E}_x, \mathcal{E}_u, \mathcal{D}_x$, and \mathcal{F} by minimizing the following loss function, analogous to (19),

$$\mathcal{L}(\mathcal{E}_x, \mathcal{E}_u, \mathcal{D}_x, \mathcal{F}) = \mathcal{L}_{\text{pred}}(\mathcal{E}_x, \mathcal{E}_u, \mathcal{F}) + \beta_2 \mathcal{L}_{\text{recon}}(\mathcal{E}_x, \mathcal{D}_x), \quad (23)$$

where $\beta_2 > 0$ is a tunable hyperparameter and $\mathcal{L}_{\text{pred}}, \mathcal{L}_{\text{recon}}$ are defined as follows,

$$\begin{aligned} \mathcal{L}_{\text{pred}}(\mathcal{E}_x, \mathcal{E}_u, \mathcal{F}) &= \frac{1}{n} \sum_{i=0}^{n-1} \left\| \mathcal{E}_x(\mathbf{x}(t_{i+1})) - \mathcal{G}(\mathcal{E}_x(\mathbf{x}(t_i)), \mathcal{E}_u(\mathbf{u}(t_i))) \right\|^2, \\ \mathcal{L}_{\text{recon}}(\mathcal{E}_x, \mathcal{D}_x) &= \frac{1}{n} \sum_{i=1}^n \left\| \mathbf{x}(t_i) - \mathcal{D}_x \circ \mathcal{E}_x(\mathbf{x}(t_i)) \right\|^2. \end{aligned} \quad (24)$$

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.

4.2 Learning control

Once we get a trained ROM of the form (21) using the method proposed in section 4.1, 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 by Saha et al. (2021) for learning control law for nonlinear systems, represented by DNNs. The core idea of the method is to hypothesize a target dynamics that is exponentially stable at the desired state and simultaneously learn a control policy to realize that target dynamics in 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 :

$$\mathbf{u} = \Pi(\mathbf{x}_R). \quad (25)$$

This control \mathbf{u} is then encoded by (trained) \mathcal{E}_u to its low-dimensional representation \mathbf{u}_R . Finally, the reduced state \mathbf{x}_R and actuation \mathbf{u}_R are fed to the (trained) ROM of (21) to get $\mathcal{F}(\mathbf{x}_R, \mathbf{u}_R)$. The overall framework for learning control is shown in Figure 3.

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

$$\mathcal{L}_{\text{ctrl}}(\mathcal{F}_s, \Pi) = \frac{1}{n} \sum_{i=1}^n \left\| \mathcal{F}(\mathcal{E}_x(\mathbf{x}(t_i)), \mathcal{E}_u \circ \Pi \circ \mathcal{E}_x(\mathbf{x}(t_i))) - \mathcal{F}_s \circ \mathcal{E}_x(\mathbf{x}(t_i)) \right\|^2. \quad (26)$$




Figure 3: 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.

To minimize the control effort, we add a regularization loss with (26), and the overall training objective for learning control is given by

$$\mathcal{L}_{\text{ctrl,reg}}(\mathcal{F}_s, \Pi) = \mathcal{L}_{\text{ctrl}}(\mathcal{F}_s, \Pi) + \beta_3 \frac{1}{n} \sum_{i=1}^n \|\Pi(\mathbf{x}_R(t_i))\|^2, \quad (27)$$

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 , \mathcal{E}_u , 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:

$$\mathbf{u} = \Pi \circ \mathcal{E}_x(\mathbf{x}) = \pi(\mathbf{x}). \quad (28)$$

As we mentioned earlier, we require the target dynamics, hypothesized by a DNN, to be exponentially stable at the desired state. Without loss of generality, we consider stability at $\mathbf{x}_R = \mathbf{0}$. As we mentioned earlier, the system can be stabilized at any desired state by adding a feedforward component to the control. Dynamics represented by a standard neural network is not stable at any equilibrium point, in general. Kolter & Manek (2019) showed that it is possible to design a DNN, by means of Lyapunov functions, to represent a dynamics that is exponentially stable at an equilibrium point. Accordingly, we represent our target dynamics as follows:

$$\frac{d\mathbf{x}_R}{dt} = \mathcal{F}_s(\mathbf{x}_R) = \mathcal{P}(\mathbf{x}_R) - \frac{\text{ReLU}(\nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R) + \alpha \mathcal{V}_R(\mathbf{x}_R))}{\|\nabla \mathcal{V}_R(\mathbf{x}_R)\|^2} \nabla \mathcal{V}_R(\mathbf{x}_R), \quad (29)$$

where α is a positive constant, $\text{ReLU}(z) = \max(0, z)$, $z \in \mathbb{R}$, and $\mathcal{V}_R : \mathbb{R}^{r_x} \rightarrow \mathbb{R}$ is a candidate Lyapunov function, i.e., satisfies the criteria similar to (3) and (5). We use

$$\mathcal{V}_R(\mathbf{x}_R) = \mathbf{x}_R^\top \mathbf{K} \mathbf{x}_R, \quad (30)$$

where $\mathbf{K} \in \mathbb{R}^{r_x \times r_x}$ is a positive definite matrix.

Though the efficacy of learning control by minimizing the difference with respect to a target dynamics is experimentally demonstrated by Saha et al. (2021), the stability of the closed-loop system subjected to

the learned control law has not been studied analytically. Here, we present a result that shows that if we can minimize $\mathcal{L}_{\text{ctrl}}$ such that the difference between the target dynamics and the closed-loop dynamics is sufficiently small for all $\mathbf{x}_R \in \mathbb{X}_R \subset \mathbb{R}^{r_x}$, then the trajectories of the closed-loop ROM starting sufficiently close to the origin remains close to the origin, i.e., *ultimately bounded* (Khalil (2002)).

Theorem 4.2.1. *Consider the target dynamics defined by (29) and the candidate Lyapunov function defined by (30). Suppose the difference between the target dynamics and the closed-loop dynamics satisfies*

$$\|\mathcal{F}(\mathbf{x}_R, \mathcal{E}_u \circ \Pi(\mathbf{x}_R)) - \mathcal{F}_s(\mathbf{x}_R)\| \leq \delta < \frac{\alpha\theta\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})} \sqrt{\frac{\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}} \eta, \quad (31)$$

for all $\mathbf{x}_R \in \mathbb{X}_R = \{\mathbf{x}_R \in \mathbb{R}^{r_x} \mid \|\mathbf{x}_R\| < \eta\}$ and $0 < \theta < 1$. Then, for all initial points satisfying $\|\mathbf{x}_R(t_0)\| < \sqrt{\frac{\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}} \eta$, the solution of the closed-loop ROM $\frac{d\mathbf{x}_R}{dt} = \mathcal{F}(\mathbf{x}_R, \mathcal{E}_u \circ \Pi(\mathbf{x}_R))$ satisfies

$$\|\mathbf{x}_R(t)\| \leq \lambda e^{-\gamma(t-t_0)} \|\mathbf{x}_R(t_0)\|, \quad \forall t_0 \leq t < t_c + t_0 \quad (32)$$

and

$$\|\mathbf{x}_R(t)\| \leq \frac{2\delta}{\alpha\theta} \lambda^3, \quad \forall t \geq t_c + t_0 \quad (33)$$

for some finite $t_c > 0$, where

$$\gamma = \frac{\alpha(1-\theta)\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})} \quad \text{and} \quad \lambda = \sqrt{\frac{\lambda_{\max}(\mathbf{K})}{\lambda_{\min}(\mathbf{K})}} \quad (34)$$

Proof Sketch. This can be proved by first deriving the Lyapunov conditions for the target dynamics (29) (Theorem 1, Kolter & Manek (2019)) and then applying the stability analysis of perturbed systems (Lemma 9.2, Khalil (2002)) and ultimate boundedness (Theorem 4.18, Khalil (2002)) on the closed-loop ROM. A unified proof is provided in appendix A.6.

5 Empirical Results

For empirical analysis, we consider modeling and controlling spatiotemporal PDE-driven systems with high-dimensional measurements over discretized space. The first example investigates a single variable actuation, whereas distributed actuation is considered for the second example. The similarity between DMDc and LAROM is demonstrated using the dynamic modes estimated in respective methods. The prediction performance of DeepROM is compared against the baseline DMDc. We compare the control performance of DeepROC with a baseline reduced order controller obtained by applying *linear quadratic regulator* (LQR) on the ROM derived from DMDc.

5.1 Reaction–diffusion system stabilization

For the first experiment, we consider the Newell–Whitehead–Segel reaction-diffusion equation with the Neumann boundary condition

$$\begin{aligned} \frac{\partial q}{\partial t} &= \sigma \nabla^2 q + q(1 - q^2) + \mathbf{1}_{\mathbb{W}} w \quad \text{in } \mathbb{I} \times \mathbb{R}^+, \\ \nabla q(\zeta_l, t) &= \nabla q(\zeta_r, t) = 0, \quad t \in \mathbb{R}^+, \end{aligned} \quad (35)$$

which is used to describe various nonlinear physical systems including Rayleigh–Bénard convection. Here, $q(\zeta, t) \in \mathbb{R}$ 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} . (35) is a bistable system with ± 1 as stable and 0 as unstable equilibria. For the control task, we consider feedback stabilization of (35) at the unstable equilibrium 0, as studied by Kalise & Kunisch (2018). 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 B.

5.1.1 Similarity with DMDc

To investigate the similarity DMDc, we first train the LAROM to minimize the objective (19) using gradient descent. As mentioned earlier, L_{recon} , defined by (18), is minimized for any invertible matrix \mathbf{C} , $\mathbf{D}_x = \mathbf{U}_{r_x} \mathbf{C}$, and $\mathbf{E}_x = \mathbf{C}^{-1} \mathbf{U}_{r_x}^\top$. 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} . We need additional constraints to encourage similarity to DMDc. For this, 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 (19).

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 Schmid (2010), not the exact modes obtained in Proctor et al. (2016). 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, Tu et al. (2014)). Figure 4 compares the dynamic modes obtained using DMDc and LAROM for the case when the dimension of the ROMs is 3. Dynamic modes of both methods are similar except for the different signs of the first two modes.




Figure 4: The first three dynamic modes of the reaction–diffusion system, obtained using DMDc and LAROM.

5.1.2 Prediction performance of DeepROM

We now compare the performance of DeepROM, trained with the objective (23), and DMDc in the prediction task. Note, this example uses low-dimensional actuation (just a single variable). Accordingly, the control encoder \mathcal{E}_u is not used here. Figure 5 shows the quantitative and qualitative comparison of the recursive multi-step predictions obtained using DMDc and DeepROM. The prediction error is computed as *normalized mean squared error* (NMSE) with respect to the solution obtained using the PDE solver. For DeepROM, the prediction error plot shows the mean error and 95% confidence interval from 100 test sequences and 3 training instances, whereas the color maps are shown for one example sequence with one training instance. Prediction error increases more quickly for DMDc than DeepROM as DMDc cannot capture the nonlinearity present in the system.

5.1.3 Control performance of DeepROC

Figure 6 shows the control performance of DeepROC and the baseline 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 metrics for comparison:

- (i) mean squared error over time between the controlled solutions and the desired profile
- (ii) differential magnitude that measures the differential changes between the profiles at consecutive time steps. In the steady state, the differential magnitude should be close to zero.
- (iii) the amount of actuation applied




Figure 5: Prediction performance of DMDc and DeepROM in the reaction–diffusion example. For DeepROM, the prediction error plot shows the mean error and 95% confidence interval from 100 test sequences and 3 training instances. One example sequence is used to visually compare the predictions with the solution from a PDE solver.




Figure 6: Control performance of DMDc+LQR and DeepROC in the reaction–diffusion example. For DeepROC, the plots show the mean values with 1-standard deviation interval from 3 training instances.

For DeepROC, the plots show the mean values with 1-standard deviation interval from 3 training instances. Both methods show similar closed-loop error profiles. However, DeepROC requires significantly less amount of actuation compared with DMDc+LQR to reach a similar steady-state error. DeepROC can account for the decaying nonlinear term $-q^3$ present in the system (35) and therefore learns to apply less actuation. Figure 7 visually compares the uncontrolled solution and the controlled solutions obtained using both methods. When uncontrolled, the system reaches the stable equilibrium at 1, whereas the feedback-controlled system is stabilized at the desired state 0 in both cases.

5.2 Vortex shedding suppression in fluid

In this experiment, we consider modeling and suppressing vortex shedding in two-dimensional incompressible flow past a circular cylinder. This is a well-known problem (Schäfer et al. (1996)) and is of great importance




Figure 7: Visual comparison of the uncontrolled solution and the controlled solutions using DeepROC and DMDc+LQR.

for many engineering applications (Williamson (1996)). The dynamics is governed by the incompressible Navier-Stokes equations given by

$$\frac{\partial \mathbf{v}}{\partial t} - \nu \nabla^2 \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{1}{\rho} \nabla p + \mathbf{1}_{\mathbb{W}} \mathbf{w}, \quad \nabla \cdot \mathbf{v} = \mathbf{0} \quad \text{in } \mathbb{I} \times \mathbb{R}^+, \quad (36)$$

where $\mathbf{v}(\boldsymbol{\zeta}, t) \in \mathbb{R}^2$ denotes the flow velocity at location $\boldsymbol{\zeta} \in \mathbb{I} \subset \mathbb{R}^2$ and time t , $p(\boldsymbol{\zeta}, t) \in \mathbb{R}$ denotes the pressure, ν denotes the kinematic viscosity and ρ denotes the density of the fluid. $\mathbf{w}(\boldsymbol{\zeta}, t)$ is the actuation/force applied to the system and $\mathbf{1}_{\mathbb{W}}(\boldsymbol{\zeta})$ is the indicator function with $\mathbb{W} \subset \mathbb{I}$. We use $\mathbb{I} = (0, 2.2) \times (0, 0.41)$ and $\mathbb{W} = (0.11, 0.77) \times (0, 0.41)$. Density and kinematic viscosity are chosen such that the Reynolds number is $Re = 50$, which is just above the cutoff for the onset of the vortex shedding (Williamson (1996)). In this case, vortices are created at the back of the cylinder and are shed periodically from the upper and lower surfaces of the cylinder forming a von Kármán vortex street (Morton et al. (2018)). We use the domain \mathbb{W} for observation and distributed actuation. The Stokes flow is used as the desired state for the control task. More details on the problem setup, dataset generation, neural network architectures, and training settings are given in appendix C.

5.2.1 Similarity with DMDc

To analyze the dynamic modes, we train the LAROM by enforcing $\mathbf{D}_{\mathbf{x}} = \mathbf{E}_{\mathbf{x}}^\top$ and adding the semi-orthogonality constraint to the learning objective, as mentioned in subsection 5.1.1. Figure 8 compares the first two dynamic modes obtained using DMDc and LAROM. Only the streamwise components are shown for brevity. Also, complex modes occur in conjugate pairs and only one from each pair is shown. Dynamic modes identified by LAROM are similar to the ones obtained from DMDc, except the real and imaginary components of the first mode are swapped.

5.2.2 Prediction performance of DeepROM

Figure 9 shows the quantitative and qualitative comparison of the recursive multi-step predictions, starting from $t = 0.1$, obtained using DMDc and DeepROM. The initial state is chosen at $t = 0.1$ because the fluid does not reach the observation region \mathbb{W} before that time. The prediction error is computed as the *mean squared error* (MSE) with respect to the solution obtained using a PDE solver. For DeepROM, the prediction error plot shows the mean error and 1-standard deviation interval from 3 training instances. DeepROM shows lower prediction error in comparison with DMDc. Moreover, unlike DeepROM, DMDc is unable to capture the shedding pattern in multi-step prediction as shown in the contour plots of the velocity magnitude.

5.2.3 Control performance of DeepROC

Figure 10 shows the control performance of DeepROC and the baseline DMDc+LQR in the task of suppressing vortex shedding. We use the same metrics as the previous example for comparison except for actuation.




Figure 8: The first two dynamic modes obtained using DMDc and LAROM for the flow past a cylinder system.




Figure 9: Prediction performance of DMDc and DeepROM in the fluid flow example. For DeepROM, the prediction error plot shows the mean error and 1-standard deviation interval from 3 training instances. Predictions at time step 2500 for the test sequence are visually compared with the solution from a PDE solver. v_m denotes the velocity magnitude.

Since distributed control is applied in this case, we use the magnitude of the actuation here. For DeepROC, the plots show the mean values with 1-standard deviation interval from 3 training instances. To reach a similar steady-state error, DeepROC takes a longer time but uses less amount of actuation. Figure 11 shows the velocity magnitude of the controlled flow for DeepROC and DMDc+LQR at different times, starting from a von Kármán vortex street pattern. Both methods accomplish a similar steady-state flow pattern where vortex shedding has been suppressed.

6 Conclusion

We presented a framework for autoencoder-based modeling and control learning for high-dimensional dynamical systems. We showed that autoencoding ROMs are capable of capturing the dominant modes that are essential in analyzing and designing control for the underlying systems. With gradient descent-based




Figure 10: Control performance of DMDc+LQR and DeepROC in the vortex shedding suppression task. For DeepROC, the plots show the mean values with 1-standard deviation interval from 3 training instances.




Figure 11: Visual comparison of the velocity magnitude of the flow over time subjected to the controllers obtained using DeepROC and DMDc+LQR.

optimization, LAROM can be used to approximate modal decomposition for high-dimensional systems where computing SVD is expensive. As we showed in experiments, DeepROM offers better prediction accuracy than a linear ROM over a relatively longer prediction horizon when applied to nonlinear systems. However, this advantage does not always translate to significant improvement in control performance. Though the used control learning method theoretically ensures ultimate boundedness for the closed-loop ROM solution, data-driven optimization of the learning objective often makes the models susceptible to distribution shift which can impact the control performance. The control learning process in the DeepROC framework can easily be replaced with other methods like model-based RL or MPC. It would be interesting for future work to investigate whether updating both the reduced model and the controller in the MPC framework ensures robustness under distribution shift and offers better control performance. Designing controllers for DNN-based models is a challenging task due to the standard difficulties associated with non-convex optimization. Nevertheless, we envision great prospects in solving many problems of control design for high-dimensional systems utilizing autoencoder-based models as they continue to demonstrate their effectiveness in the analysis and prediction of such systems.

References

- Arthur Albert. *Regression and the Moore-Penrose Pseudoinverse*. Academic Press, 1972.
- Jeanne A Atwell, Jeffrey T Borggaard, and Belinda B King. Reduced order controllers for burgers’ equation with a nonlinear observer. *International Journal of Applied Mathematics and Computer Science*, 11(6): 1311–1330, 2001.
- Ibrahim Ayed, Emmanuel de Bézenac, Arthur Pajot, Julien Brajard, and Patrick Gallinari. Learning dynamical systems from partial observations. *arXiv preprint arXiv:1902.11136*, 2019.
- Pierre Baldi and Kurt Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58, 1989.
- Gerben Beintema, Alessandro Corbetta, Luca Biferale, and Federico Toschi. Controlling rayleigh–bénard convection via reinforcement learning. *Journal of Turbulence*, 21(9-10):585–605, 2020.
- Katharina Bieker, Sebastian Peitz, Steven L Brunton, J Nathan Kutz, and Michael Dellnitz. Deep model predictive flow control with limited sensor data and online learning. *Theoretical and computational fluid dynamics*, 34:577–591, 2020.
- Oumayma Bounou, Jean Ponce, and Justin Carpentier. Online learning and control of dynamical systems from sensory input. In *NeurIPS 2021-Thirty-fifth Conference on Neural Information Processing Systems Year*, 2021.
- Kaixuan Chen, Jin Lin, Yiwei Qiu, Feng Liu, and Yonghua Song. Deep learning-aided model predictive control of wind farms for agc considering the dynamic wake effect. *Control Engineering Practice*, 116: 104925, 2021.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Emmanuel De Bézenac, Arthur Pajot, and Patrick Gallinari. Deep learning for physical processes: Incorporating prior scientific knowledge. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12): 124009, 2019.
- Jérémy Donà, Jean-Yves Franceschi, Sylvain Lamprier, and Patrick Gallinari. Pde-driven spatiotemporal disentanglement. *arXiv preprint arXiv:2008.01352*, 2020.
- Hamidreza Eivazi, Hadi Veisi, Mohammad Hossein Naderi, and Vahid Esfahanian. Deep neural networks for nonlinear model order reduction of unsteady flows. *Physics of Fluids*, 32(10):105104, 2020.
- N Benjamin Erichson, Michael Muehlebach, and Michael W Mahoney. Physics-informed autoencoders for lyapunov-stable fluid flow prediction. *arXiv preprint arXiv:1905.10866*, 2019.
- Paul Garnier, Jonathan Viquerat, Jean Rabault, Aurélien Larcher, Alexander Kuhnle, and Elie Hachem. A review on deep reinforcement learning for fluid mechanics. *Computers & Fluids*, 225:104973, 2021.
- Philipp Holl, Nils Thuerey, and Vladlen Koltun. Learning to control pdes with differentiable physics. In *International Conference on Learning Representations*, 2020.
- Jer-Nan Juang and Richard S Pappa. An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of guidance, control, and dynamics*, 8(5):620–627, 1985.
- Jer-Nan Juang, Minh Phan, Lucas G Horta, and Richard W Longman. Identification of observer/kalman filter markov parameters-theory and experiments. *Journal of Guidance, Control, and Dynamics*, 16(2): 320–329, 1993.
- Dante Kalise and Karl Kunisch. Polynomial approximation of high-dimensional hamilton–jacobi–bellman equations and applications to feedback control of semilinear parabolic pdes. *SIAM Journal on Scientific Computing*, 40(2):A629–A652, 2018.

- Hassan K. Khalil. *Nonlinear systems*. Prentice Hall, third edition, 2002.
- Mohammad Amin Khodkar, Pedram Hassanzadeh, and Athanasios 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.
- J Zico Kolter and Gaurav Manek. Learning stable deep dynamics models. *Advances in neural information processing systems*, 32, 2019.
- Ian Lenz, Ross A Knepper, and Ashutosh Saxena. Deepmpc: Learning deep latent features for model predictive control. In *Robotics: Science and Systems*, volume 10. Rome, Italy, 2015.
- Anders Logg, Kent-Andre Mardal, and Garth Wells. *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84. Springer Science & Business Media, 2012.
- Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. Pde-net: Learning pdes from data. In *International Conference on Machine Learning*, pp. 3208–3216. PMLR, 2018.
- Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1):4950, 2018.
- Pingchuan Ma, Yunsheng Tian, Zherong Pan, Bo Ren, and Dinesh Manocha. Fluid directed rigid body control using deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 37(4):1–11, 2018.
- Jan R Magnus and Heinz Neudecker. Symmetry, 0-1 matrices and jacobians: A review. *Econometric Theory*, 2(2):157–190, 1986.
- George Matsaglia and George PH Styan. Equalities and inequalities for ranks of matrices. *Linear and multilinear Algebra*, 2(3):269–292, 1974.
- Jeremy Morton, Antony Jameson, Mykel J Kochenderfer, and Freddie Witherden. Deep dynamical modeling and control of unsteady fluid flows. *Advances in Neural Information Processing Systems*, 31, 2018.
- Zuowei Ping, Zhun Yin, Xiuting Li, Yefeng Liu, and Tao Yang. Deep koopman model predictive control for enhancing transient stability in power grids. *International Journal of Robust and Nonlinear Control*, 31(6):1964–1978, 2021.
- Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016.
- Jean Rabault, Miroslav Kuchta, Atle Jensen, Ulysse Réglade, and Nicolas Cerardi. Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *Journal of fluid mechanics*, 865:281–302, 2019.
- Maziar Raissi. Deep hidden physics models: Deep learning of nonlinear partial differential equations. *The Journal of Machine Learning Research*, 19(1):932–955, 2018.
- Xiaoli Ren, Xiaoyong Li, Kaijun Ren, Junqiang Song, Zichen Xu, Kefeng Deng, and Xiang Wang. Deep learning-based weather prediction: a survey. *Big Data Research*, 23:100178, 2021.
- Clarence W Rowley, Igor Mezić, Shervin Bagheri, Philipp Schlatter, and Dan S Henningson. Spectral analysis of nonlinear flows. *Journal of fluid mechanics*, 641:115–127, 2009.
- Priyabrata Saha, Magnus Egerstedt, and Saibal Mukhopadhyay. Neural identification for control. *IEEE Robotics and Automation Letters*, 6(3):4648–4655, 2021.
- Michael Schäfer, Stefan Turek, Franz Durst, Egon Krause, and Rolf Rannacher. *Benchmark computations of laminar flow around a cylinder*. Springer, 1996.
- Sebastian Scher. Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with deep learning. *Geophysical Research Letters*, 45(22):12–616, 2018.

- Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.
- Sungyong Seo, Chuizheng Meng, and Yan Liu. Physics-aware difference graph networks for sparsely-observed dynamics. In *International Conference on Learning Representations*, 2019.
- Eduardo D Sontag. *Mathematical control theory: deterministic finite dimensional systems*, volume 6. Springer Science & Business Media, 2013.
- Prem A Srinivasan, L Guastoni, Hossein Azizpour, PHILIPP Schlatter, and Ricardo Vinuesa. Predictions of turbulent shear flows using deep neural networks. *Physical Review Fluids*, 4(5):054603, 2019.
- Tetsuya Takahashi, Junbang Liang, Yi-Ling Qiao, and Ming C Lin. Differentiable fluids with solid coupling for learning and control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35(7), pp. 6138–6146, 2021.
- Naoya Takeishi, Yoshinobu Kawahara, and Takehisa Yairi. Learning koopman invariant subspaces for dynamic mode decomposition. *Advances in neural information processing systems*, 30, 2017.
- Hongwei Tang, Jean Rabault, Alexander Kuhnle, Yan Wang, and Tongguang Wang. Robust active flow control over a range of reynolds numbers using an artificial neural network trained through deep reinforcement learning. *Physics of Fluids*, 32(5):053605, 2020.
- Jonathan H. Tu, , Clarence W. Rowley, Dirk M. Luchtenburg, Steven L. Brunton, and J. Nathan Kutz and. On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1(2): 391–421, 2014. doi: 10.3934/jcd.2014.1.391.
- Karen Willcox and Jaime Peraire. Balanced model reduction via the proper orthogonal decomposition. *AIAA journal*, 40(11):2323–2330, 2002.
- Charles HK Williamson. Vortex dynamics in the cylinder wake. *Annual review of fluid mechanics*, 28(1): 477–539, 1996.
- SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pp. 802–810, 2015.
- Yuting Yang, Junyu Dong, Xin Sun, Estanislau Lima, Quanquan Mu, and Xinhua Wang. A cfcc-lstm model for sea surface temperature prediction. *IEEE Geoscience and Remote Sensing Letters*, 15(2):207–211, 2017.
- Enoch Yeung, Soumya Kundu, and Nathan Hodas. Learning deep neural network representations for koopman operators of nonlinear dynamical systems. In *2019 American Control Conference (ACC)*, pp. 4832–4839. IEEE, 2019.
- Ruiyang Zhang, Yang Liu, and Hao Sun. Physics-informed multi-lstm networks for metamodeling of nonlinear structures. *Computer Methods in Applied Mechanics and Engineering*, 369:113226, 2020.

Appendices

A Proofs

This section details the proofs for the results presented in section 4. The proof of theorem 4.1.1 uses the following properties of the rank (denoted by $\text{rank}(\cdot)$), the Kronecker product (denoted by \otimes) and vectorization of matrices (denoted by $\text{vec}(\cdot)$). All the definitions and properties are presented in the context of matrices over real numbers.

For any conformable matrices \mathbf{D} and \mathbf{E} such that \mathbf{E} has full row-rank,

$$\text{rank}(\mathbf{D}\mathbf{E}) = \text{rank}(\mathbf{D}). \quad (37a)$$

For any real matrix \mathbf{D} ,

$$\text{rank}(\mathbf{D}^\top \mathbf{D}) = \text{rank}(\mathbf{D}\mathbf{D}^\top) = \text{rank}(\mathbf{D}^\top) = \text{rank}(\mathbf{D}). \quad (37b)$$

For any matrices (of compatible dimensions) $\mathbf{D}, \mathbf{E}, \mathbf{F}$, and \mathbf{H} ,

$$\text{vec}(\mathbf{D}\mathbf{E}\mathbf{F}^\top) = (\mathbf{F} \otimes \mathbf{D})\text{vec}(\mathbf{E}), \quad (38a)$$

$$(\mathbf{D} \otimes \mathbf{E})^\top = \mathbf{D}^\top \otimes \mathbf{E}^\top, \quad (38b)$$

$$(\mathbf{D} \otimes \mathbf{E})(\mathbf{F} \otimes \mathbf{H}) = (\mathbf{D}\mathbf{F} \otimes \mathbf{E}\mathbf{H}), \quad (38c)$$

whenever these quantities are defined. Furthermore, if \mathbf{D} and \mathbf{E} are symmetric and positive semidefinite (resp. positive definite), then $\mathbf{D} \otimes \mathbf{E}$ is symmetric and positive semidefinite (resp. positive definite), i.e.,

$$\mathbf{D} \succeq 0, \mathbf{E} \succeq 0 \implies (\mathbf{D} \otimes \mathbf{E}) \succeq 0; \quad \mathbf{D} \succ 0, \mathbf{E} \succ 0 \implies (\mathbf{D} \otimes \mathbf{E}) \succ 0. \quad (38d)$$

Proofs of (37) and (38) can be found in (Matsaglia & PH Styan (1974)) and (Magnus & Neudecker (1986)), respectively.

To derive the results presented in corollary (4.1.1.1), we use the following definitions of the Moore-Penrose inverse of a matrix (denoted by $(\cdot)^+$). For any matrix \mathbf{D} and its (full) SVD, i.e., $\mathbf{D} = \mathbf{U}_\mathbf{D} \boldsymbol{\Sigma}_\mathbf{D} \mathbf{V}_\mathbf{D}^\top$,

$$\mathbf{D}^+ = (\mathbf{D}^\top \mathbf{D})^{-1} \mathbf{D}^\top, \quad \text{when } (\mathbf{D}^\top \mathbf{D})^{-1} \text{ exists}, \quad (39a)$$

$$\mathbf{D}^+ = \mathbf{D}^\top (\mathbf{D}\mathbf{D}^\top)^{-1}, \quad \text{when } (\mathbf{D}\mathbf{D}^\top)^{-1} \text{ exists}, \quad (39b)$$

$$\mathbf{D}^+ = \mathbf{V}_\mathbf{D} \boldsymbol{\Sigma}_\mathbf{D}^+ \mathbf{U}_\mathbf{D}^\top, \quad (39c)$$

$$\mathbf{D}^+ = \lim_{\varepsilon \rightarrow 0} (\mathbf{D}^\top \mathbf{D} + \varepsilon^2 \mathbf{I})^{-1} \mathbf{D}^\top = \lim_{\varepsilon \rightarrow 0} \mathbf{D}^\top (\mathbf{D}\mathbf{D}^\top + \varepsilon^2 \mathbf{I})^{-1}, \quad (39d)$$

where \mathbf{I} is the identity matrix of compatible dimension. The proof of (39d) can be found in (Albert (1972)).

To prove Theorem 4.1.1, we use some well-known results, summarized as the following lemma in (Baldi & Hornik (1989)), for linear least-squares optimization.

Lemma A.0.1. *The quadratic function $L(\mathbf{z}) = \|\mathbf{y} - \mathbf{M}\mathbf{z}\|^2 = \mathbf{y}^\top \mathbf{y} - 2\mathbf{y}^\top \mathbf{M}\mathbf{z} + \mathbf{z}^\top \mathbf{M}^\top \mathbf{M}\mathbf{z}$ is convex, and a point \mathbf{z} globally minimizes L if and only if $\nabla L(\mathbf{z}) = 0$, or equivalently, $\mathbf{M}^\top \mathbf{M}\mathbf{z} = \mathbf{M}^\top \mathbf{y}$. Furthermore, if $\mathbf{M}^\top \mathbf{M} \succ 0$, i.e., positive definite, then L is strictly convex and reaches its unique minimum for $\mathbf{z} = (\mathbf{M}^\top \mathbf{M})^{-1} \mathbf{M}^\top \mathbf{y}$.*

A.1 Proof of theorem 4.1.1

Theorem 4.1.1. *Consider the following objective function*

$$L_{\text{pred}}(\mathbf{E}_\mathbf{x}, \mathbf{G}) = \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_\mathbf{x} \mathbf{x}(t_{i+1}) - \mathbf{G} \mathbf{E}_{\mathbf{x}\mathbf{u}} \boldsymbol{\omega}(t_i)\|^2, \quad (13)$$

where $\mathbf{G} = [\mathbf{A}_R \ \mathbf{B}_R] \in \mathbb{R}^{r_x \times (r_x + d_u)}$, $\mathbf{E}_{xu} = \begin{bmatrix} \mathbf{E}_x & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix} \in \mathbb{R}^{(r_x + d_u) \times (d_x + d_u)}$, \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

$$\mathbf{G}\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top = \mathbf{E}_x\mathbf{Y}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top, \quad (14)$$

where \mathbf{Y} and $\mathbf{\Omega}$ are the data matrices as defined in section (3.3). If \mathbf{E}_x has full rank r_x , and $\mathbf{\Omega}\mathbf{\Omega}^\top$ is non-singular, then L_{pred} is strictly convex and has a unique minimum for

$$\mathbf{G} = [\mathbf{A}_R \ \mathbf{B}_R] = \mathbf{E}_x\mathbf{Y}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top(\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top)^{-1}. \quad (15)$$

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

$$\begin{aligned} L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) &= \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x\mathbf{x}(t_{i+1}) - \mathbf{G}\mathbf{E}_{xu}\boldsymbol{\omega}(t_i)\|^2 \\ &= \|\text{vec}(\mathbf{E}_x\mathbf{Y}) - \text{vec}(\mathbf{G}\mathbf{E}_{xu}\mathbf{\Omega})\|^2 \\ &= \|\text{vec}(\mathbf{E}_x\mathbf{Y}) - (\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})\text{vec}(\mathbf{G})\|^2. \end{aligned} \quad (40)$$

The third equality is obtained using (38a). For fixed \mathbf{E}_x , we can apply Lemma A.0.1 to (40): (40) is convex in coefficient of \mathbf{G} , and \mathbf{G} corresponds to a global minimum of L_{pred} if and only if

$$(\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})^\top(\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})\text{vec}(\mathbf{G}) = (\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})^\top\text{vec}(\mathbf{E}_x\mathbf{Y}). \quad (41)$$

Using (38b) and (38c), we can write (41) as

$$(\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})\text{vec}(\mathbf{G}) = (\mathbf{E}_{xu}\mathbf{\Omega} \otimes \mathbf{I}_{r_x})\text{vec}(\mathbf{E}_x\mathbf{Y}). \quad (42)$$

Applying (38a) on (42), we get $\mathbf{G}\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top = \mathbf{E}_x\mathbf{Y}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top$, i.e., (14).

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 $\mathbf{\Omega}\mathbf{\Omega}^\top \in \mathbb{R}^{(d_x + d_u) \times (d_x + d_u)}$ is non-singular, then $\mathbf{\Omega}$ has full row-rank $(d_x + d_u)$. Consequently, using (37a) and (37b), we have

$$\text{rank}(\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top) = \text{rank}(\mathbf{E}_{xu}\mathbf{\Omega}) = \text{rank}(\mathbf{E}_{xu}) = r_x + d_u. \quad (43)$$

Hence the symmetric positive semidefinite matrix $\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top$ has full rank and therefore positive definite. Using (38b), (38c), and (38d), we can see that $(\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})^\top(\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) = (\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x})$ is positive definite as well. Therefore, by Lemma A.0.1, (40) is strictly convex in the coefficients of \mathbf{G} and has a unique minimum. Since $\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top \succ 0$, it is invertible. Hence, from (14), we can say that the unique minimum of (40) is reached at $\mathbf{G} = \mathbf{E}_x\mathbf{Y}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top(\mathbf{E}_{xu}\mathbf{\Omega}\mathbf{\Omega}^\top\mathbf{E}_{xu}^\top)^{-1}$, i.e., (15). ■

A.2 An alternative representation of (15)

Here we provide a possible alternative representation of (15) required to prove corollary 4.1.1.1.

Lemma A.2.1. Consider the (full) SVD of the data matrix $\mathbf{\Omega}$ given by $\mathbf{\Omega} = \mathbf{U}_\Omega\boldsymbol{\Sigma}_\Omega\mathbf{V}_\Omega^\top$, 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 $\mathbf{V}_\Omega \in \mathbb{R}^{n \times n}$. (15) can be expressed as

$$\mathbf{G} = \lim_{\varepsilon \rightarrow 0} \mathbf{E}_x\mathbf{Y}\mathbf{V}_\Omega(\boldsymbol{\Sigma}_\Omega^\top\mathbf{U}_\Omega^\top\mathbf{E}_{xu}^\top\mathbf{E}_{xu}\mathbf{U}_\Omega\boldsymbol{\Sigma}_\Omega + \varepsilon^2\mathbf{I}_n)^{-1}\boldsymbol{\Sigma}_\Omega^\top\mathbf{U}_\Omega^\top\mathbf{E}_{xu}^\top. \quad (44)$$

Proof. Replacing $\mathbf{\Omega}$ with its SVD in (15) we get,

$$\begin{aligned} \mathbf{G} &= \mathbf{E}_x\mathbf{Y}\mathbf{V}_\Omega\boldsymbol{\Sigma}_\Omega^\top\mathbf{U}_\Omega^\top\mathbf{E}_{xu}^\top(\mathbf{E}_{xu}\mathbf{U}_\Omega\boldsymbol{\Sigma}_\Omega\mathbf{V}_\Omega^\top\mathbf{V}_\Omega\boldsymbol{\Sigma}_\Omega^\top\mathbf{U}_\Omega^\top\mathbf{E}_{xu}^\top)^{-1} \\ &= \mathbf{E}_x\mathbf{Y}\mathbf{V}_\Omega\boldsymbol{\Sigma}_\Omega^\top\mathbf{U}_\Omega^\top\mathbf{E}_{xu}^\top(\mathbf{E}_{xu}\mathbf{U}_\Omega\boldsymbol{\Sigma}_\Omega\boldsymbol{\Sigma}_\Omega^\top\mathbf{U}_\Omega^\top\mathbf{E}_{xu}^\top)^{-1} \\ &= \mathbf{E}_x\mathbf{Y}\mathbf{V}_\Omega(\mathbf{E}_{xu}\mathbf{U}_\Omega\boldsymbol{\Sigma}_\Omega)^\top \end{aligned} \quad (45)$$

The second equality is due to the orthogonality of \mathbf{V}_Ω . The third equality is obtained using (39b). Substituting $(\mathbf{E}_{xu}\mathbf{U}_\Omega\Sigma_\Omega)^+$ with the *limit definition* (39d) of the Moore-Penrose inverse, we get

$$\mathbf{G} = \lim_{\varepsilon \rightarrow 0} \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega (\Sigma_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega \Sigma_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \Sigma_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top. \quad (46)$$

A.3 Proof of Corollary 4.1.1.1

Corollary 4.1.1.1. *Consider the (full) SVD of the data matrix Ω given by $\Omega = \mathbf{U}_\Omega \Sigma_\Omega \mathbf{V}_\Omega^\top$, where $\mathbf{U}_\Omega \in \mathbb{R}^{(d_x+d_u) \times (d_x+d_u)}$, $\Sigma_\Omega \in \mathbb{R}^{(d_x+d_u) \times n}$, and $\mathbf{V}_\Omega \in \mathbb{R}^{n \times n}$. If $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$ and $\Omega\Omega^\top$ is non-singular, then the solution for $\mathbf{G} = [\mathbf{A}_R \ \mathbf{B}_R]$ corresponding to the unique minimum of L_{pred} can be expressed as*

$$\mathbf{A}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \Sigma^* \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y, \quad \text{and} \quad \mathbf{B}_R = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega \Sigma^* \mathbf{U}_{\Omega,2}^\top, \quad (16)$$

where $[\mathbf{U}_{\Omega,1}^\top \ \mathbf{U}_{\Omega,2}^\top] = \mathbf{U}_\Omega^\top$ 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 $\Sigma^* = \lim_{\varepsilon \rightarrow 0} (\Sigma_\Omega^\top \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \Sigma_\Omega + \Sigma_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \Sigma_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \Sigma_\Omega^\top$.

Proof. By the definition of truncated SVD, the columns of $\hat{\mathbf{U}}_Y$ are orthonormal. Therefore, $\hat{\mathbf{U}}_Y^\top$ has full row-rank r_x . Hence, by theorem 4.1.1 and lemma A.2.1, if $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$, and $\Omega\Omega^\top$ is non-singular, then the unique minimum of L_{pred} , is reached when

$$\mathbf{G} = \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega (\mathbf{E}_{xu} \mathbf{U}_\Omega \Sigma_\Omega)^+ = \lim_{\varepsilon \rightarrow 0} \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega (\Sigma_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega \Sigma_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \Sigma_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top. \quad (47)$$

Now, substituting $\mathbf{E}_x = \hat{\mathbf{U}}_Y^\top$ in \mathbf{E}_{xu} , and using the partition $\mathbf{U}_\Omega^\top = [\mathbf{U}_{\Omega,1}^\top \ \mathbf{U}_{\Omega,2}^\top]$, 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

$$\mathbf{E}_{xu} \mathbf{U}_\Omega = \begin{bmatrix} \hat{\mathbf{U}}_Y^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_u} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{\Omega,1} \\ \mathbf{U}_{\Omega,2} \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \\ \mathbf{U}_{\Omega,2} \end{bmatrix}, \quad (48)$$

and

$$\mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top \mathbf{E}_{xu} \mathbf{U}_\Omega = \begin{bmatrix} \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y & \mathbf{U}_{\Omega,2}^\top \end{bmatrix} \begin{bmatrix} \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \\ \mathbf{U}_{\Omega,2} \end{bmatrix} = \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} + \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2}. \quad (49)$$

Plugging (48) and (49) into (47) leads to

$$\mathbf{G} = \lim_{\varepsilon \rightarrow 0} \hat{\mathbf{U}}_Y^\top \mathbf{Y} \mathbf{V}_\Omega (\Sigma_\Omega^\top \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \Sigma_\Omega + \Sigma_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \Sigma_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \Sigma_\Omega^\top \begin{bmatrix} \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y & \mathbf{U}_{\Omega,2}^\top \end{bmatrix}. \quad (50)$$

Defining $\Sigma^* \triangleq \lim_{\varepsilon \rightarrow 0} (\Sigma_\Omega^\top \mathbf{U}_{\Omega,1}^\top \hat{\mathbf{U}}_Y \hat{\mathbf{U}}_Y^\top \mathbf{U}_{\Omega,1} \Sigma_\Omega + \Sigma_\Omega^\top \mathbf{U}_{\Omega,2}^\top \mathbf{U}_{\Omega,2} \Sigma_\Omega + \varepsilon^2 \mathbf{I}_n)^{-1} \Sigma_\Omega^\top$, we can split (50) into

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

which is (16). ■

A.4 The case when $\Omega\Omega^\top$ not invertible

When the covariance matrix $\Omega\Omega^\top$ is not invertible, which is always true if $n < d_x + d_u$, the matrix $\mathbf{E}_{xu} \Omega \Omega^\top \mathbf{E}_{xu}^\top$ is not guaranteed to be invertible. In that case, the minimum of L_{pred} corresponds to infinitely many solutions for \mathbf{G} . However, minimizing L_{pred} with added ℓ_2 regularization, i.e., $L_{\text{pred,reg}}(\mathbf{E}_x, \mathbf{G}) = L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) + \beta \|\text{vec}(\mathbf{G})\|^2$ provides a unique solution for \mathbf{G} , for a fixed \mathbf{E}_x . We have the following result.

Theorem A.4.1. *For any fixed matrix \mathbf{E}_x and $\beta > 0$, the objective function $L_{\text{pred,reg}}(\mathbf{E}_x, \mathbf{G}) = L_{\text{pred}}(\mathbf{E}_x, \mathbf{G}) + \beta \|\text{vec}(\mathbf{G})\|^2$ is strictly convex in the coefficients of \mathbf{G} , and the global minimum of $L_{\text{pred,reg}}$ corresponds to the unique solution for \mathbf{G} , given by*

$$\mathbf{G} = \mathbf{E}_x \mathbf{Y} \Omega^\top \mathbf{E}_{xu}^\top (\mathbf{E}_{xu} \Omega \Omega^\top \mathbf{E}_{xu}^\top + \beta \mathbf{I}_{r_x+d_u})^{-1}. \quad (51)$$

Proof. $L_{\text{pred,reg}}(\mathbf{E}_x, \mathbf{G})$ can be written as, using (38a-c),

$$\begin{aligned} L_{\text{pred,reg}}(\mathbf{E}_x, \mathbf{G}) &= \|\text{vec}(\mathbf{E}_x \mathbf{Y}) - (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{G})\|^2 + \beta \|\text{vec}(\mathbf{G})\|^2 \\ &= \text{vec}(\mathbf{E}_x \mathbf{Y})^\top \text{vec}(\mathbf{E}_x \mathbf{Y}) - 2 \text{vec}(\mathbf{E}_x \mathbf{Y})^\top (\boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\mathbf{G}) \\ &\quad + \text{vec}(\mathbf{G})^\top (\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x} + \beta \mathbf{I}_{r_x(r_x+d_u)}) \text{vec}(\mathbf{G}) \end{aligned}$$

$\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top$ is a symmetric positive semidefinite matrix, irrespective of whether it has full rank or not. Hence, by (38d), $\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x}$ is symmetric positive semidefinite. Consequently, for any $\beta > 0$, $\mathbf{E}_{xu} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top \mathbf{E}_{xu}^\top \otimes \mathbf{I}_{r_x} + \beta \mathbf{I}_{r_x(r_x+d_u)}$ is positive definite. According to lemma A.0.1, $L_{\text{pred,reg}}$ is therefore strictly convex in the coefficients of \mathbf{G} and globally minimized when $\nabla L_{\text{pred,reg}} = 0$. The unique solution of (51) can be derived in the same manner as theorem 4.1.1. ■

Remark. Replacing $\boldsymbol{\Omega}$ with its SVD in (51) we get,

$$\mathbf{G} = \mathbf{E}_x \mathbf{Y} \mathbf{V}_\Omega \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top (\mathbf{E}_{xu} \mathbf{U}_\Omega \boldsymbol{\Sigma}_\Omega \boldsymbol{\Sigma}_\Omega^\top \mathbf{U}_\Omega^\top \mathbf{E}_{xu}^\top + \beta \mathbf{I}_{r_x+d_u})^{-1}. \quad (52)$$

In the limit $\beta \rightarrow 0^+$, (52) converges to (45).

A.5 DMDc through a linear autoencoding structure

Here we present a linear autoencoding structure that leads to a linear ROM exactly resembling the DMDc solution when $\mathbf{E}_x = \widehat{\mathbf{U}}_Y^\top$. However, its DNN-based nonlinear counterpart does not actually offer dimensionality reduction.

Theorem A.5.1. *Consider the following objective function*

$$L_{\text{pred,alt}}(\mathbf{E}_x, \tilde{\mathbf{G}}) = \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x \mathbf{x}(t_{i+1}) - \tilde{\mathbf{G}} \boldsymbol{\omega}(t_i)\|^2, \quad (53)$$

where $\tilde{\mathbf{G}} \in \mathbb{R}^{r_x \times (d_x+d_u)}$. For any fixed matrix \mathbf{E}_x , the objective function $L_{\text{pred,alt}}$ is convex in the coefficients of $\tilde{\mathbf{G}}$ and attains its minimum for any $\tilde{\mathbf{G}}$ satisfying

$$\tilde{\mathbf{G}} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top, \quad (54)$$

where \mathbf{Y} and $\boldsymbol{\Omega}$ are the data matrices as defined in section (3.3). If $\boldsymbol{\Omega} \boldsymbol{\Omega}^\top$ is non-singular, then $L_{\text{pred,alt}}$ is strictly convex and has a unique minimum for

$$\tilde{\mathbf{G}} = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top (\boldsymbol{\Omega} \boldsymbol{\Omega}^\top)^{-1}. \quad (55)$$

Proof. The proof is very similar to the proof of theorem 4.1.1. Using (38a), we can write $L_{\text{pred,alt}}(\mathbf{E}_x, \tilde{\mathbf{G}})$ as follows,

$$\begin{aligned} L_{\text{pred,alt}}(\mathbf{E}_x, \tilde{\mathbf{G}}) &= \frac{1}{n} \sum_{i=0}^{n-1} \|\mathbf{E}_x \mathbf{x}(t_{i+1}) - \tilde{\mathbf{G}} \boldsymbol{\omega}(t_i)\|^2 \\ &= \|\text{vec}(\mathbf{E}_x \mathbf{Y}) - \text{vec}(\tilde{\mathbf{G}} \boldsymbol{\Omega})\|^2 \\ &= \|\text{vec}(\mathbf{E}_x \mathbf{Y}) - (\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\tilde{\mathbf{G}})\|^2. \end{aligned} \quad (56)$$

For fixed \mathbf{E}_x , applying Lemma A.0.1 to (56), we can say $L_{\text{pred,alt}}$ is convex in the coefficients of $\tilde{\mathbf{G}}$, and $\tilde{\mathbf{G}}$ corresponds to a global minimum of $L_{\text{pred,alt}}$ if and only if

$$(\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_x})^\top (\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_x}) \text{vec}(\tilde{\mathbf{G}}) = (\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_x})^\top \text{vec}(\mathbf{E}_x \mathbf{Y}). \quad (57)$$

Using (38a-c), we can write (57) as $\tilde{\mathbf{G}} \boldsymbol{\Omega} \boldsymbol{\Omega}^\top = \mathbf{E}_x \mathbf{Y} \boldsymbol{\Omega}^\top$, which is (54).

If $\boldsymbol{\Omega}\boldsymbol{\Omega}^\top$ is non-singular, then it is symmetric positive definite. Using (38b-d), we can see that $(\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_{\mathbf{x}}})^\top (\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_{\mathbf{x}}}) = (\boldsymbol{\Omega}\boldsymbol{\Omega}^\top \otimes \mathbf{I}_{r_{\mathbf{x}}})$ is positive definite as well. Therefore, by Lemma A.0.1, (56) is strictly convex in coefficient in $\tilde{\mathbf{G}}$ and has a unique minimum. In that case, from (54), we can say that the unique minimum of (56) is reached at $\tilde{\mathbf{G}} = \mathbf{E}_{\mathbf{x}} \mathbf{Y} \boldsymbol{\Omega}^\top (\boldsymbol{\Omega}\boldsymbol{\Omega}^\top)^{-1}$, i.e., (55). ■

Corollary A.5.1.1. *Consider the (full) SVD of the data matrix $\boldsymbol{\Omega}$ given by $\boldsymbol{\Omega} = \mathbf{U}_{\boldsymbol{\Omega}} \boldsymbol{\Sigma}_{\boldsymbol{\Omega}} \mathbf{V}_{\boldsymbol{\Omega}}^\top$, where $\mathbf{U}_{\boldsymbol{\Omega}} \in \mathbb{R}^{(d_{\mathbf{x}}+d_{\mathbf{u}}) \times (d_{\mathbf{x}}+d_{\mathbf{u}})}$, $\boldsymbol{\Sigma}_{\boldsymbol{\Omega}} \in \mathbb{R}^{(d_{\mathbf{x}}+d_{\mathbf{u}}) \times n}$, and $\mathbf{V}_{\boldsymbol{\Omega}} \in \mathbb{R}^{n \times n}$. If $\mathbf{E}_{\mathbf{x}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top$ and $\boldsymbol{\Omega}\boldsymbol{\Omega}^\top$ is non-singular, then the solution for $\tilde{\mathbf{G}}$ corresponding to the unique minimum of $L_{\text{pred,alt}}$ can be expressed as*

$$\tilde{\mathbf{G}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \mathbf{V}_{\boldsymbol{\Omega}} \boldsymbol{\Sigma}_{\boldsymbol{\Omega}}^+ \mathbf{U}_{\boldsymbol{\Omega}}^\top. \quad (58)$$

Proof. By theorem A.5.1, if $\mathbf{E}_{\mathbf{x}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top$, and $\boldsymbol{\Omega}\boldsymbol{\Omega}^\top$ is non-singular, then the unique minimum of $L_{\text{pred,alt}}$ is reached when

$$\tilde{\mathbf{G}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \boldsymbol{\Omega}^\top (\boldsymbol{\Omega}\boldsymbol{\Omega}^\top)^{-1} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \boldsymbol{\Omega}^+ \quad (59)$$

The second equality is due to (39b). Substituting $\boldsymbol{\Omega}^+$ with its SVD definition (39c) into (59), we get $\hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \mathbf{V}_{\boldsymbol{\Omega}} \boldsymbol{\Sigma}_{\boldsymbol{\Omega}}^+ \mathbf{U}_{\boldsymbol{\Omega}}^\top$, which is (58). ■

Remark. From (53), it can be seen that $\tilde{\mathbf{G}}$ maps the concatenated vector, $\boldsymbol{\omega}(t_i)$, of full state and actuation to the next reduce state $\mathbf{x}_{\text{R}}(t_{i+1})$. We can partition (58) as $\tilde{\mathbf{G}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \mathbf{V}_{\boldsymbol{\Omega}} \boldsymbol{\Sigma}_{\boldsymbol{\Omega}}^+ [\mathbf{U}_{\boldsymbol{\Omega},1}^\top \quad \mathbf{U}_{\boldsymbol{\Omega},2}^\top] = [\tilde{\mathbf{A}} \quad \tilde{\mathbf{B}}]$ to separate out the blocks corresponding to state and actuation. Here, $\mathbf{U}_{\boldsymbol{\Omega},1}, \mathbf{U}_{\boldsymbol{\Omega},2}$ are the same as defined in corollary 4.1.1.1, and $\tilde{\mathbf{A}} \in \mathbb{R}^{r_{\mathbf{x}} \times d_{\mathbf{x}}}$, $\tilde{\mathbf{B}} \in \mathbb{R}^{r_{\mathbf{x}} \times d_{\mathbf{u}}}$. Now, if we post-multiply $\tilde{\mathbf{A}}$ with $\mathbf{E}_{\mathbf{x}}^\top = \hat{\mathbf{U}}_{\mathbf{Y}} \in \mathbb{R}^{d_{\mathbf{x}} \times r_{\mathbf{x}}}$, we get a ROM

$$\tilde{\mathbf{A}}_{\text{R}} = \tilde{\mathbf{A}} \hat{\mathbf{U}}_{\mathbf{Y}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \mathbf{V}_{\boldsymbol{\Omega}} \boldsymbol{\Sigma}_{\boldsymbol{\Omega}}^+ \mathbf{U}_{\boldsymbol{\Omega},1}^\top \hat{\mathbf{U}}_{\mathbf{Y}}, \quad \tilde{\mathbf{B}}_{\text{R}} = \tilde{\mathbf{B}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \mathbf{V}_{\boldsymbol{\Omega}} \boldsymbol{\Sigma}_{\boldsymbol{\Omega}}^+ \mathbf{U}_{\boldsymbol{\Omega},2}^\top, \quad (60)$$

which maps the current reduced state $\mathbf{x}_{\text{R}}(t_i)$ and actuation $\mathbf{u}(t_i)$ to the next reduced state $\mathbf{x}_{\text{R}}(t_{i+1})$. It can be verified easily that if we use the truncated SVD (as defined by 10), instead of the full SVD, for $\boldsymbol{\Omega}$ in (59) and follow the similar steps afterward, we get an approximation of (60):

$$\hat{\mathbf{A}}_{\text{R}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \hat{\mathbf{V}}_{\boldsymbol{\Omega}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\Omega}}^{-1} \hat{\mathbf{U}}_{\boldsymbol{\Omega},1}^\top \hat{\mathbf{U}}_{\mathbf{Y}} = \mathbf{A}_{\text{R,DMDc}}; \quad \hat{\mathbf{B}}_{\text{R}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top \mathbf{Y} \hat{\mathbf{V}}_{\boldsymbol{\Omega}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\Omega}}^{-1} \hat{\mathbf{U}}_{\boldsymbol{\Omega},2}^\top = \mathbf{B}_{\text{R,DMDc}}.$$

In summary, the aforementioned method can be carried out using gradient descent-based optimization and leads to the same ROM as DMDc, when $\mathbf{E}_{\mathbf{x}} = \hat{\mathbf{U}}_{\mathbf{Y}}^\top$. However, in this method, the benefit of dimensionality reduction is realized only when linear networks are used. A nonlinear counterpart (a DNN in the context of this paper) of $\tilde{\mathbf{A}}_{\text{R}}$, i.e., a nonlinear mapping from $\mathbb{R}^{r_{\mathbf{x}}}$ to $\mathbb{R}^{r_{\mathbf{x}}}$, cannot be pre-computed from a nonlinear counterpart of $\tilde{\mathbf{G}}$, unlike the linear case (60). Consequently, we lose the benefit of dimensionality reduction when nonlinear networks are used.

A.6 Proof of theorem 4.2.1

Theorem 4.2.1. *Consider the target dynamics defined by (29) and the candidate Lyapunov function defined by (30). Suppose the difference between the target dynamics and the closed-loop dynamics satisfies*

$$\|\mathcal{F}(\mathbf{x}_{\text{R}}, \mathcal{E}_{\mathbf{u}} \circ \Pi(\mathbf{x}_{\text{R}})) - \mathcal{F}_{\text{s}}(\mathbf{x}_{\text{R}})\| \leq \delta < \frac{\alpha\theta\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})} \sqrt{\frac{\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}} \eta, \quad (31)$$

for all $\mathbf{x}_{\text{R}} \in \mathbb{X}_{\text{R}} = \{\mathbf{x}_{\text{R}} \in \mathbb{R}^{r_{\mathbf{x}}} \mid \|\mathbf{x}_{\text{R}}\| < \eta\}$ and $0 < \theta < 1$. Then, for all initial points satisfying $\|\mathbf{x}_{\text{R}}(t_0)\| < \sqrt{\frac{\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}} \eta$, the solution of the closed-loop ROM $\frac{d\mathbf{x}_{\text{R}}}{dt} = \mathcal{F}(\mathbf{x}_{\text{R}}, \mathcal{E}_{\mathbf{u}} \circ \Pi(\mathbf{x}_{\text{R}}))$ satisfies

$$\|\mathbf{x}_{\text{R}}(t)\| \leq \lambda e^{-\gamma(t-t_0)} \|\mathbf{x}_{\text{R}}(t_0)\|, \quad \forall t_0 \leq t < t_c + t_0 \quad (32)$$

and

$$\|\mathbf{x}_{\text{R}}(t)\| \leq \frac{2\delta}{\alpha\theta} \lambda^3, \quad \forall t \geq t_c + t_0 \quad (33)$$

for some finite $t_c > 0$, where

$$\gamma = \frac{\alpha(1-\theta)\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})} \quad \text{and} \quad \lambda = \sqrt{\frac{\lambda_{\max}(\mathbf{K})}{\lambda_{\min}(\mathbf{K})}} \quad (34)$$

Proof. From the definition of \mathcal{V}_R , we have

$$\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 \leq \mathcal{V}_R(\mathbf{x}_R) \leq \lambda_{\max}(\mathbf{K})\|\mathbf{x}_R\|^2, \quad \forall \mathbf{x}_R \in \mathbb{R}^{r_\mathbf{x}}, \quad (61)$$

where $\lambda_{\min}(\mathbf{K})$ and $\lambda_{\max}(\mathbf{K})$ denote the smallest and largest eigenvalues, respectively, of \mathbf{K} and have positive values since the matrix \mathbf{K} is positive definite. Moreover, the definition of the target dynamics (29) implies

$$\begin{aligned} \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{F}_s(\mathbf{x}_R) &= \begin{cases} \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R), & \text{if } \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R) \leq -\alpha \mathcal{V}_R(\mathbf{x}_R) \\ \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R) - \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \frac{\nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R) + \alpha \mathcal{V}_R(\mathbf{x}_R)}{\|\nabla \mathcal{V}_R(\mathbf{x}_R)\|^2} \nabla \mathcal{V}_R(\mathbf{x}_R), & \text{otherwise} \end{cases} \\ &= \begin{cases} \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R), & \text{if } \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{P}(\mathbf{x}_R) \leq -\alpha \mathcal{V}_R(\mathbf{x}_R) \\ -\alpha \mathcal{V}_R(\mathbf{x}_R), & \text{otherwise} \end{cases} \\ &\leq -\alpha \mathcal{V}_R(\mathbf{x}_R) \\ &\leq -\alpha \lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2, \quad \forall \mathbf{x}_R \in \mathbb{R}^{r_\mathbf{x}}. \end{aligned} \quad (62)$$

The last inequality is due to (61).

Now, assume $\mathcal{F}(\mathbf{x}_R, \mathcal{E}_u \circ \Pi(\mathbf{x}_R)) = \mathcal{H}(\mathbf{x}_R) = \mathcal{F}_s(\mathbf{x}_R) + \mathcal{J}(\mathbf{x}_R)$ for some function $\mathcal{J} : \mathbb{R}^{r_\mathbf{x}} \rightarrow \mathbb{R}^{r_\mathbf{x}}$ and consider $\mathcal{V}_R(\mathbf{x}_R) = \mathbf{x}_R^\top \mathbf{K} \mathbf{x}_R$ as a candidate Lyapunov function for

$$\frac{d\mathbf{x}_R}{dt} = \mathcal{H}(\mathbf{x}_R) = \mathcal{F}_s(\mathbf{x}_R) + \mathcal{J}(\mathbf{x}_R). \quad (63)$$

We have $\|\nabla \mathcal{V}_R(\mathbf{x}_R)\| = \|2\mathbf{K}\mathbf{x}_R\| \leq 2\lambda_{\max}(\mathbf{K})\|\mathbf{x}_R\|$. The time-derivative of \mathcal{V}_R along the trajectories of (63) satisfies

$$\begin{aligned} \frac{d\mathcal{V}_R}{dt} &= \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{F}_s(\mathbf{x}_R) + \nabla \mathcal{V}_R(\mathbf{x}_R)^\top \mathcal{J}(\mathbf{x}_R) \\ &\leq -\alpha \lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 + \|\nabla \mathcal{V}_R(\mathbf{x}_R)\| \|\mathcal{J}(\mathbf{x}_R)\| \\ &\leq -\alpha \lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 + 2\lambda_{\max}(\mathbf{K})\|\mathbf{x}_R\| \delta, \quad \forall \|\mathbf{x}_R\| < \eta \\ &= -\alpha(1-\theta)\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 - \alpha\theta\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 + 2\lambda_{\max}(\mathbf{K})\|\mathbf{x}_R\| \delta, \quad 0 < \theta < 1, \forall \|\mathbf{x}_R\| < \eta \\ &\leq -\alpha(1-\theta)\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 < 0, \quad \text{when } \eta > \|\mathbf{x}_R\| \geq \frac{2\delta\lambda_{\max}(\mathbf{K})}{\alpha\theta\lambda_{\min}(\mathbf{K})} \triangleq \mu. \end{aligned} \quad (64)$$

The second inequality is obtained using (62) and the third inequality is obtained using (31). Clearly, we have a non-empty region where $\frac{d\mathcal{V}_R}{dt} < 0$ only when

$$\delta < \frac{\alpha\theta\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})} \eta. \quad (65)$$

Let $b = \lambda_{\min}(\mathbf{K})\eta^2$ and $c = \lambda_{\max}(\mathbf{K})\mu^2$. Consider the sublevel sets $\chi_b = \{\mathbf{x}_R \in \mathbb{R}^{r_\mathbf{x}} \mid \mathcal{V}_R(\mathbf{x}_R) < b\}$ and $\chi_c = \{\mathbf{x}_R \in \mathbb{R}^{r_\mathbf{x}} \mid \mathcal{V}_R(\mathbf{x}_R) \leq c\}$. It can be easily verified that if $\delta < \frac{\alpha\theta\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})} \sqrt{\frac{\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}} \eta$, then $c < b$, which implies $\chi_c \subset \chi_b$. Note, this condition satisfies the necessary condition (65) for the non-empty region since $\lambda_{\min}(\mathbf{K}) \leq \lambda_{\max}(\mathbf{K})$.

For any \mathbf{x}_R inside χ_b , using (61), we have

$$\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R\|^2 \leq \mathcal{V}_R(\mathbf{x}_R) < b = \lambda_{\min}(\mathbf{K})\eta^2, \quad (66)$$

implying $\|\mathbf{x}_R\| < \eta$. Similarly, for any \mathbf{x}_R on the boundary or outside of χ_c , we have

$$\lambda_{\max}(\mathbf{K})\mu^2 = c \leq \mathcal{V}_R(\mathbf{x}_R) \leq \lambda_{\max}(\mathbf{K})\|\mathbf{x}_R\|^2, \quad (67)$$

which implies $\|\mathbf{x}_R\| \geq \mu$.

Combining (66) and (67) we can say for any \mathbf{x}_R outside (including the boundary) of χ_c , but inside χ_b , (64) holds true. For such \mathbf{x}_R (i.e. $\mathbf{x}_R \in \chi_b \setminus \chi_c$) we have

$$\frac{d\mathcal{V}_R}{dt} \leq -\frac{\alpha(1-\theta)\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}\mathcal{V}_R(\mathbf{x}_R) \triangleq -2\gamma\mathcal{V}_R(\mathbf{x}_R), \quad (68)$$

using (61) and (64).

If the initial point (at time t_0) satisfies $\|\mathbf{x}_R(t_0)\| < \sqrt{\frac{\lambda_{\min}(\mathbf{K})}{\lambda_{\max}(\mathbf{K})}}\eta$, then by (61),

$$b = \lambda_{\min}(\mathbf{K})\eta^2 > \lambda_{\max}(\mathbf{K})\|\mathbf{x}_R(t_0)\|^2 \geq \mathcal{V}_R(\mathbf{x}_R(t_0)),$$

which implies the initial point $\mathbf{x}_R(t_0)$ is inside χ_b . Assuming an initial point in $\chi_b \setminus \chi_c$, and integrating (68) in time interval $[t_0, t]$, we get

$$\mathcal{V}_R(\mathbf{x}_R(t)) \leq \mathcal{V}_R(\mathbf{x}_R(t_0))e^{-2\gamma(t-t_0)}. \quad (69)$$

Hence, $\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R(t)\|^2 \leq \mathcal{V}_R(\mathbf{x}_R(t)) \leq \mathcal{V}_R(\mathbf{x}_R(t_0))e^{-2\gamma(t-t_0)} \leq \lambda_{\max}(\mathbf{K})\|\mathbf{x}_R(t_0)\|^2e^{-2\gamma(t-t_0)}$ as long as $\mathbf{x}_R(t)$ remains outside of χ_c . Since $\frac{d\mathcal{V}_R}{dt}$ is always negative outside of χ_c , any trajectory starting outside of it, must enter χ_c in finite time. Let the trajectory starting at $\mathbf{x}_R(t_0)$ enters χ_c for the first time at time $t_c + t_0$. Then, we have

$$\|\mathbf{x}_R(t)\| \leq \sqrt{\frac{\lambda_{\max}(\mathbf{K})}{\lambda_{\min}(\mathbf{K})}}e^{-\gamma(t-t_0)}\|\mathbf{x}_R(t_0)\| = \lambda e^{-\gamma(t-t_0)}\|\mathbf{x}_R(t_0)\|, \quad \forall t_0 \leq t < t_c + t_0. \quad (70)$$

Once a trajectory enters χ_c , it cannot escape χ_c because $\frac{d\mathcal{V}_R}{dt}$ is negative on the boundary. Therefore, all points of a trajectory after $t \geq t_c + t_0$ satisfies $\lambda_{\min}(\mathbf{K})\|\mathbf{x}_R(t)\|^2 \leq \mathcal{V}_R(\mathbf{x}_R(t)) \leq c$, equivalently,

$$\|\mathbf{x}_R(t)\| \leq \sqrt{\frac{\lambda_{\max}(\mathbf{K})}{\lambda_{\min}(\mathbf{K})}}\mu = \frac{2\delta}{\alpha\theta} \left(\frac{\lambda_{\max}(\mathbf{K})}{\lambda_{\min}(\mathbf{K})} \right)^{3/2} = \frac{2\delta}{\alpha\theta}\lambda^3, \quad \forall t \geq t_c + t_0. \quad (71)$$

From (68), (70) and (71), we have $\gamma = \frac{\alpha(1-\theta)\lambda_{\min}(\mathbf{K})}{2\lambda_{\max}(\mathbf{K})}$ and $\lambda = \sqrt{\frac{\lambda_{\max}(\mathbf{K})}{\lambda_{\min}(\mathbf{K})}}$. ■

B Details on reaction–diffusion system experiment

B.1 Dataset

We use FEniCS (Logg et al. (2012)), an open-source computing platform for solving PDEs using the finite element method, with Python interface to generate the dataset. For the reaction-diffusion system of (35), we generate 100 training sequences of length 50 with time step size 0.01 and 256 nodes in \mathbb{I} . The initial conditions and actuations of these sequences are given by

$$q(\zeta, 0) = |a| \sum_{k=0}^4 b_k T_k(\zeta), \quad \zeta \in \mathbb{I}, \quad (72)$$

and

$$w(t_i) = 10g_i \max_{\zeta} |q(\zeta, t_{i-1})|, \quad i = 1, 2, \dots, 49, \quad (73)$$

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.




Figure 12: Architectures for all the DNN modules used in the reaction–diffusion experiment. The ‘Copy’ operation denotes the reuse of the same DNN block for zero and nonzero actuation. The ‘Concat’ operator concatenates the input features along the last dimension. The ‘Integrator’ performs the numerical integration for (22). The ‘Stability Criterion’ block implements (29).

B.2 DNN architectures

Figure 12 shows the DNN architectures used for different modules in the reaction–diffusion experiment. The state encoder comprises 1D convolutional layers, followed by fully connected layers. The state decoder has the reversed order with convolutional layers replaced by transposed convolutional layers. The ROM is designed by breaking the function \mathcal{F} into two components: $\mathcal{F}(\mathbf{x}_R, \mathbf{u}_R) = \mathcal{F}_{\text{auto}}(\mathbf{x}_R) + \mathcal{F}_{\text{forced}}(\mathbf{x}_R, \mathbf{u}_R) - \mathcal{F}_{\text{forced}}(\mathbf{x}_R, \mathbf{0})$. $\mathcal{F}_{\text{auto}}$ 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. Two multilayer perceptrons (MLPs) are used to implement $\mathcal{F}_{\text{auto}}$ and $\mathcal{F}_{\text{forced}}$. The output of the ROM is integrated using a numerical integrator to get the next state. The controller is implemented using an MLP. The target dynamics is implemented using another MLP, followed by a stability criterion in the form of (29).

B.3 Training settings

We use $r_x = 5$ in the prediction task and $r_x = 2$ in the control task for both DeepROM and DMDc. 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 (23). For learning control, we use $\beta_3 = 0.2$ in (27), $\alpha = 0.2$ in (29), and $\mathbf{K} = 0.5\mathbf{I}_{r_x}$ in (30). Since the learned ROMs from one training instance to another can vary, the hyperparameter pair (α, β_3) may require re-tuning accordingly.

C Details on vortex shedding suppression experiment

C.1 Dataset

For the flow past a circular cylinder problem, the geometry and physical parameters of the system are taken from the DFG 2D-2 benchmark (Schäfer et al. (1996)). The geometry is shown in Figure 13. We use the blue-shaded region for observation and actuation. Following the DFG 2D-2 benchmark, we use the no-slip boundary condition of zero velocity for the walls and the cylinder boundary, zero outlet pressure, and the inflow velocity profile (at the inlet) as

$$\mathbf{v}(\boldsymbol{\zeta}, t) = \left(1.5 \frac{4\zeta_2(0.41 - \zeta_2)}{0.41^2}, 0 \right), \quad (74)$$

where ζ_1 and ζ_2 denote the horizontal and vertical coordinates, respectively, of $\boldsymbol{\zeta}$. We use kinematic viscosity $\nu = 0.002$ and density $\rho = 1$ leading to the Reynolds number $Re = 50$. The training sequence of length 5000 is generated in FEniCS with a time step size 0.001 and applying actuations

$$\mathbf{w}(\boldsymbol{\zeta}, t) = a \sum_{k=0}^4 \begin{bmatrix} \sin(k\pi(\zeta_1 - 0.11)/0.66) & \sin(k\pi\zeta_2/0.41) \end{bmatrix} \begin{bmatrix} b_{k,1,1} & b_{k,2,1} \\ b_{k,1,2} & b_{k,2,2} \end{bmatrix}, \quad \boldsymbol{\zeta} \in \mathbb{W}, \quad (75)$$

where $a \sim \mathcal{U}(0, 1)$ and $b_{k,i,j} \sim \mathcal{U}(-1, 1)$, $i, j = 1, 2$ are chosen randomly. Similarly, a test sequence is generated to evaluate the prediction performance. For learning control, we use the Stokes flow or creeping flow as the desired state, which can be obtained by solving the Stokes equations

$$\nu \nabla^2 \mathbf{v} - \frac{1}{\rho} \nabla p = \mathbf{0}, \quad \nabla \cdot \mathbf{v} = 0 \quad \text{in } \mathbb{I} \times \mathbb{R}^+. \quad (76)$$

For training, the flow velocity data from the observation region (blue shaded in Figure 13) are interpolated onto a rectangular uniform grid of size 32×48 so that it can be used in standard CNNs.




Figure 13: Geometry of the flow past a circular cylinder set-up.

C.2 DNN architectures

Figure 14 shows the DNN architectures used for different modules in the vortex shedding control experiment. The architectures for the ROM and target dynamics are the same as in the previous example. Moreover, the state encoder and decoder have similar architectures as the previous example except for the 1D convolutions and transposed convolutions are replaced by their 2D counterparts. Here, an additional module is used: the control encoder for encoding the distributed control/actuation. It has the same architecture as the state encoder. To learn the distributed actuation, we design the controller as a linear combination of space-dependent polynomial basis functions. One MLP is used to learn these space-dependent polynomial basis functions given the locations of the actuation nodes and another MLP is used to learn the corresponding coefficients. We use this architecture instead of a standard convolutional one because the PDE solver takes the actuation input in a triangular mesh, not in a uniform rectangular grid. The polynomial basis architecture can be used to compute actuation in both uniform rectangular grid during training and triangular mesh during evaluation.




Figure 14: Architectures for all the DNN modules used in the fluid flow experiment. The ‘Split’ operator splits the input features into two vectors, along the last dimension.

C.3 Training settings

We use $r_x = 5$ in both the prediction task and control task for both DeepROM and DMDc. 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 (23). For learning control, we use $\beta_3 = 2$ in (27), $\alpha = 0.1$ in (29), and $\mathbf{K} = 0.5\mathbf{I}_{r_x}$ in (30). Since the learned ROMs from one training instance to another can vary, the hyperparameter pair (α, β_3) may require re-tuning accordingly.