Research Article

Long-Time Predictive Modeling of Nonlinear Dynamical Systems Using Neural Networks

Shaowu Pan and Karthik Duraisamy

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, USA

Correspondence should be addressed to Karthik Duraisamy; kdur@umich.edu

Received 31 May 2018; Revised 27 August 2018; Accepted 9 September 2018; Published 2 December 2018

Academic Editor: Francisco Chinesta

Copyright © 2018 Shaowu Pan and Karthik Duraisamy. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

We study the use of feedforward neural networks (FNN) to develop models of nonlinear dynamical systems from data. Emphasis is placed on predictions at long times, with limited data availability. Inspired by global stability analysis, and the observation of strong correlation between the local error and the maximal singular value of the Jacobian of the ANN, we introduce Jacobian regularization in the loss function. This regularization suppresses the sensitivity of the prediction to the local error and is shown to improve accuracy and robustness. Comparison between the proposed approach and sparse polynomial regression is presented in numerical examples ranging from simple ODE systems to nonlinear PDE systems including vortex shedding behind a cylinder and instability-driven buoyant mixing flow. Furthermore, limitations of feedforward neural networks are highlighted, especially when the training data does not include a low dimensional attractor. Strategies of data augmentation are presented as remedies to address these issues to a certain extent.

1. Introduction

The need to model dynamical behavior from data is pervasive across science and engineering. Applications are found in diverse fields such as in control systems [1], time series modeling [2], and describing the evolution of coherent structures [3]. While data-driven modeling of dynamical systems can be broadly classified as a special case of system identification [4], it is important to note certain distinguishing qualities: the learning process may be performed off-line, physical systems may involve very high dimensions, and the goal may involve the prediction of long-time behavior from limited training data.

Artificial neural networks (ANN) have attracted considerable attention in recent years in domains such as image recognition in computer vision [5, 6] and in control applications [7]. The success of ANNs arises from their ability to effectively learn low-dimensional representations from complex data and in building relationships between features and outputs. Neural networks with a single hidden layer and nonlinear activation function are guaranteed to be able to predict

any Borel measurable function to any degree of accuracy on a compact domain [8].

The idea of leveraging neural networks to model dynamical systems has been explored since the 1990s. ANNs are prevalent in the system identification and time series modeling community [9–12], where the mapping between inputs and outputs is of prime interest. Billings et al. [13] explored connections between neural networks and the nonlinear autoregressive moving average model (NARMAX) with exogenous inputs. It was shown that neural networks with one hidden layer and sigmoid activation function represent an infinite series consisting of polynomials of the input and state units. Elanayar and Shin [14] proposed the approximation of nonlinear stochastic dynamical systems using radial basis feedforward neural networks. Early work using neural networks to forecast multivariate time series of commodity prices [15] demonstrated its ability to model stochastic systems without knowledge of the underlying governing equations. Tsung and Cottrell [16] proposed learning the dynamics in phase space using a feedforward neural network with time-delayed coordinates.

Paez and Urbina [17–19] modeled a nonlinear hardening oscillator using a neural network-based model combined with dimension reduction using canonical variate analysis (CVA). Smaoui [20–22] pioneered the use of neural networks to predict fluid dynamic systems such as the unstable manifold model for bursting behavior in the 2-D Navier-Stokes and the Kuramoto-Sivashinsky equations. The dimensionality of the original PDE system is reduced by considering a small number of proper orthogonal decomposition (POD) coefficients [23]. Interestingly, similar ideas of using principal component analysis for dimension reduction can be traced back to work in cognitive science by Elman [24]. Elman also showed that knowledge of the intrinsic dimensions of the system can be very helpful in determining the structure of the neural network. However, in the majority of the results [20–22], the neural network model is only evaluated a few time steps from the training set, which might not be a stringent performance test if longer time predictions are of interest.

ANNs have also been applied to chaotic nonlinear systems that are challenging from a data-driven modeling perspective, especially if long time predictions are desired. Instead of minimizing the pointwise prediction error, Bakker et al. [25] satisfied the Diks’ criterion in learning the chaotic attractor. Later, Lin et al. [26] demonstrated that even the simplest feedforward neural network for nonlinear chaotic hydrodynamics can show consistency in the time-averaged characteristics, power spectra, and Lyapunov exponent between the measurements and the model.

A major difficulty in modeling dynamical systems is the issue of memory. It is known that even for a Markovian system, the corresponding reduced-dimensional system could be non-Markovian [27, 28]. In general, there are two main ways of introducing memory effects in neural networks. First, a simple workaround for feedforward neural networks (FNN) is to introduce time delayed states in the inputs [11]. However, the drawback is that this could potentially lead to an unnecessarily large number of parameters [29]. To mitigate this, Bakker [25] considered following Broomhead and King [30] in reducing the dimension of the delay vector using weighted principal component analysis (PCA). The second approach uses output or hidden units as additional feedback. As an example, Elman’s network [29] is a recurrent neural network (RNN) that incorporates memory in a dynamic fashion.

Miyoshi et al. [31] demonstrated that recurrent RBF networks have the ability to reconstruct simple chaotic dynamics. Sato and Nagaya [32] showed that evolutionary algorithms can be used to train recurrent neural networks to capture the Lorenz system. Bailer-Jones et al. [33] used a standard RNN to predict the time derivative in discrete or continuous form for simple dynamical systems; this can be considered an RNN extension to Tsung’s phase space learning [16]. Wang et al. [34] proposed a framework combining POD for dimension reduction and long-short-term memory (LSTM) recurrent neural networks and applied it to a fluid dynamic system.

We limit ourselves to feedforward neural networks, since there are still many unanswered questions about modeling

dynamical systems even in this simplest form. It is known that time delayed FNNs closely resemble simple RNNs trained with teacher forcing [35]. Further, RNNs are not easy to train since standard training algorithms (e.g., back propagation through time [36]) are likely to introduce stronger overfitting than FNN due to vanishing gradients [35]. Recently, sparse regression (SINDy) [3, 4] has gained popularity as a tool for data-driven modeling. The idea is to search for a sparse representation of a linear combination of functions selected from a library. In this work, we will compare it with FNN-based models and highlight some differences.

The paper is organized as follows: the problem description is provided in Section 2 and the mathematical formulation of standard and Jacobian-regularized FNNs is presented in Section 3. Results and discussion are presented in Section 4. We first present a comparison with SINDy for simple dynamical systems. Then, we highlight the importance of stabilization to control the global error of predicted trajectory and the impact of Jacobian regularization. Finally, we apply the model in a nonlinear PDE system where a low dimensional attractor is not realized and discuss the limitations of black-box modeling of dynamical system and propose data augmentation as remedies. Conclusions are drawn in Section 5.

2. Problem Description

Consider a dynamical system in Euclidean space $\mathbb{R}^M$ which is described by a continuously differentiable function $\phi : \mathbb{R} \times \mathbb{R}^M \mapsto \mathbb{R}^M$, where $\phi(t, \mathbf{x}) = \phi_t(\mathbf{x})$. The state $\mathbf{x} \in \mathbb{R}^M$ satisfies the composition relation $\phi_{t+s} = \phi_t \circ \phi_s$ for $t, s \in \mathbb{R}$ and $\phi_0 : \mathbb{R}^M \mapsto \mathbb{R}^M$ is the identity function. ϕ_t is the t map of the flow described by a vector function $\mathbf{F}_c : \mathbb{R}^M \mapsto \mathbb{R}^M$ as

$$\dot{\mathbf{x}} = \mathbf{F}_c(\mathbf{x}). \quad (1)$$

Similarly, one can define a discrete dynamical system induced by the above smooth dynamical system by considering a constant time step $\Delta t \in \mathbb{R}$ and a state transition map $F_d(\mathbf{x}) = \phi_{\Delta t}(\mathbf{x}) : \mathbb{R}^M \mapsto \mathbb{R}^M$ such that

$$\mathbf{x}^{n+1} = F_d(\mathbf{x}^n). \quad (2)$$

Equivalently, one can rewrite the above system as

$$\frac{(\mathbf{x}^{n+1} - \mathbf{x}^n)}{\Delta t} = \frac{(F_d(\mathbf{x}^n) - \mathbf{x}^n)}{\Delta t} \triangleq \mathbf{F}_r(\mathbf{x}^n), \quad (3)$$

where $\mathbf{F}_r : \mathbb{R}^M \mapsto \mathbb{R}^M$ resembles a first order solution [33] to $\mathbf{F}_c$.

Our goal is to find an approximation to the dynamics, either (i) in a discrete sense F_d , given the data $\mathcal{D} = \{\mathbf{x}^i\}_{i=0}^{N-1}$ uniformly sampled from a trajectory given initial condition $\mathbf{x}^0 = \mathbf{x}(t=0)$ or (ii) in a continuous sense $\mathbf{F}_c$, given the data $\mathcal{D} = \{(\mathbf{x}^i, \dot{\mathbf{x}}^i)\}_{i=0}^{N-1}$, where N is the number of data points. It must be mentioned that—as highlighted in the result section—data does not have to be collected on the same trajectory.

Depending on the way one defines the training and testing set, two types of problems are considered in the current work.

- (1) Prediction of a certain trajectory starting from an initial condition that is different from the training trajectories
- (2) Prediction of the future trajectory given the past information of the trajectory as training data

Conservatively speaking, the success of tackling the first of the above problems requires the trajectories in the training data to be a representative of the distribution in the region of interest, which may or may not be feasible depending on how informative the data is. In the context of modeling dynamical systems, it is often implied in previous literature [22] that the initial condition of unseen testing data is not far away from the training data. The second problem can also be difficult since it will challenge the effectiveness of the model as past information might not be sufficient for the model to be predictive on unseen data. Again, it is often implied in the previous works [20, 34, 37] that successful predictions are often accompanied by an underlying low dimensional attractor so the past states as training data can be collected until it becomes a representative of the future.

3. Mathematical Modeling Framework

In this section, we first define performance metrics of the approximation to the dynamics f ; then, introduce the standard FNN model and the Jacobian-regularized FNN model. Finally, techniques to mitigate overfitting are described.

3.1. Definitions of Error Metrics. To measure the prediction error for each sample in an a priori sense (i.e., given exact $\mathbf{x}^i$), we define the local error vector $\xi_{\text{local}}^i \in \mathbb{R}^M$ for the i -th sample $(\mathbf{x}^i, \mathbf{y}^i)$ as

$$\xi_{\text{local}}^i = \mathbf{y}^i - f(\mathbf{x}^i), \quad (4)$$

where $\mathbf{y}^i \in \mathbb{R}^M$ is the i -th target to learn from the i -th feature $\mathbf{x}^i \in \mathbb{R}^M$. For example, the feature is the state vector at the i -th step, $\mathbf{x}^i$ and the target can be $\mathbf{x}^{i+1}$ for discrete dynamical system or $\dot{\mathbf{x}}^i$ for continuous dynamical system.

Then, we can define local error at the i -th sample by

$$e_{\text{local}}^i = \left\| \xi_{\text{local}}^i \right\|_2 = \left\| \mathbf{y}^i - f(\mathbf{x}^i) \right\|_2, \quad (5)$$

where $\|\cdot\|_2 : \mathbb{R}^M \mapsto [0, +\infty)$ is the vector 2-norm, i.e., l^2 norm, and $|\cdot| : \mathbb{R} \mapsto [0, \infty)$ is the absolute value.

We can further define the local error of the i -th sample for the j -th component as shown by

$$e_{\text{local},j}^i = \left| y_j^i - f(\mathbf{x}^i)_j \right|. \quad (6)$$

The local error assumes that the i -th input feature $\mathbf{x}_i$ is predicted accurately. On the other hand, the global error

vector is defined by equation (7), in which $\hat{\mathbf{x}}^i$ is obtained by *iterative prediction*, i.e., a posteriori *evaluation*, at the i -th step from an initial condition through either time integration or transition function as a discrete map. That is, $\hat{\mathbf{x}}^i$ is obtained from $f(\hat{\mathbf{x}}^{i-1})$ in a recursive sense as follows:

$$\xi_{\text{global}}^i = \mathbf{x}^i - \hat{\mathbf{x}}^i. \quad (7)$$

Similarly, the global error is defined by

$$e_{\text{global}}^i = \left\| \xi_{\text{global}}^i \right\|_2 = \left\| \mathbf{x}^i - \hat{\mathbf{x}}^i \right\|_2, \quad (8)$$

and for the j -th component specifically by

$$e_{\text{global},j}^i = \left| x_j^i - \hat{x}_j^i \right|. \quad (9)$$

Further, to obtain a holistic view of the model performance in feature space, if F_d or F_c is known, either in the continuous or discrete case, we can define stepwise error as

$$e_{\text{stepwise}}(\mathbf{x}) = \left\| F_{c,d}(\mathbf{x}) - f(\mathbf{x}) \right\|_2. \quad (10)$$

Note that e_{stepwise} is not restricted by the training or testing trajectory, but it can be evaluated arbitrarily in the region of interest.

Finally, we consider the uniform averaged coefficient of determination R^2 as a scalar metric for measuring regression performance

$$R^2 = \frac{1}{M} \sum_{j=0}^{M-1} R_j^2, \quad (11)$$

where R_j^2 is given by

$$R_j^2 = R^2(y_j, \tilde{y}_j) = 1 - \frac{\sum_{i=0}^{n_{\text{sample}}-1} (y_j^i - \tilde{y}_j^i)^2}{\sum_{i=0}^{n_{\text{sample}}-1} (y_j^i - \bar{y}_j)^2}, \quad (12)$$

where n_{sample} is the number of samples in the validation data, $\bar{y}_j = 1/n_{\text{sample}} \sum_{i=0}^{n_{\text{sample}}-1} y_j^i$ and $\tilde{y}^i = f(\mathbf{x}^i)$ is the prediction of f based on the i -th feature $\mathbf{x}^i$.

3.2. Feedforward Neural Network Model

3.2.1. Basic Model: Densely Connected Feedforward Neural Network. The basic model approximates F_c in equation (1) for the continuous case and F_r in equation (3) in the discrete case using a feedforward neural network. The existence of an arbitrarily accurate feedforward neural network approximation to any Borel measurable function given the enough number of hidden units is guaranteed from the property of the universal approximator [38]. It should be noted that our basic model is related to Tsung's phase-space-learning model [16]. If the Markovian assumption is adopted, the

training feature matrix snapshots X and training target matrix snapshots Y are as follows:

$$X = \begin{bmatrix} x_1^0 & x_2^0 & x_3^0 & \dots & x_M^0 \\ x_1^1 & x_2^1 & x_3^1 & \dots & x_M^1 \\ \dots & \dots & \dots & \dots & \dots \\ x_1^{N-1} & x_2^{N-1} & x_3^{N-1} & \dots & x_M^{N-1} \end{bmatrix} \in \mathbb{R}^{N \times M}, \quad (13)$$

and

$$Y = \begin{bmatrix} y_1^0 & y_2^0 & y_3^0 & \dots & y_M^0 \\ y_1^1 & y_2^1 & y_3^1 & \dots & y_M^1 \\ \dots & \dots & \dots & \dots & \dots \\ y_1^{N-1} & y_2^{N-1} & y_3^{N-1} & \dots & y_M^{N-1} \end{bmatrix} \in \mathbb{R}^{N \times M}, \quad (14)$$

where M is the dimension of the state, N is the total number of snapshots of training data, learning target Y is the time derivative, and the subscript stands for the index of the component. Note that each component of the feature and target are normalized to zero mean and unit variance for better training performance in the neural network.

By generally constructing a densely connected feedforward neural network $f(\cdot): \mathbb{R}^M \mapsto \mathbb{R}^M$ with $L-1$ hidden layers and output layer as linear, the following recursive expression is defined for each hidden layer:

$$\eta^l = \sigma_l(W_l \eta^{l-1} + b_l), \quad l = 1, \dots, L-1, \quad (15)$$

where η^0 stands for the input of the neural network $\mathbf{x}$, $\eta^l \in \mathbb{R}^{n_l}$, $n_l \in \mathbb{N}^+$ is the number of hidden units in layer l , and σ_l is the activation function of layer l . Note that the output layer is linear, i.e., $\sigma_L(x) = x$:

$$f(\mathbf{x}; \mathbf{W}^L, \mathbf{b}^L) = \eta^L = W_L \eta^{L-1} + b_L, \quad (16)$$

where the parameters of the neural network are $\mathbf{W}^L = \{W_i\}_{i=1,\dots,L}$, $\mathbf{b}^L = \{b_i\}_{i=1,\dots,L}$.

For example, if we consider using two hidden layers where $L=3$ and the number of hidden units are the same, the full expression for the neural network model is given by

$$\hat{\mathbf{y}} = f(\mathbf{x}; \mathbf{W}, \mathbf{b}) = W_3 \sigma(W_2 \sigma(W_1 \mathbf{x} + b_1) + b_2) + b_3, \quad (17)$$

where $\mathbf{x} \in \mathbb{R}^M$ is the state of the dynamical system, i.e., the input to the neural network and $\hat{\mathbf{y}} \in \mathbb{R}^M$ is the modeling target, i.e., the output of neural network. $\sigma(\cdot): \mathbb{R} \mapsto \mathbb{R}$ is a nonlinear activation function. $W_1 \in \mathbb{R}^{n_h \times M}$, $W_2 \in \mathbb{R}^{n_h \times n_h}$, $W_3 \in \mathbb{R}^{M \times n_h}$, $b_1 \in \mathbb{R}^{n_h}$, $b_2 \in \mathbb{R}^{n_h}$, $b_3 \in \mathbb{R}^M$. Sets of weights and biases are $\mathbf{W}^3 = \{W_1, W_2, W_3\}$ and $\mathbf{b}^3 = \{b_1, b_2, b_3\}$. The problem is to find the set of parameters of $\mathbf{W}^3$ and $\mathbf{b}^3$ that result in the best approximation of the underlying ground truth ($\mathbf{F}_e$, $\mathbf{F}_d$, or $\mathbf{F}_r$). Under the framework of statistical learning, it is standard to perform empirical risk minimization (ERM) with mean-square-error loss. The set

up and parameters corresponding to the desired solution $f(\mathbf{x}, \mathbf{W}^*, \mathbf{b}^*)$ can be written as

$$\mathbf{W}^*, \mathbf{b}^* = \arg \min_{\mathbf{W}^3, \mathbf{b}^3} \frac{1}{|I_{\text{train}}|} \sum_{i \in I_{\text{train}}} \|f(\mathbf{x}^i; \mathbf{W}^3, \mathbf{b}^3) - \mathbf{y}^i\|_2^2, \quad (18)$$

where I_{train} is the index set of training data, and $\mathbf{x}^i$ and $\mathbf{y}^i$ correspond to the i -th feature target pair.

To deal with the high dimensionality of the optimization problem, we employ Adam [39], a gradient-based algorithm, which is essentially a mixture between momentum acceleration and rescaling parameters. The weights are initialized using a truncated normal distribution to potentially avoid saturation and use the automatic differentiation (AD) provided by Tensorflow [40] to compute the gradients. The neural network model is implemented in Python using the Tensorflow library [40]. Due to the nonconvex nature of equation (18), for such a high degree of freedom of parameters, one can only afford to find a local minimum. In practice, however, a good local minimum is usually satisfactory [35]. Hyperparameters considered in the current work for the basic model are the number of units for each hidden layer n_h and activation function $\sigma(\cdot)$.

Model selection for neural networks is an active research area [41–43]. Well-known methods involve grid search/random search [41]/Tree of Parzen Estimators (TPE) [42]/Bayesian optimization [43] with cross validation. We pursue the following trial-error strategy:

- (1) Given the number of training points, computing the number of equations to satisfy if the network overfits all the training data
- (2) Pick a neural network with uniform hidden layer structure to overfit the training data with the number of parameters in the network no more than 10% to 50% of the number in step 1
- (3) Keep reducing the size of neural network by either decreasing the hidden units or number of layers until the training and validation error are roughly the same order
- (4) For the choice of other hyperparameters, we simply perform grid search

3.2.2. Jacobian Regularized Model. In standard FNNs, minimizing mean-squared-error on the training data only guarantees model performance in terms of the local training error. It does not guarantee the reconstruction of even training trajectory in the *a posteriori* sense.

Here, we take a closer look at the error propagation in a dynamical system for the FNN model when evaluated in an iterative fashion, i.e., *a posteriori* sense. Without any loss of generality, considering the discrete case, after we obtain the model f , we can predict $\hat{\mathbf{x}}^{i+1}$ given $\hat{\mathbf{x}}^i$

$$\hat{\mathbf{x}}^{i+1} = f(\hat{\mathbf{x}}^i). \quad (19)$$

FIGURE 1: VDP case. (a) Distribution of training/testing sets. (b) Learning rate.

Moreover, given $\mathbf{F}_d$, we can find the $\xi_{\text{global}}^{i+1}$ given $\hat{\mathbf{x}}^i$ and ξ_{global}^i as follows

$$\xi_{\text{global}}^{i+1} = \mathbf{x}^{i+1} - \hat{\mathbf{x}}^{i+1} = \mathbf{F}_d(\mathbf{x}^i) - f(\hat{\mathbf{x}}^i) = \mathbf{F}_d(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i) - f(\hat{\mathbf{x}}^i), \quad (20)$$

$$\begin{aligned} \xi_{\text{global}}^{i+1} &= \mathbf{F}_d(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i) - f(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i) \\ &\quad + f(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i) - f(\hat{\mathbf{x}}^i). \end{aligned} \quad (21)$$

Consider a Taylor expansion of $f(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i)$ about $\hat{\mathbf{x}}^i$, we have

$$\begin{aligned} \xi_{\text{global}}^{i+1} &= \mathbf{F}_d(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i) - f(\hat{\mathbf{x}}^i + \xi_{\text{global}}^i) + \left. \frac{\partial f}{\partial \mathbf{x}} \right|_{\mathbf{x}=\hat{\mathbf{x}}^i} \xi_{\text{global}}^i \\ &\quad + \frac{1}{2} \xi_{\text{global}}^{i,T} H \xi_{\text{global}}^i + \dots, \end{aligned} \quad (22)$$

where H is the Hessian matrix evaluated at some point between $\hat{\mathbf{x}}^i$ and $\hat{\mathbf{x}}^i + \xi_{\text{global}}^i$.

Assuming $\|\xi_{\text{global}}^i\|_2 \ll 1$, $\|H\|_2$ is bounded, and the high order terms are negligible compared to the Jacobian term, we have

$$\begin{aligned} e_{\text{global}}^{i+1} &\leq e_{\text{local}}^i + \left\| \frac{\partial f}{\partial \mathbf{x}} \right\|_{\mathbf{x}=\hat{\mathbf{x}}^i} \|e_{\text{global}}^i\|_2 + o(e_{\text{global}}^i) \\ &\leq e_{\text{local}}^i + \left\| \frac{\partial f}{\partial \mathbf{x}} \right\|_{\mathbf{x}=\hat{\mathbf{x}}^i} \|e_{\text{global}}^i\|_F + o(e_{\text{global}}^i). \end{aligned} \quad (23)$$

TABLE 1: Hyperparameter configuration of the basic model: VDP case.

Layer structure	Activation function	Loss function	Optimizer	Learning rate
2-8-8-2	Swish	MSE	Adam	0.002

Similarly, in the continuous case, we have

$$\begin{aligned} e_{\text{global}}^{i+1} &\leq e_{\text{global}}^i + \int_{i\Delta t}^{(i+1)\Delta t} e_{\text{local}}^{\tau/\Delta t} d\tau + \int_{i\Delta t}^{(i+1)\Delta t} \left\| \frac{\partial f}{\partial \mathbf{x}} \right\|_{\mathbf{x}=\hat{\mathbf{x}}^i} e_{\text{global}}^{\tau/\Delta t} d\tau \\ &\quad + o(e_{\text{global}}^i \Delta t), \end{aligned} \quad (24)$$

$$\begin{aligned} e_{\text{global}}^{i+1} &\leq \left(1 + \Delta t \left\| \frac{\partial f}{\partial \mathbf{x}} \right\|_{\mathbf{x}=\hat{\mathbf{x}}^i} \right) e_{\text{global}}^i + \int_{i\Delta t}^{(i+1)\Delta t} e_{\text{local}}^{\tau/\Delta t} d\tau \\ &\quad + o(e_{\text{global}}^i \Delta t). \end{aligned} \quad (25)$$

The right hand sides of equations (23) and (25) contain contributions from the global error and accumulation of local error. Optimization as in equation (18) can minimize the latter term, but not necessarily the former. This suggests that manipulating the eigenspectrum of the Jacobian might be beneficial for stabilization by suppressing the growth of the error. Due to the simplicity of computing the Frobenius norm compared to the 2-norm, we consider penalizing the Frobenius norm of the Jacobian of the neural network model. In the context of improving generalization performance of input-output neural network models, similar regularization has been also proposed by Rifai et al. [44]. It should be noted that our purpose is to achieve better error dynamics in a temporal sense,

FIGURE 2: A priori and a posteriori result of the basic model on VDP case. (a, b) Training data. (c, d) Testing data. (a, c) A priori. (b, d) A posteriori.

which differs from the generalization goal in deep learning. Thus, one may seek a locally optimal solution that can suppress the growth in global error while minimizing the local error.

The regularized loss function inspired from the above discussion is thus

$$\mathbf{W}^*, \mathbf{b}^* = \arg \min_{\mathbf{W}, \mathbf{b}} \frac{1}{N_{\text{train}}} \sum_{i \in I_{\text{train}}} \|f(\mathbf{x}^i; \mathbf{W}, \mathbf{b}) - \mathbf{y}^i\|_2^2 + \lambda \|J(\mathbf{x}^i; \mathbf{W}, \mathbf{b})\|_F^2, \quad (26)$$

where J is the Jacobian of the neural network output with respect to the input and λ is a hyperparameter. On one hand, it should be noted that regularizing the Frobenius norm of the eigenspectrum of the Jacobian indirectly suppresses the magnitude of the eigenvalue of the Jacobian.

On the other hand, excessive weighting on the magnitude of the eigenvalue would lead to less weighting on local error, which might result in an undesirably large local error. Thus, λ should be set as a relatively small value without strongly impacting the model performance in an a priori sense.

3.3. Reducing Overfitting. Overfitting is a common issue in the training of machine learning models, and it arises when models tend to memorize the training data instead of generalizing true functional relations. In neural networks, overfitting can occur from poor local minima and is partially due to the unavoidable nonconvexity of an artificial neural network. Overfitting cannot be completely eliminated for most problems, given the NP-hard nature of the problem. Generally, overfitting can be controlled by three kinds of regularization techniques. The first follows the Occam's razor

FIGURE 3: Variation of local error and global error: basic model on VDP case. (a, b) Local error. (c, d) Global error. (e, f) Maximal singular value of the Jacobian evaluated in a priori. (a, c, e) Training data. (b, d, f) Testing data.

principle, e.g., L1 sparsity regularization [3]. However, there is no guarantee that Occam's razor is appropriate for all cases, and finding the optimal sparsity level is often iterative. The second is to smooth the function, e.g., using weight decay [35]. The third type is especially suitable in iterative learning,

e.g., early stopping, which is a widely used strategy in the deep learning community [35]. In this work, we found validation-based early stopping to be sufficient. We split the data further into pure training and validation sets, and then monitor overfitting by measuring R^2 .

FIGURE 4: Stepwise error contour of the basic model on VDP case. (a) Training data. (b) Testing data. Blue dot: ground truth. Orange dot: prediction. White arrow: direction of target vector of ground truth. Red arrow: direction of target vector of prediction.

FIGURE 5: A priori and a posteriori result of SINDy on VDP case. (a, b) Training data. (c, d) Testing data. (a, c) A priori. (b, d) A posteriori.

FIGURE 6: Stepwise error contour of SINDy on VDP case. (a) Training data. (b) Testing data. Blue dot: ground truth. Orange dot: prediction. White arrow: direction of target vector of ground truth. Red arrow: direction of target vector of prediction.

FIGURE 7: Nonrational nonpolynomial case. (a) Data distribution of training/testing feature. (b) Learning curve.

4. Results and Discussion

Given sequential training data, the capability of the basic FNN is first evaluated in two-dimensional dynamical systems with polynomial nonlinearities in Section 4.1 and nonpolynomial nonrational dynamics in Section 4.2. The basic model is compared with SINDy [3], which is a method that directly aims at learning functional models using L_1 sparse regression on a dictionary of candidate basis functions. In Section 4.3, we demonstrate that the basic model performs better than SINDy on the problem of incompressible flow behind a cylinder, in spite of the explicit addition of quadratic terms to the dictionary. In

TABLE 2: Hyperparameter configuration of the basic model: nonrational nonpolynomial case.

Layer structure	Activation function	Loss function	Optimizer	Learning rate
2-8-8-2	elu	MSE	Adam	0.005

addition, the local error is found to be strongly correlated with the maximal singular value of the Jacobian, thus serving as an inspiration for Jacobian regularization. In Section 4.4, we demonstrate the stabilizing aspect of

FIGURE 8: A priori and a posteriori result of the basic model on nonrational nonpolynomial case. (a, b) Training data. (c, d) Testing data. (a, c) A priori. (b, d) A posteriori.

Jacobian regularization for the problem of laminar wake behind a cylinder, where the system exhibits a low dimensional attractor. In Section 4.5, we assess the ability of our regularized FNN model to approximate a dynamically evolving high-dimensional buoyancy-driven mixing flow system that is a characteristic of flow physics driven by instabilities. The results show that, for systems that do not exhibit a low dimensional attractor, it is difficult for a black-box model to have satisfactory long-time prediction capabilities. In Section 4.6, we show that predictive properties can be improved by data augmentation in the state space of interest.

4.1. 2D Polynomial System: Van der Pol Oscillator. The first order forward discretized scheme of the Van der Pol (VDP) system is given by

$$\begin{pmatrix} x_1^{n+1} \\ x_2^{n+1} \end{pmatrix} = \begin{pmatrix} x_1^n \\ x_2^n \end{pmatrix} + \Delta t \begin{pmatrix} x_2^n \\ (\mu(1-x_1^n)x_2^n - x_1^n) \end{pmatrix}, \quad (27)$$

where $\Delta t = 0.1$ and $\mu = 2.0$. The modeling target is

$$\mathbf{y}^n = \begin{pmatrix} \frac{(x_1^{n+1} - x_1^n)}{\Delta t} \\ \frac{(x_2^{n+1} - x_2^n)}{\Delta t} \end{pmatrix} = \begin{pmatrix} x_2^n \\ (\mu(1-x_1^n)x_2^n - x_1^n) \end{pmatrix} = F_r(\mathbf{x}^n). \quad (28)$$

Our goal is to reproduce the dynamics governed by F_r from on data collected from a single trajectory. The loss

FIGURE 9: Variation of local error and global error: basic model on nonrational, nonpolynomial case. (a, b) Local error. (c, d) Global error. (e, f) Maximal singular value of the Jacobian evaluated in a priori. (a, c, e) Training data. (b, d, f) Testing data.

function of the basic model in equation (18) is optimized using training data from a single trajectory, containing 399 data points. Test data containing 599 points is generated using a different initial condition. The data distribution of the training and testing features is shown in Figure 1. Note that initially, a few test points (orange)

are away from the training data (blue) which require the model to perform extrapolation. Configuration of hyperparameters is shown in Table 1. Data is normalized to zero mean and unit standard deviation for each component. We use minibatch training with batch size=64 and 80,000 epochs. The basic model consists of two hidden

FIGURE 10: Stepwise error contour of the basic model on nonrational nonpolynomial case. (a) training data. (b) testing data. Blue dot: ground truth. Orange dot: prediction. White arrow: direction of target vector of ground truth. Red arrow: direction of target vector of prediction.

FIGURE 11: A priori and a posteriori result of SINDy on nonrational nonpolynomial case. (a, b) Training data. (c, d) Testing data. (a, c) A priori. (b, d) A posteriori.

FIGURE 12: Stepwise error contour of SINDy on nonrational nonpolynomial case. (a) Training data. (b) Testing data. Blue dot: ground truth. Orange dot: prediction. White arrow: direction of target vector of ground truth. Red arrow: direction of target vector of prediction.

layers with each layer containing 8 hidden units. Two hidden layers are accompanied by Swish nonlinear activation as $\sigma(x) = x \cdot \text{sigmoid}(\beta x)$, where in practice, β is fixed as unity [45]. The output layer is linear. Randomly 20% of training data is used as a validation set, and we monitor the performance on the validation set as a warning of overfitting. In Figure 1, the learning curve suggests that the model is well-trained and overfitting is not observed.

Results of a priori and a posteriori prediction are shown in Figure 2. The basic model predicts the $\mathbf{F}_r$ at each training point very well a priori, but slight phase lag is observed a posteriori in testing, which originates from the extrapolation of the testing data initially.

The variation of the local and global error together with the maximal singular value of the Jacobian is shown in Figure 3. For training data, e_{local} is observed to be relatively uniform as expected, since the objective optimized is MSE uniformly across all training data points. For testing data, e_{local} exhibits peak values near the beginning of the trajectory as expected, since the first few points are far away from the training data shown in Figure 1. Moreover, it is interesting to observe that in Figure 3, the peak of the temporal history of local/global error shows a strong correlation with the maximal singular value of the Jacobian.

Stepwise error contours are displayed in Figure 4. The region of large error close to red (implying the difference of the stepwise vector between neural network prediction and ground truth is large) is located near the corner of figure, where there is a dearth of training points. The model performs well near the training points as expected. In this case, since testing data is not very far away from the training data, good performance of extrapolation can be expected. However, we would like to note that there is a moderate amount of error associated with the vector direction in Figure 4 not only at the corners but also near the origin. This implies that a feedforward neural network can generalize to some extent, but with no guarantees, even in regions

enclosed by training data. The results also confirm that the known result for a dynamical system with an attractor, the neural network can reproduce the dynamics near the attractor [13, 25, 33, 37, 46].

With the prior knowledge that the system is polynomial in nature, one can use polynomial basis functions to extract the ground truth. To illustrate this, results obtained from SINDy [3] are shown in Figure 5, with threshold parameter as 2×10^{-4} , maximal polynomial order as 3, and no validation data set considered. As displayed in Figure 6, the excellent result of SINDy shows the advantage of finding the global features where the parameters obtained are not restricted to the scope of training data since the ground truth is governed by sparse polynomials.

4.2. 2D Nonpolynomial System: A Nonrational Nonpolynomial Oscillator. The success of SINDy is a consequence of the fact that the underlying system can be represented as a sparse vector in a predefined basis library such as that consisting of polynomial or rational functions [4]. Here, we choose a different case: a nonrational nonpolynomial oscillator with $\Delta t = 0.004$

$$\begin{pmatrix} x_1^{n+1} \\ x_2^{n+1} \end{pmatrix} = \begin{pmatrix} x_1^n \\ x_2^n \end{pmatrix} + \Delta t \left(2.5 - 100 \frac{x_1^n x_2^n}{1 + (x_2^n/0.52)^4} - 200 \frac{x_1^n x_2^n}{1 + (x_2^n/0.52)^4} + 9.2 - 2.3x_2^n - 1.28|x_2^n|^{3/2} \right). \quad (29)$$

Here, the basic model in equation (18) is optimized using 1199 data points of a single trajectory. Testing data contains 1799 points. Randomly, 20% of the training data is taken as the validation set, but also included in later evaluation. The feature distribution in phase space is shown in Figure 7. Hyperparameters are listed in Table 2 and 128

FIGURE 13: Flow in a cylinder wake. (a) Data distribution of x_1 vs. x_2 . (b) Data distribution of x_1 and x_3 . (c) Data distribution of x_2 and x_3 . (d) Learning curve.

minibatches and 20,000 epochs are used. The training error and validation error is also shown in Figure 7.

Results for a priori and a posteriori performance on training and testing data are shown in Figure 8. The training trajectory is perfectly reconstructed while the predictions show slight deviation.

The distribution of the local and global error is shown in Figure 9. Again, we observe that maximal local/global error correlates with the peaks of the maximal singular value of the Jacobian. It is interesting to note that the highest local testing error occurs at the peak of the maximal singular value of the Jacobian, instead of at the points close to the initial condition.

The error contour in Figure 10 shows that the stepwise error around the training trajectory is below 0.1. It is important to note that model performance deteriorates at places far

TABLE 3: Hyperparameter configuration of the basic model: flow in a cylinder wake.

Layer Structure	Activation Function	Loss Function	Optimizer	Learning rate
2-20-20-2	elu	MSE	Adam	0.001

away from the training trajectory, especially at the right corner shown in Figure 10.

For SINDy, the polynomial order is set to three and threshold as 2×10^{-4} . A priori and a posteriori validation for training and testing is shown in Figure 11. Correspondingly, the stepwise error contour displayed in Figure 12 shows the misfit for the region of interests ranging from -1 to 3 for both two components. Because there is no sparsity

FIGURE 14: A posteriori comparison for flow in a cylinder wake between the basic model and SINDy. Red: first POD coefficient. Green: second POD coefficient. Blue: shift-mode coefficient. (a, b) Basic model. (c, d) SINDy. (a, c) Training data. (b, d) Testing data.

in polynomial basis in this case, it is expected that SINDy cannot reconstruct the dynamics correctly and would perform worse than the basic model of FNN. The implication is that for strongly nonpolynomial systems, neural networks are far more flexible compared to SINDy.

4.3. Nonlinear PDE System: Flow behind a Cylinder. In this section, we compare the basic model with SINDy in reconstructing the flow in a cylinder wake. The data is from Brunton et al. [3] which comes from an immersed boundary method solution [47] of the 2D incompressible N-S equations with $Re=100$ based on the cylinder diameter. The computational domain consists of a nonuniform grid with near-wall refinement. The inlet condition is uniform flow and the outlet is a convective boundary condition to allow the vorticity to exit the domain freely. Testing data is generated as a temporal extension of states that lie on a limit cycle at the boundary of training data, which indicates that this is not an extrapolation task. To work with such a high-dimensional nonlinear PDE system, we use the coefficients of two POD modes [23] and one “shift mode”, which

represents the shift of short-term averaged flow away from the POD space of the first two harmonic modes to reduce the spatial dimension. More details on POD and “shift-modes” are provided in references [23, 48]. Training and testing data is the same as in Brunton et al. [3] where the first 2999 snapshots in time are used for training, and a later 2994 snapshots used for testing. A random 10% of training snapshots is considered as validation set but also included in later evaluation. The distribution of training data and testing data is shown in Figure 13.

Hyperparameters of the basic model are shown in Table 3 with 40,000 epochs. For SINDy, the hyperparameters are the same as in the previous work [3]. As shown in Figure 14, for training data, SINDy reconstructs a smaller growth rate of oscillating behavior, while the basic model accurately reconstructs both the shift mode and two POD modes. For testing data, SINDy contains an observable phase lag for the time period concerned, while the basic model achieves an almost perfect match. This implies that the model obtained from SINDy, although much easier to interpret than neural network, is not the best model for this dynamical system in

FIGURE 15: A posteriori comparison between basic and regularized models for cylinder wake. Red: first POD coefficient. Green: second POD coefficient. Blue: shift-mode coefficient. (a, b) Basic model. (c, d) Regularized model. (a, c) Training data. (b, d) Testing data.

terms of accuracy. However, we note that from the data distribution in Figure 13, the basic model performs as expected, as the training data covers the attractor well.

4.4. Stabilizing the Neural Network with Jacobian Regularization. Due to the nonconvexity of the optimization problem that arises in the solution of the basic model in equation (18), employing a stochastic gradient-descent type method might lead to a solution corresponding to a local minimum, which is often undesirable and difficult to avoid. Most works in the field of deep learning for feedforward neural networks focus on decreasing the impact of poor local minima to promote generalizability. However, in the context of modeling a dynamical system, as it is often assumed that the trajectory of interest is stable with respect to small disturbances [25], the model should be able to approximately reconstruct the training trajectory in the presence of local errors that arise at each step. This would require regularizing instabilities that could arise in a posteriori prediction. To have meaningful comparisons, random number seeds are fixed for initialization of weights and training data shuffling. Nevertheless, we observe that in some cases, for example, in

the previous case of the cylinder wake, an inappropriate choice of neural network configuration of the basic model, e.g., number of hidden units and type of activation function, can potentially lead to instability in a posteriori evaluation. Such instabilities may materialize even while reconstructing the training trajectory, while the corresponding a priori prediction is almost perfect. Previous work [16] explicitly ensured stability by simply adding more adjacent trajectories. Here, we take a different approach by adding a Jacobian regularization term in the cost function in equation (26).

In our numerical experiments, with a certain fixed random seed, it is observed that, when the layer structure is 2-20-20-2 with $\tanh$ as activation function instead of elu , the basic model becomes numerically unstable after 2000 steps for training data which is displayed in Figure 15. Similar numerical instability is also observed in testing evaluations. However, for the same fixed random seed, the regularized model with $\lambda = 5 \times 10^{-5}$ shows numerically stable results with the same neural network configuration for both training and testing data.

The effectiveness of Jacobian regularization may be attributed to finding a balance between lowering the prediction

FIGURE 16: Flow in a cylinder wake. Comparison of eigenvalue of the Jacobian between regularized and basic model on training data. (a, b) Maximal real eigenvalue. (c, d) Maximal imaginary eigenvalue. (a, c) A priori. (b, d) A posteriori.

error, i.e., MSE, and suppressing the sensitivity of the prediction of the future state to the current local error. As shown in Figures 16 and 17, on average, the maximal eigenvalue of the Jacobian is smaller for the regularized model than for the basic model. Furthermore, the distribution of the eigenvalues of the Jacobian is shown in Figure 18 in the form of a linear stability diagram with explicit 5th order Runge-Kutta time integration. It is clear that the model with Jacobian regularization has significantly smaller positive real eigenvalues. Note that, due to the Frobenius norm, negative real eigenvalues are also decreased in magnitude.

4.5. Nonlinear PDE System: Instability-Driven Buoyant Mixing Flow. The test problems thus far have served to assess the performance of the basic and Jacobian regularized models on nonlinear dynamical systems that either evolve on or towards an attractor. Such systems, even if high-

dimensional, are amenable for projection onto a lower dimensional subspace, using, for instance, POD techniques. In this section, we consider the Boussinesq buoyant mixing flow [49, 50], also known as the unsteady lock-exchange problem [51] which exhibits strong shear and Kelvin-Helmholtz instability phenomena driven by the temperature gradient. Compared to the cylinder flow that evolves on a low-dimensional attractor approaching a limit cycle, the Boussinesq flow is highly convective and instability driven. Consequently, such a system state cannot be represented by a compact set of POD modes from the spatial-temporal field of nondimensionalized velocity and temperature. Rather, the low-dimensional manifold itself evolves with time. Further, any noise in the initial data can produce unexpected deviations that makes such systems challenging to model, even using equation-driven reduced order models such as POD-Galerkin [51].

FIGURE 17: Flow in a cylinder wake. Comparison of eigenvalue of the Jacobian between regularized and basic model on testing data. (a, b) Maximal real eigenvalue. (c, d) Maximal imaginary eigenvalue. (a, c) A priori. (b, d) A posteriori.

The data set is generated by solving the dimensionless form of the two-dimensional incompressible Boussinesq equations [51], as shown in equation (30) on a rectangular domain that is $0 < x < 8$ and $0 < y < 1$.

$$\frac{\partial u}{\partial x} + \frac{\partial u}{\partial y} = 0, \quad (30a)$$

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} = -\frac{\partial P}{\partial x} + \frac{1}{\text{Re}} \nabla^2 u, \quad (30b)$$

$$\frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} = -\frac{\partial P}{\partial y} + \frac{1}{\text{Re}} \nabla^2 v + \text{Ri} \theta, \quad (30c)$$

$$\frac{\partial \theta}{\partial t} + u \frac{\partial \theta}{\partial x} + v \frac{\partial \theta}{\partial y} = \frac{1}{\text{Re Pr}} \nabla^2 \theta, \quad (30d)$$

where u , v , and θ are the horizontal, vertical velocity, and temperature components, respectively. The dimensionless parameters Re , Ri , and Pr are the Reynolds number, Richardson number, and Prandtl number, respectively, with values chosen as follows: $\text{Re} = 1000$, $\text{Ri} = 4.0$, and $\text{Pr} = 1.0$. These equations are discretized on a 256×33 grid. Initially, fluids at two different temperatures are separated by a vertical line at $x = 4$. The bounding walls are treated as adiabatic with the no-slip condition. A fourth-order compact finite difference scheme is used to compute the derivatives in equation (30). The evolution of the thermal field over the simulation time interval of 32 seconds is shown in Figure 19 and illustrates the highly transient nature of the dynamics. To reduce the dimensionality of the system, POD modes are extracted from the entire data set consisting of 1600 snapshots. The reduced feature set consisting of ten POD weights captures nearly 97%

FIGURE 18: Flow in a cylinder wake. Comparison of linear stability diagram between regularized and basic model. (a) A priori on training data. (b) A posteriori on training data. (c) A priori on testing data. (d) A posteriori on testing data.

of the total energy is used to train the model and predict the trajectory.

For the setup of training and testing, future state prediction is pursued with the first 70% states of the trajectory treated as training data and the rest for testing. For such a system in 10 dimensions, it is observed that the problem of the a posteriori instability in the basic model becomes more pronounced and difficult to avoid. Challenges of numerical instability were observed even for reconstruction for a wide range of network configurations, and thus, results from the basic model are not reported.

The Jacobian regularized model is employed with hyperparameters shown in Table 4, with Figure 20 showing a posteriori evaluation on training data. The reconstruction is successful, but the performance deteriorates on testing data because the trajectory of the system does not exhibit a low dimensional attractor as in the cylinder case. Therefore, the training data is not informative for predictions on the test

set. For a black-box machine learning model, this phenomena can be expected to be more pronounced in high dimensional space due to data scarcity. Specifically, we discuss this problem in the following section.

4.6. Improving Model Predictability by Data Augmentation.

In this section, we consider two scenarios of data augmentation: (i) augmenting the information in the data by spreading training locations randomly following a uniform distribution provided that one has access to F_c or F_d at any desired location; (ii) augmenting the data by assembling several trajectories generated from different initial conditions.

4.6.1. Random Uniform Sampling in Phase Space. Recall that, in the two-dimensional problems in Section 4.1 and Section 4.2, the stepwise error contour shows that local error increases on testing scenarios located far away from the training data which was highly concentrated in a compact region

FIGURE 19: Time evolution of the temperature field for a two-dimensional Buoyant mixing flow.

TABLE 4: Hyperparameter configuration of the Jacobian regularized model for buoyant mixing flow.

Layer structure	Activation function	Loss function	Optimizer	Learning rate	λ
10-20-20-10	Penalized tanh [56]	MSE	Adam	0.001	5e-4

FIGURE 20: A posteriori comparison between prediction of Jacobian regularized model and ground truth for POD coefficient of buoyant mixing flow. Dashed: model. Solid: ground truth.

of phase space. Without any knowledge of system behavior, it is sensible to start with training data from a random uniform distribution in a compact region of phase space corresponding to interesting dynamics. To conduct a thorough stepwise error contour evaluation of the training target in phase space, the VDP system is chosen to illustrate this idea.

Determining the most informative data samples would potentially involve specific knowledge of the underlying system and the models used and is beyond the scope of the current work. Here, we simply consider uniform random sampling in the phase space in a finite domain: $[-3, 3]$ for the first component and $[-5, 5]$ for the second component. We obtained a new set of 399 training data points using random uniform sampling in phase space while retaining the same testing data as in Section 4.1.

The performance of the basic model with the same hyperparameter setting as in Section 4.1 on randomly distributed training data is shown in Figure 21. While the number of data points has not been changed, the contour error of the resulting model decreased significantly compared to training with the same number of data points in a single trajectory, which indicates an improved generalizability with the same amount of training data.

4.6.2. Training with Multiple Trajectories with Random Initialization. Training data can also be augmented by multiple trajectories with different initial conditions. Here, we take

the one-dimensional viscous Burgers equation shown in equation (31) as example.

The initial conditions are generated following a specific energy spectrum [28, 52] shown in equation (32).

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2}, \quad (31)$$

where $x \in [0, 2\pi]$ is a periodic domain discretized using 2048 uniformly distributed grid points, and $t \in [0, 20]$, $\nu = 0.01$.

$$u(x, 0) = \sum_{k=1}^{k_c} \frac{1}{\pi} \sqrt{2AE(k)} \sin(kx + \beta_k), \quad (32)$$

where for each k , β_k is a random number drawn from a uniform distribution on $[-\pi, \pi]$, $E(k) = 5^{-5/3}$ if $1 \leq k \leq 5$, $A = 25$, and $E(k) = k^{-5/3}$ if $k > 5$. Multiple trajectories are generated using different seeds for random numbers to obtain the trajectories of the full-order system. To fully resolve the system as a DNS, equation (31) is solved using a standard pseudospectral method with SSP-RK3 [53] for time stepping. Here, we choose $k_c = 2$. Discrete cosine transformation (DCT) is used to reduce the dimension of the full system to first 4 cosine modes in the system where around 97% of kinetic energy is preserved. For simplicity, we seek a closed Markovian

FIGURE 21: A priori and a posteriori result of the basic model on the VDP system, training with random sampled data distribution. (a) A priori evaluation. (b) A posteriori evaluation. (c) Learning curve. (d) Stepwise error contour. Blue dot: ground truth. Orange dot: prediction. White arrow: direction of target vector of ground truth. Red arrow: direction of target vector of prediction.

reduced-order-system, whereas the underlying dynamics is clearly non-Markovian [28, 54].

Since the first component of the DCT is constant, the remaining components of feature space are shown in Figure 22. The training data is far away from the testing data initially, whereas the data converges at a later stage. This is because of the presence of a spiral fixed point attractor resulting from the viscous dissipative nature of the system. Therefore, if the model is only trained from a single trajectory, it will be very difficult for the model to generalize well in the phase space especially where the state of the system is not near an attractor.

Many dynamical systems in nature exhibit attractors in the asymptotic sense. From the viewpoint of data-driven modeling of such dynamics, data scarcity is encountered

at the start of trajectory where the number of trajectories required to provide enough information to cover the region of interest grows exponentially. Much research on applying neural network-based models for dynamical systems [16, 20, 34] demonstrate problems starting on limit cycles or chaotic attractors in a low-dimensional feature space, where the issue of initial data scarcity is not significant, or can be easily alleviated by a small increase in available data. However, for the purpose of modeling phenomena such as turbulent fluid flow, which can be high dimensional even after dimension reduction, the model would likely fail for long-time prediction due to data scarcity. Such a situation may be realized in regions of phase space where the state has not arrived at the low manifold attractor. Therefore, the training data might not be representative of testing

FIGURE 22: Feature distribution of training/testing data for component x_2, x_3, x_4 . (a) With one training trajectory. (b) With 18 training trajectory. Black line: testing data. Red line: training data.

TABLE 5: Hyperparameter configuration of the basic model: VBE system with four modes.

Layer structure	Activation function	Loss function	Optimizer	Learning rate
4-30-30-30-4	elu	MSE	Adam	0.0005

data which violates the fundamental assumption of a well-posed machine learning problem [55]. Moreover, data scarcity will shrink the region of generalizability of the model as the dimension of the system increases.

A key benefit of using a neural network model is its linear growth in complexity with dimension of the system, in contrast to traditional polynomial regression methods [35]. However, initial data scarcity would limit the generalizability of an ANN in modeling a high dimensional dynamical system that does not exhibit a low dimensional attractor. We believe that this phenomenon of data scarcity observed from this simple nonlinear PDE example also applies to other non-linear dynamical systems.

To alleviate the initial data scarcity issue, a solution is to augment the training data with more trajectories with different random number seeds in generating the initial condition, while keeping the energy spectrum the same across all cases. In this case, we choose 18 such trajectories. Each trajectory contains states of 1000 snapshots equally spaced in time. For testing data, we simply consider one DNS result with an initial condition different from all training trajectories. The corresponding training and testing trajectories are visualized in phase space as shown in Figure 22.

The basic model is trained with hyperparameters in Table 5 and 1000 epochs. The resulting learning curve and a posteriori evaluation are shown in Figure 23. Relatively large discrepancy is observed near the initial condition as the initial data scarcity is not completely eliminated due to limited number of additional trajectories. Increasing the number of additional trajectories may be unaffordable for

very high dimensional systems. Moreover, the result also shows that the error decreases once the trajectory falls on the fixed point attractor. Thus, if the model starts in the low dimensional attractor where the information is well-preserved in the training data, better performance might be expected. This hypothesis is consistent with the previous work [34], where successful prediction of future states starts at the time when the states converge to a low dimensional attractor.

5. Conclusions

This work investigated the modeling of dynamical systems using feedforward neural networks (FNN), with a focus on long time prediction. It was shown that neural networks have advantages over sparse polynomial regression in terms of adaptability, but with a trade-off in training cost and difficulty in extrapolation, which is a natural barrier for almost all supervised learning. From the perspective of global error analysis, and the observation of the strong correlation between the local error and maximal singular value of the Jacobian, we propose the suppression of the Frobenius norm of the Jacobian as regularization. This showed promise in improving the robustness of the basic FNN model given limited data, or when the model has a nonideal architecture, or when the model is unstable. The effectiveness of Jacobian regularization is attributed to finding a balance between lowering the prediction error and suppressing the sensitivity of the prediction of the future state to the current local error. In terms of modeling dynamical systems that do not involve

FIGURE 23: Result of the basic model for one-dimensional VBE turbulence case. (a) Learning curve. (b) A posteriori result with augmented training trajectories. Red: x_1 . Green: x_2 . Blue: x_3 . Cyan: x_4 . Dashed: model prediction. Solid: ground truth.

low-dimensional attractors, limitations of FNNs, and perhaps all local ML methods, was demonstrated in a buoyant mixing flow. Challenges were noted in the example of the reduced-order viscous burgers system, where significant initial data scarcity is present. Augmenting the data either by altering the distribution of training data in phase space or by simply adding multiple trajectories from different initial conditions resulted in improvement of the performance of FNN model to some extent. However, these remedies require a significant amount of additional sampling in phase space, especially for high dimensional systems for the period of time without the apparent low-dimensional attractor which suffers data sparsity from the curse of dimensionality.

Data Availability

Most of the data contained in the current study are synthetic data which can be easily generated through the configuration described. In addition, the cylinder flow data is directly downloadable in the public repo of SINDy [41]: faculty.washington.edu/sbrunton/sparsedynamics.zip. And the instability-driven buoyant mixing flow data were supplied by Professor Balaji Jayaraman under license and so cannot be made freely available. Requests for access to these data should be made to balaji.jayaraman@okstate.edu. Furthermore, all of the data used to support the findings of this study are available from the corresponding author upon request. Ref: [1] Brunton, Steven L., Joshua L. Proctor, and J. Nathan Kutz. "Discovering governing equations from data by sparse identification of nonlinear dynamical systems." *Proceedings of the National Academy of Sciences* 113.15 (2016): 3932–3937.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

The authors acknowledge Professor Balaji Jayaraman for his helpful discussion and comments and for providing the data for the two-dimensional instability-driven Buoyant mixing flow. This work was supported by AFOSR and AFRL under grants FA9550-16-1-0309 and FA9550-17-1-0195.

References

- [1] M. Tanaskovic, L. Fagiano, C. Novara, and M. Morari, "Data-driven control of nonlinear systems: an on-line direct approach," *Automatica*, vol. 75, pp. 1–10, 2017.
- [2] R. H. Shumway and D. S. Stoffer, "Time series analysis and its applications," *Studies in Informatics and Control*, vol. 9, no. 4, pp. 375–376, 2000.
- [3] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [4] N. M. Mangan, S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Inferring biological networks by sparse identification of nonlinear dynamics," *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, vol. 2, no. 1, pp. 52–63, 2016.
- [5] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Largescale video classification with convolutional neural networks," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 1725–1732, Columbus, OH, USA, 2014.
- [6] O. Russakovsky, J. Deng, H. Su et al., "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [7] T. Duriez, S. L. Brunton, and B. R. Noack, *Machine Learning Control – Taming Nonlinear Dynamics and Turbulence*, Springer, 2017.

- [8] K. Hornik, "Approximation capabilities of multilayer feedforward networks," *Neural Networks*, vol. 4, no. 2, pp. 251–257, 1991.
- [9] J. G. Kuschewski, S. Hui, and S. H. Zak, "Application of feedforward neural networks to dynamical system identification and control," *IEEE Transactions on Control Systems Technology*, vol. 1, no. 1, pp. 37–49, 1993.
- [10] K. S. Narendra and K. Parthasarathy, "Identification and control of dynamical systems using neural networks," *IEEE Transactions on Neural Networks*, vol. 1, no. 1, pp. 4–27, 1990.
- [11] K. S. Narendra and K. Parthasarathy, "Neural networks and dynamical systems," *International Journal of Approximate Reasoning*, vol. 6, no. 2, pp. 109–131, 1992.
- [12] M. M. Polycarpou and P. A. Ioannou, "Identification and control of nonlinear systems using neural network models: design and stability analysis," University of Southern California, 1991.
- [13] S. A. Billings, H. B. Jamaluddin, and S. Chen, "Properties of neural networks with applications to modelling nonlinear dynamical systems," *International Journal of Control*, vol. 55, no. 1, pp. 193–224, 1992.
- [14] V. T. S. Elanayar and Y. C. Shin, "Radial basis function neural network for approximation and estimation of nonlinear stochastic dynamic systems," *IEEE Transactions on Neural Networks*, vol. 5, no. 4, pp. 594–603, 1994.
- [15] K. Chakraborty, K. Mehrotra, C. K. Mohan, and S. Ranka, "Forecasting the behavior of multivariate time series using neural networks," *Neural Networks*, vol. 5, no. 6, pp. 961–970, 1992.
- [16] F.-S. Tsung and G. W. Cottrell, "Phase-space learning," in *Advances in Neural Information Processing Systems*, pp. 481–488, MIT Press, 1995.
- [17] T. L. Paez and N. F. Hunter, "Dynamical system modeling via signal reduction and neural network simulation," Sandia National Labs, Albuquerque, NM (United States), 1997.
- [18] T. L. Paez and N. F. Hunter, "Nonlinear system modeling based on experimental data," Technical report, Sandia National Labs., Albuquerque, NM (US); Sandia National Labs., Livermore, CA (US), 2000.
- [19] A. Urbina, N. F. Hunter, and T. L. Paez, "Characterization of nonlinear dynamic systems using artificial neural networks," Technical report, Sandia National Labs, Albuquerque, NM (United States), 1998.
- [20] N. Smaoui, "Artificial neural network-based low-dimensional model for spatio-temporally varying cellular flames," *Applied Mathematical Modelling*, vol. 21, no. 12, pp. 739–748, 1997.
- [21] N. Smaoui, "A model for the unstable manifold of the bursting behavior in the 2D navier–stokes flow," *SIAM Journal on Scientific Computing*, vol. 23, no. 3, pp. 824–839, 2001.
- [22] N. Smaoui and S. Al-Enezi, "Modelling the dynamics of nonlinear partial differential equations using neural networks," *Journal of Computational and Applied Mathematics*, vol. 170, no. 1, pp. 27–58, 2004.
- [23] G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annual Review of Fluid Mechanics*, vol. 25, no. 1, pp. 539–575, 1993.
- [24] J. L. Elman, "Finding structure in time," *Cognitive Science*, vol. 14, no. 2, pp. 179–211, 1990.
- [25] R. Bakker, J. C. Schouten, C. L. Giles, F. Takens, and C. M. van den Bleek, "Learning chaotic attractors by neural networks," *Neural Computation*, vol. 12, no. 10, pp. 2355–2383, 2000.
- [26] H. Y. Lin, W. Chen, and A. Tsutsumi, "Long-term prediction of nonlinear hydrodynamics in bubble columns by using artificial neural networks," *Chemical Engineering and Processing: Process Intensification*, vol. 42, no. 8–9, pp. 611–620, 2003.
- [27] A. Chorin and O. H. Hald, *Stochastic Tools in Mathematics and Science*, vol. 3 of Surveys and Tutorials in the Applied Mathematical Sciences, Springer, 2009.
- [28] E. Parish and K. Duraisamy, "Reduced order modeling of turbulent flows using statistical coarse-graining," in *46th AIAA Fluid Dynamics Conference*, Washington, D.C., USA, 2016.
- [29] T. Koskela, M. Lehtokangas, J. Saarinen, and K. Kaski, "Time series prediction with multilayer perceptron, FIR and Elman neural networks," in *Proceedings of the World Congress on Neural Networks*, pp. 491–496, 1996, CiteSeer.
- [30] D. S. Broomhead and G. P. King, "Extracting qualitative dynamics from experimental data," *Physica D: Nonlinear Phenomena*, vol. 20, no. 2–3, pp. 217–236, 1986.
- [31] T. Miyoshi, H. Ichihashi, S. Okamoto, and T. Hayakawa, "Learning chaotic dynamics in recurrent RBF network," in *Proceedings of ICNN'95 - International Conference on Neural Networks*, vol. 1, pp. 588–593, Perth, WA, Australia, 1995.
- [32] Y. Sato and S. Nagaya, "Evolutionary algorithms that generate recurrent neural networks for learning chaos dynamics," in *Proceedings of IEEE International Conference on Evolutionary Computation*, pp. 144–149, Nagoya, Japan, 1996.
- [33] C. A. L. Bailer-Jones, D. J. C. MacKay, and P. J. Withers, "A recurrent neural network for modelling dynamical systems," *Network: Computation in Neural Systems*, vol. 9, no. 4, pp. 531–547, 1998.
- [34] Z. Wang, D. Xiao, F. Fang, R. Govindan, C. C. Pain, and Y. Guo, "Model identification of reduced order fluid dynamics systems using deep learning," *International Journal for Numerical Methods in Fluids*, vol. 86, no. 4, pp. 255–268, 2018.
- [35] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT press, 2016.
- [36] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [37] A. P. Trischler and G. M. T. D'Eleuterio, "Synthesis of recurrent neural networks for dynamical system simulation," *Neural Networks*, vol. 80, pp. 67–78, 2016.
- [38] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989.
- [39] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, <https://arxiv.org/abs/1412.6980>.
- [40] M. Abadi, A. Agarwal, P. Barham et al., "Large-scale machine learning on heterogeneous distributed systems," CoRR, 2016, <https://arxiv.org/abs/1603.04467>.
- [41] J. Bergstra and Y. Bengio, "Random search for hyperparameter optimization," *Journal of Machine Learning Research*, vol. 13, pp. 281–305, 2012.
- [42] J. Bergstra, D. Yamins, and D. D. Cox, "Making a science of model search: hyperparameter optimization in hundreds of dimensions for vision architectures," in *Proceedings 30th International Conference on Machine Learning*, pp. 115–123, 2013.
- [43] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," in *Advances in Neural Information Processing Systems*, pp. 2951–2959, NIPS Foundation, 2012.

- [44] S. Rifai, X. Glorot, Y. Bengio, and P. Vincent, "Adding noise to the input of a model trained with a regularized objective," 2011, <https://arxiv.org/abs/1104.3250>.
- [45] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," CoRR, 2017, <https://arxiv.org/abs/1710.05941>.
- [46] R. Yu, S. Zheng, and Y. Liu, "Learning chaotic dynamics using tensor recurrent neural networks," in *Proceedings of the ICML 17 workshop on deep structured prediction*, vol. 70, Sydney, Australia, 2017PMLR.
- [47] K. Taira and T. Colonius, "The immersed boundary method: a projection approach," *Journal of Computational Physics*, vol. 225, no. 2, pp. 2118–2137, 2007.
- [48] B. R. Noack, K. Afanasiev, M. A. R. E. K. MORZYNSKI, G. Tadmor, and F. Thiele, "A hierarchy of low-dimensional models for the transient and post-transient cylinder wake," *Journal of Fluid Mechanics*, vol. 497, pp. 335–363, 2003.
- [49] J.-G. Liu, C. Wang, and H. Johnston, "A fourth order scheme for incompressible Boussinesq equations," *Journal of Scientific Computing*, vol. 18, no. 2, pp. 253–285, 2003.
- [50] E. Weinan and C.-w. Shu, "Small-scale structures in Boussinesq convection," *Physics of Fluids*, vol. 6, no. 1, pp. 49–58, 1994.
- [51] O. San and J. Borggaard, "Principal interval decomposition framework for pod reduced-order modeling of convective Boussinesq flows," *International Journal for Numerical Methods in Fluids*, vol. 78, no. 1, pp. 37–62, 2015.
- [52] Y. Li and Z. J. Wang, "A priori and a posteriori evaluations of sub-grid scale models for the burgers equation," *Computers & Fluids*, vol. 139, pp. 92–104, 2016.
- [53] S. Gottlieb, C.-W. Shu, and E. Tadmor, "Strong stability-preserving high-order time discretization methods," *SIAM Review*, vol. 43, no. 1, pp. 89–112, 2001.
- [54] S. Pan and K. Duraisamy, "Data-driven discovery of closure models," 2018, <https://arxiv.org/abs/1803.09318>.
- [55] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning, Volume 1*, Springer series in statistics New York, Springer-Verlag, New York, NY, USA, 2001.
- [56] B. Xu, R. Huang, and M. Li, "Revise saturated activation functions," 2016, <https://arxiv.org/abs/1602.05980>.

Submit your manuscripts at
www.hindawi.com