Data-driven prediction of a multi-scale Lorenz 96 chaotic system using deep learning methods: Reservoir computing, ANN, and RNN-LSTM

Ashesh Chattopadhyay[1], Pedram Hassanzadeh[1,2], and Devika Subramanian[3,4]

[1]Department of Mechanical Engineering, Rice University
[2]Department of Earth Environmental and Planetary Sciences, Rice University
[3]Department of Electrical and Computer Engineering, Rice University
[4]Department of Computer Science, Rice University

December 5, 2019

Abstract

In this paper, the performance of three deep learning methods for predicting short-term evolution and for reproducing the long-term statistics of a multi-scale spatio-temporal Lorenz 96 system is examined. The methods are: echo state network (a type of reservoir computing, RC-ESN), deep feed-forward artificial neural network (ANN), and recurrent neural network with long short-term memory (RNN-LSTM). This Lorenz 96 system has three tiers of nonlinearly interacting variables representing slow/large-scale (X), intermediate (Y), and fast/small-scale (Z) processes. For training or testing, only X is available; Y and Z are never known or used. We show that RC-ESN substantially outperforms ANN and RNN-LSTM for short-term prediction, e.g., accurately forecasting the chaotic trajectories for hundreds of numerical solver's time steps, equivalent to several Lyapunov timescales. The RNN-LSTM and ANN show some prediction skills as well; RNN-LSTM bests ANN. Furthermore, even after losing the trajectory, data predicted by RC-ESN and RNN-LSTM have probability density functions (PDFs) that closely match the true PDF, even at the tails. The PDF of the data predicted using ANN, however, deviates from the true PDF. Implications, caveats, and applications to data-driven and data-assisted surrogate modeling of complex nonlinear dynamical systems such as weather/climate are discussed.

1 Introduction

Various components of the Earth system involve multi-scale, multi-physics processes and high-dimensional chaotic dynamics. These processes are often modeled using sets of strongly-coupled nonlinear partial differential equations (PDEs), which are solved numerically on supercomputers. As we aim to simulate such systems with increasing levels of fidelity, we need to increase the numerical resolutions and/or incorporate more physical processes from a wide range of spatio-temporal scales into the models. For example, in atmospheric modeling for predicting the weather and climate systems, we need to account for the nonlinear interactions across the scales of cloud microphysics processes, gravity waves, convection, baroclinic waves, synoptic eddies, and large-scale circulation, just to name a few, not to mention the fast/slow processes involved in feedbacks from the ocean, land, and cryosphere [?, ?, ?, ?, ?].

Solving the coupled systems of PDEs representing such multi-scale processes is computationally prohibitive for many practical applications. As a result, to make the simulations tractable, a few strategies have been developed, which mainly involve only solving for slow/large-scale variables and accounting for the fast/small-scale processes through surrogate models. In weather/climate models, for example, semi-empirical physics-based parameterizations are often used to represent the effects of processes such as gravity waves and moist convection in the atmosphere or sub-mesoscale eddies in the ocean [?, ?, ?, ?, ?]. A more advanced approach is "super-parameterization", which for example, involves solving the PDEs of moist convection on

a high-resolution grid inside each grid point of large-scale atmospheric circulation, whose governing PDEs (the Navier-Stokes equations) are solved on a coarse grid [**?**]. While computationally more expensive, super-parameterized climate models have been shown to outperform parameterized models in simulating some aspects of climate variability and extremes [**?, ?, ?**].

More recently, "inexact computing" has received attention from the weather/climate community. This approach involves reducing the computational cost of each simulation by decreasing the *precision* of some of the less-critical calculations [**?, ?**], thus allowing the saved resources to be *reinvested*, for example, in more simulations (e.g., for probabilistic predictions) and/or higher resolutions for critical processes [**?, ?, ?, ?**]. One type of inexact computing involves solving the PDEs of some of the processes with single- or half-precision arithmetic, which requires less computing power and memory compared to the conventional double-precision calculations [**?, ?, ?**].

In the past few years, the rapid algorithmic advances in artificial intelligence (AI), and in particular data-driven modeling, have been explored for improving the simulation and prediction of nonlinear dynamical systems [**?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?**]. One appeal of the data-driven approach is that fast/accurate data-driven surrogate models that are trained on data from high-fidelity, computationally demanding simulations can be used to accelerate and/or improve the prediction and simulation of complex dynamical systems. Furthermore, for some poorly understood processes for which observational data are available (e.g., clouds), data-driven surrogate models built using such data might potentially outperform physics-based surrogate models [**?, ?**]. Recent studies have shown promising results in using AI to build data-driven parameterizations for modeling of some atmospheric and oceanic processes [**?, ?, ?, ?, ?, ?, ?, ?**]. In the turbulence and dynamical systems communities, similarly encouraging outcomes have been reported [**?, ?, ?, ?, ?, ?, ?, ?, ?, ?**] .

The objective of the current study is to make progress toward addressing the following question: Which AI-based data-driven technique(s) can best predict the spatio-temporal evolution of a multi-scale chaotic system, when only the slow/large-scale variables are known (during training) and are of interest (to predict during testing)? We emphasize that unlike many other studies, our focus is not on reproducing long-term statistics of the underlying dynamical system (although that will be examined too), but on predicting the short-term trajectory from a given initial condition. Furthermore, we emphasize that, following the problem formulation introduced by [**?**], the system's state-vector is only partially known and the fast/small-scale variables are unknown even during training.

Our objective is more clearly demonstrated using the canonical chaotic system that we will use as a testbed for the data-driven methods: A multi-scale Lorenz 96 system:

$$\frac{dX_k}{dt} = X_{k-1}\left(X_{k+1} - X_{k-2}\right) + F - \frac{hc}{b}\Sigma_j Y_{j,k} \tag{1}$$

$$\frac{dY_{j,k}}{dt} = -cbY_{j+1,k}\left(Y_{j+2,k} - Y_{j-1,k}\right) - cY_{j,k} + \frac{hc}{b}X_k - \frac{he}{d}\Sigma_i Z_{i,j,k} \tag{2}$$

$$\frac{dZ_{i,j,k}}{dt} = edZ_{i-1,j,k}\left(Z_{i+1,j,k} - Z_{i-2,j,k}\right) - geZ_{i,j,k} + \frac{he}{d}Y_{j,k} \tag{3}$$

This set of coupled nonlinear ODEs is a 3-tier extension of Lorenz's original model [**?**] and has been proposed by Thornes *et al.* [**?**] as a fitting prototype for multi-scale chaotic variability of the weather/climate system and a useful testbed for novel methods. In these equations, $F = 20$ is a large-scale forcing that makes the system highly chaotic; $b = c = e = d = g = 10$ and $h = 1$ are tuned to produce appropriate spatio-temporal variability in the three variables (see below). The indices $i, j, k = 1, 2, \ldots 8$, thus X has 8 elements while Y and Z have 64 and 512 elements, respectively. Figure 1 shows examples of the chaotic temporal evolution of X, Y, and Z obtained from directly solving Eqs. (1)-(3). The examples demonstrate that X has large amplitudes and slow variability; Y has relatively small amplitudes, high-frequency variability, and intermittency; and Z has small amplitudes and high-frequency variability. In the context of atmospheric circulation, the slow variable X can represent the low-frequency variability of the large-scale circulation while the intermediate variable Y and fast variable Z can represent synoptic/baroclinic eddies and small-scale convection, respectively. Similar analogies in the context of ocean circulation and various other natural or engineering systems can be found, making this multi-scale Lorenz 96 system a useful prototype to focus on.

Our objective is to predict the spatio-temporal evolution of $X(t)$ using a data-driven model that is trained on past observations of $X(t)$ (Figure 2). In the conventional approach of solving Eqs. (1)-(3) numerically,

Figure 1: Time evolution of (A) X_k at grid point $k = 1$, (B) $Y_{j,1}$ at grid point $j = 1$, and (C) $Z_{i,1,1}$ at grid point $i = 1$. The time series show chaotic behavior in X, which has large amplitude and slow variability; Y, which has relatively small amplitudes, high frequency variability, and intermittency; Z, which has small amplitudes and high frequency variability. The x-axis is in model time unit (MTU) which is related to the time step of the numerical solver (Δt) and largest positive Lyapunov exponent (λ_{max}) as 1 MTU $= 200\Delta t \approx 4.5/\lambda_{max}$ (see section 2.1). Note that here we are presenting raw data, which have not been standardized for prediction/testing yet.

the governing equations have to be known, initial conditions for $Y(t)$ and $Z(t)$ have to be available, and the numerical resolution is dictated by the fast/small-scale variable Z, leading to high computational costs. In the fully data-driven approach that is our objective here, the governing equations do not have to be known, $Y(t)$ and $Z(t)$ do not have to be observed at any time, and evolution of $X(t)$ is predicted just from knowledge of the past observations of X, leading to low computational costs. To successfully achieve this objective, a data-driven method should be able to

1. Accurately predict the evolution of a chaotic system along a trajectory for some time,

2. Account for the effects of $Y(t)$ and $Z(t)$ on the evolution of $X(t)$.

Inspired by several recent studies (that are discussed below), we have focused on evaluating the performance of three deep learning techniques in accomplishing (1) and (2). These data-driven methods are

- RC-ESN: Echo state network (ESN), a specialized type of recurrent neural network (RNN), which belongs to the family of reservoir computing (RC),

- ANN: A deep feed-forward artificial neural network,

- RNN-LSTM: An RNN with long short-term memory (LSTM).

We have focused on these three methods because they have either shown promising performance in past studies (RC-ESN and ANN), or they are considered state of the art in learning from sequential data (RNN-LSTM). There is a growing number of studies focused on using deep learning techniques for data-driven modeling of chaotic and turbulent systems, for example to improve weather/climate modeling and
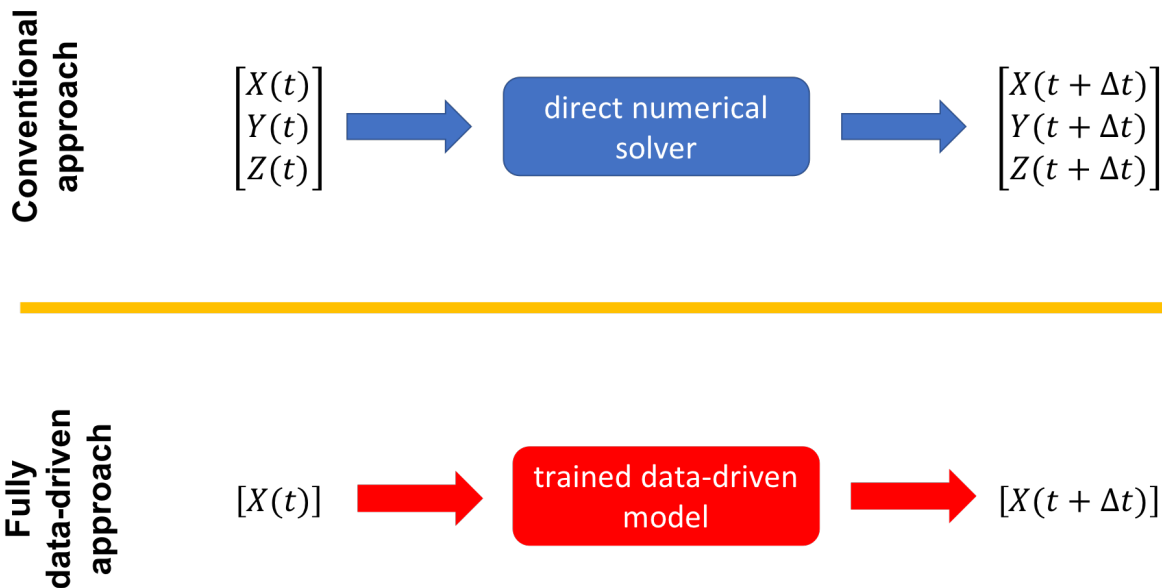
Figure 2: The conventional approach of solving the full governing equations numerically versus our data-driven approach for predicting the spatio-temporal evolution of the slow/large-scale variable X. For the direct numerical solution, the governing equations have to be known and the numerical resolution is dictated by the fast and small-scale variable Z, resulting in high computational costs. In the data-driven approach, the governing equations do not have to be known, Y and Z do not have to be known, and evolution of X is predicted just from knowing the past observations of X, leading to low computational costs.

prediction. Some of these studies have been referenced above. Below, we briefly describe three sets of studies with closest relevance to our objective and approach. Pathak and co-workers [**?**, **?**, **?**] have recently shown promising results with RC-ESN for predicting short-term spatio-temporal evolution (item 1) and in replicating attractors for the Lorenz 63 and Kuramoto-Sivashinsky equations. The objective of our paper (items 1-2) and the employed multi-scale Lorenz 96 system are identical to that of [**?**], who reported their ANN to have some skills in data-driven prediction of the spatio-temporal evolution of X. Finally, Vlachas *et al.* [**?**] found an RNN-LSTM to have skills (though limited) for predicting the short-term spatio-temporal evolution (item 1) of a number of chaotic toy models such as the original Lorenz 96 system.

Here we aim to build on these pioneering studies and examine, *side by side*, the performance of RC-ESN, ANN, and RNN-LSTM in achieving (1) and (2) for the chaotic multi-scale Lorenz 96 system. We emphasize the need for such a side-by-side comparison that uses the exact same system as the testbed and metrics for assessing performance.

The structure of the paper is as follows: in section 2, the multi-scale Lorenz 96 system and the three deep learning methods are discussed; results on how these methods predict the short-term spatio-temporal evolution of X and reproduce the long-term statistics of X are presented in section 3; key findings and future work are discussed in section 4.

2 Materials and Methods

2.1 The multi-scale Lorenz 96 system

2.1.1 Numerical solution

We have used a 4th-order Runge-Kutta solver with time step $\Delta t = 0.005$ to solve the system of Eqs. (1)-(3). The system has been integrated for 100 million time steps to generate a large dataset for training, testing, and

examining the robustness of the results. In the Results section, we often report time in terms of model time units (MTUs), where 1 MTU = $200\Delta t$. In terms of Lyapunov timescale, 1 MTU in this system is $\approx 4.5/\lambda_{max}$ [?], where λ_{max} is the largest positive Lyapunov exponent. In terms of the e-folding decorrelation timescale (τ) of the leading principal component (PC1) of X [?], we estimate 1 MTU $\approx 6.9\tau_{\mathrm{PC1}}$. Finally, as discussed in Thornes *et al.* [?], 1 MTU in this system corresponds to ≈ 3.75 Earth days in the real atmosphere.

2.1.2 Training and testing datasets

To build the training and testing datasets from the numerical solution, we have sampled $X \in \Re^8$ at every Δt and then standardized the data by subtracting the mean and dividing by the standard deviation. We then construct a training set consisting of N sequential samples and a testing consisting of the next 2000 sequential samples from point $N + 1$ to $N + 2000$. We have randomly chosen 100 such training/testing sets, each of length $(N + 2000)\Delta t$ starting from a random point and separated from the next training/testing set by at least $2000\Delta t$. There is never any overlap between the training and the testing sequences in any of the 100 training/testing sets.

2.2 Reservoir computing-echo state network (RC-ESN)

2.2.1 Architecture

The RC-ESN [?, ?] is an RNN that has a reservoir with D neurons, which are sparsely connected in an Erdős-Rényi graph configuration (see Figure 3). The connectivity of the reservoir neurons is represented by the adjacency matrix $\mathbf{A}$ of size $D \times D$ whose values are drawn from a uniform random distribution on the interval $[-1, 1]$. The state of the reservoir, representing the activations of its constituent neurons, is a vector $r(t) \in \Re^D$. The typical reservoir size used in this study is $D = 5000$ (but we have experimented with D as large as 20000, as discussed later). The other two components of the RC-ESN are an input-to-reservoir layer with weight matrix $\mathbf{W_{in}}$, and a reservoir-to-output layer with weight matrix $\mathbf{W_{out}}$. The inputs for training are N sequential samples of $X(t) \in \Re^8$. At the beginning of the training phase, $\mathbf{A}$ and $\mathbf{W_{in}}$ are chosen randomly and are fixed; i.e., they do not change during training or testing. During training, only the weights of the output-to-reservoir layer ($\mathbf{W_{out}}$) are updated.

$\mathbf{W_{out}}$ is the only trainable matrix in this network. This architecture yields a simple training process that has two key advantages:

- It does not suffer from the vanishing and the exploding gradient problem, which has been a major difficulty in training RNNs, especially before the advent of LSTMs [?],

- $\mathbf{W_{out}}$ can be computed in one shot (see below), thus this algorithm is orders of magnitude faster than the backpropagation through time (BPTT) algorithm [?], which is used for training general RNNs and RNN-LSTMs (see section 2.4).

The equations governing the RC-ESN training process are as follows:

$$r(t + \Delta t) = \tanh\left(\mathbf{A}r(t) + \mathbf{W_{in}}X(t)\right), \tag{4}$$

$$\mathbf{W_{out}} = \arg\min_{\mathbf{W_{out}}} ||\mathbf{W_{out}}\tilde{r}(t) - X(t)|| + \alpha||\mathbf{W_{out}}||. \tag{5}$$

Here $||\cdot||$ is the L_2-norm of a vector and α is the L_2 regularization (ridge regression) constant. Equation (4) maps the observable $X(t) \in \Re^8$ to the higher dimensional reservoir's state $r(t + \Delta t) \in \Re^D$ (reminder: D is $O(1000)$ for this problem). Note that Eq. (5) contains $\tilde{r}(t)$ rather than $r(t)$. As observed by Pathak *et al.* [?] for the Kuramoto-Sivashinsky equation, and observed by us in this work, the columns of the matrix $\tilde{\mathbf{r}}$ should be chosen as nonlinear combinations of the columns of the reservoir state matrix $\mathbf{r}$ (each row of the matrix $\mathbf{r}$ is a snapshot $r(t)$ of size D). For example, following Pathank *et al.* [?], we compute $\tilde{\mathbf{r}}$ as having the same even columns as $\mathbf{r}$ while its odd columns are the square of the odd columns of $\mathbf{r}$ (algorithm T_1 hereafter). As shown in Appendix A, we have found that a nonlinear transformation (between $\mathbf{r}$ and $\tilde{\mathbf{r}}$) is essential for skillful predictions, while several other transformation algorithms yield similar results as T_1. The choices of these transformations (T_2 and T_3, see Appendix A), although not based on a rigorous mathematical analysis, are inspired from the nature of the quadratic nonlinearity that is present in Eqs. (1)-(3).

Figure 3: A schematic of RC-ESN. $D \times D$ is the size of the adjacency matrix of the reservoir $\mathbf{A}$. $X(t)$ ($-T \leq t \leq 0$) is the input during training. $\mathbf{W_{in}}$ and $\mathbf{A}$ are random matrices that are chosen at the beginning of training and fixed during training and testing. Only $\mathbf{W_{out}}$ is computed during training. During testing ($t > 0$), $X(t + \Delta t)$ is predicted from a given $X(t)$ that is either known from the initial condition or has been previously predicted.

The prediction process is governed by:

$$v(t + \Delta t) = \mathbf{W_{out}} \tilde{r}(t + \Delta t), \tag{6}$$
$$X(t + \Delta t) = v(t + \Delta t). \tag{7}$$

$\tilde{r}(t + \Delta t)$ in Eq. (6) is computed by applying one of the T_1, T_2 or T_3 algorithms on $r(t + \Delta t)$, which itself is calculated via Eq. (4) from $X(t)$ that is either known from initial condition or has been previously predicted.

See [?, ?, ?, ?] and references therein for further discussions on RC-ESNs, and [?, ?, ?, ?, ?, ?, ?, ?, ?] for examples of recent applications to dynamical systems.

2.2.2 Training and Prediction

During training ($-T \leq t \leq 0$), $\mathbf{W_{in}}$ and $\mathbf{A}$ are initialized with random numbers, which stay fixed during the training (and testing) process, The state matrix $\mathbf{r}$ is initialized to 0. During initialization, we ensure that the spectral radius of $\mathbf{A}$ is less than unity by first dividing the matrix $\mathbf{A}$ with its largest eigenvalue Λ and further multiplying it by a scalar ($\rho \leq 1$). Then the state matrix $\mathbf{r}$ is computed using Eq. (4) for the training set, and $\mathbf{W_{out}}$ is computed using Eq. (5). During prediction (i.e., testing) corresponding to $t > 0$, the computed $\mathbf{W_{out}}$ is used to march $v(t)$ (and thus $X(t)$) forward in time (Eqs. (6)-(7)) while as mentioned earlier, $r(t)$ keeps getting updated using the predicted $X(t)$ (Eq. (4)). A non-linear transformation is used to compute $\tilde{r}(t)$ from $r(t)$ before using Eq. (6).

Our RC-ESN architecture and training/prediction procedure are similar to the ones used in [?]. There is no overfitting in the training phase because the final training and testing accuracies are the same. Our code is developed in Python and is made publicly available (see Code and data availability).

2.3 Feed-forward artificial neural network (ANN)

2.3.1 Architecture

We have developed a deep ANN that has the same architecture as the one used in Duben and Bauer [?] (which they employed to conduct prediction on the same multi-scale Lorenz 96 studied here). The ANN has 4 hidden layers with 100 neurons each, and 8 neurons in the input and output layers (Figure 4). It must be highlighted that unlike RC-ESN (and RNN-LSTM), this ANN is stateless, i.e., there is no hidden variable such as $r(t)$ that tracks temporal evolution. Furthermore, unlike the RC-ESN, for which the input and output are $X(t)$ and $X(t + \Delta t)$, following Duben and Bauer [?], the ANN's inputs/outputs are chosen to be pairs of $X(t)$ and $\Delta X(t) = X(t + \Delta t) - X(t)$ (Figure 4). During prediction, using $X(t)$ that is known from initial

Figure 4: A schematic of ANN, which has 4 hidden layers each with 100 neurons, and an input and an output layer each with 8 neurons (for better illustration only few neurons from each layer are shown). Training is performed on pairs of $X(t)$ and $\Delta X(t) = X(t + \Delta t) - X(t)$, and the weights in each layer are learned through the backpropagation algorithm. During testing (i.e., prediction) $\Delta X(t)$ is predicted for a given $X(t)$ that is already predicted or is known from initial condition. Knowing $X(t)$, $\Delta X(t)$, and $\Delta X(t - \Delta t)$, $X(t + \Delta t)$ is then computed using Adams-Bashforth integration scheme.

condition or previously calculated, $\Delta X(t)$ is predicted. $X(t + \Delta t)$ is then computed via Adams-Bashforth integration scheme

$$X(t + \Delta t) = X(t) + \frac{1}{2}\big[3\Delta X(t) - \Delta X(t - \Delta t)\big]. \tag{8}$$

More details on the ANN architecture has been reported in Appendix B.

2.3.2 Training and Prediction

Training is performed on N pairs of sequential $(X(t), \Delta X(t))$ from the training set. During training, the weights of the network are computed using backpropagation optimized by the stochastic gradient descent algorithm. During prediction, as mentioned above, $\Delta X(t)$ is predicted from $X(t)$ that is known from initial condition or has been previously predicted, and $X(t + \Delta t)$ is then calculated using Eq. (8).

Our ANN architecture and training/prediction procedure are similar to the ones used in Duben and Bauer [?] (we have optimized, by trial and error, the hyperparameters for this particular network). There is no overfitting in the training phase because the final training and testing accuracies are the same. Our code is developed in Keras and is made publicly available (see Code and data availability).

2.4 Recurrent neural network with long short-term memory (RNN-LSTM)

2.4.1 Architecture

The RNN-LSTM [?] is a deep learning algorithm most suited for prediction of sequential data such as time series, and has received a lot of attention in recent years [?]. Variants of RNN-LSTMs are the best performing models for time series modeling in areas such as stock pricing [?], supply chain [?], natural language processing [?], and speech recognition [?]. A major improvement over regular RNNs, which have issues with the vanishing and the exploding gradients [?], LSTMs have become the state-of-the-art approach for training RNNs in the deep learning community. Unlike regular RNNs, the RNN-LSTMs have gates that control the information flow into the neural network from previous time steps of the time series. The

RNN-LSTM used here, like the RC-ESN but unlike the ANN, is stateful (i.e., actively maintains state). Our RNN-LSTM has 50 hidden layers in each RNN-LSTM cell. More details on our RNN-LSTM are presented in C. RNN-LSTMs have many tunable parameters and are trained with the expensive BPTT algorithm [**?**].

2.4.2 Training and Prediction

The input to the RNN-LSTM is a time-delay-embedded matrix of $X(t)$ with embedding dimension q (also known as lookback) [**?**]. An extensive hyperparameter optimization (by trial and error) is performed to find the optimal value of q for which the network has the largest prediction horizon (explooring $q = 1 - 22$, we found $q = 3$ to yield the best performance). The RNN-LSTM predicts $X(t + \Delta t)$ from the previous q time steps of $X(t)$. This is in contrast with RC-ESN, which only uses $X(t)$ and reservoir state $r(t)$ to predict $X(t + \Delta t)$, and ANN, which only uses $X(t)$ and no state to predict $X(t + \Delta t)$ (via predicting $\Delta X(t)$). The weights of the LSTM layers are determined during the training process (see C). During testing, $X(t + \Delta t)$ is predicted using the past q observables $[X(t - (q-1)\Delta t) \ldots X(t - \Delta t), X(t)]$ that are either known from initial condition or have been previously predicted. We have found the best results with a stateless LSTM (which means that the state of the LSTM gets refreshed during the beginning of each batch during training, see C) that outperforms a stateful LSTM (where the states gets carried over to the next batch during training).

The architecture of our RNN-LSTM is similar to the one used in Vlachas *et al.* [**?**]. There is no overfitting in the training phase because the final training and testing accuracies are the same. Our code is developed in Keras and is made publicly available (see Code and data availability).

3 Results

3.1 Short-term prediction: Comparison of the RC-ESN, ANN, and RNN-LSTM performances

The short-term prediction skills of the three deep learning methods for the same training/testing sets are compared in Figure 5. Given the chaotic nature of the system, the performance of the methods depends on the initial condition from which the prediction is conducted. To give the readers a comprehensive view of the performance of these methods, Figures 5(A) and (B) show examples of the predicted trajectories (for one element of $X(t)$) versus the true trajectory for two specific initial conditions (from the 100 initial conditions we used): the one for which RC-ESN shows the best performance (Figure 5(A)), and the one for which RC-ESN shows the worst performance (Figure 5(B)).

As seen in Figure 5(A), RC-ESN accurately predicts the time series for over 2.3 MTU, which is equivalent to $460\Delta t$ and over 10.35 Lyapunov timescales. Closer examination shows that the RC-RSN prediction follows the true trajectory well even up to ≈ 4 MTU. The RNN-LSTM has the next best prediction performance (up to around 0.9 MTU or $180\Delta t$). The prediction from ANN is for around 0.6 MTU or $120\Delta t$. For the example in Figure 5(B), all methods have shorter prediction horizons, but RC-ESN still has the best performance (accurate prediction up to ≈ 0.7 MTU), followed by ANN and RNN-LSTM with similar prediction accuracies (≈ 0.3 MTU).

Searching through all 100 initial conditions, the best prediction with RNN-LSTM is up to ≈ 1.7 MTU (equivalent to $340\Delta t$), and the best prediction with ANN is up to ≈ 1.2 MTU (equivalent to $240\Delta t$),

To compare the results over all 100 randomly chosen initial conditions, we have defined an averaged relative L_2 error between the true and predicted trajectories

$$e(t) = \left[\frac{||X_{true}(t) - X_{pred}(t)||}{\langle ||X_{true}(t)|| \rangle} \right]. \tag{9}$$

Here $[\cdot]$ and $\langle \cdot \rangle$ indicate, respectively, averaging over 100 initial conditions and over $2000\Delta t$. To be clear, $X_{true}(t)$ refers to the data at time t obtained from the numerical solution while $X_{pred}(t)$ refers to the predicted value at t using one of the deep learning methods. Figure 5(C) compares $e(t)$ for the three methods. It is clear that RC-ESN significantly outperforms ANN and RNN-LSTM for short-term prediction in this multi-scale chaotic testbed. Figure 6 shows an example of the spatio-temporal evolution of X_{pred} (from RC-ESN), X_{true},

Figure 5: Comparison of the short-term forecast skills among the three deep learning methods. The lines show truth (black), RC-ESN (red), ANN (blue), and RNN-LSTM (cyan). (A) and (B) show examples (differing in initial condition) for which RC-ESN yields the longest and shortest prediction horizons, respectively. Vertical dashed lines approximately show the prediction horizon of RC-ESN. (C) Relative L_2 error averaged over 100 randomly chosen initial conditions $e(t)$ (see Eq. (9)). The dashed horizontal line marks $e = 0.3$. Time is in MTU, where 1 MTU = $200\Delta t \approx 4.5/\lambda_{max}$. Training size is $N = 500000$.

and their difference, which further demonstrates the capabilities of RC-ESN for short-term spatio-temporal prediction.

3.2 Short-term prediction: Scaling of RC-ESN and ANN performance with training size N

How the performance of deep learning techniques scales with the size of the training set is of significant practical importance as the amount of available data is often limited in many problems. Given that currently there is no theoretical understanding of such scaling for these deep learning techniques, we have empirically examined how the quality of short-term predictions scales with N. We have conducted the scaling analysis for $N = 10^4$ to $N = 2 \times 10^6$ for the three methods. Two metrics for the quality of prediction are used: the prediction horizon, defined as when the averaged L_2 error $e(t)$ reaches 0.3, and the prediction error E, defined as the average of $e(t)$ between $0 - 0.5$ MTU:

$$E = \frac{1}{100\Delta t} \sum_{i=0}^{i=100} e(i\Delta t). \tag{10}$$

Figure 7(A) shows that for all methods, the prediction horizon increases monotonically but nonlinearly as we increase N. The prediction horizons of RC-ESN and ANN appear to saturate after $N = 10^6$, although

Figure 6: Performance of RC-ESN for short-term spatio-temporal prediction. We remind the readers that only the slow/large-scale variable X has been used during training/testing, and the fast/small-scale variables Y and Z have not been used at any point during training or testing. RC-ESN has substantial forecast skills, providing accurate predictions up to ≈ 2 MTU, which is around $400\Delta t$ or 9 Lyapunov timescales. $N = 500000$, algorithm T_2, $D = 5000$, and $\rho = 0.1$ are used.

the RC-ESN has a more complex step-like scaling curve that needs further examination in future studies. The prediction horizon of RC-ESN exceeds that of ANN by factors ranging from 3 (for high N) to 9 (for low N). In the case of RNN-LSTM, the factor ranges from 1.2 (for high N) to 2 (for low N). Figure 7(B) shows that for all methods, the average error E decreases as N increases (as expected), most notably for ANN when N is small.

Overall, compared to both RNN-LSTM and ANN, the prediction horizon and accuracy of RC-ESN have a weaker dependence on the size of the training set, which is a significant advantage for RC-ESN when the dataset available for training is short, which is common in many practical problems.

3.3 Short-term prediction: Scaling of RC-ESN performance with reservoir size D

Given the superior performance of RC-ESN for short-term prediction, here we focus on one concern with this method: the need for large reservoirs, which can be computationally demanding. This issue has been suggested as a potential disadvantage of ESNs versus LSTMs for training RNNs [?]. Aside from the observations here that RC-ESN significantly outperforms RNN-LSTM for short-term predictions, the problem of reservoir size can be tackled in at least two ways. First, Pathak *et al.* [?] have proposed, and shown the effectiveness of, using a set of parallelized reservoirs, which allowed for easily dealing with high-dimensional chaotic toy models.

Second, Figure 8 shows that E rapidly declines by a factor of around 3 as D is increased from 500 to 5000, decreases slightly as D is further doubled to 10000, and then barely changes as D is doubled again to 20000. Training the RC-ESN with $D = 20000$ versus $D = 5000$ comes with a higher computational cost (16 GB

Figure 7: Comparison of the short-term prediction quality for RC-ESN (blue circles), RNN-LSTM (black stars), and ANN (red squared) as the size of the training set N is changed from $N = 10^4$ to $N = 2 \times 10^6$. (A) Prediction horizon (when $e(t)$ reaches 0.3). (B) Average error E (see Eq. (10)). The hyperparameters in each method are optimized for each N.

memory and ≈ 6 CPU hours for $D = 5000$ and 64 GB memory and ≈ 18 CPU hours for $D = 20000$), while little improvement in accuracy is gained. Thus, concepts from inexact computing can be used to choose D such that precision is traded for large savings in computational resources, which can be then reinvested into more simulations, higher resolutions for critical processes etc. [**?**, **?**, **?**, **?**, **?**].

3.4 Long-term statistics: Comparison of RC-ESN, ANN, and RNN-LSTM performance

All the data-driven predictions discussed earlier eventually diverge from the true trajectory (as would even predictions using the numerical solver). Still, it is interesting to examine whether the freely predicted spatio-temporal data have the same long-term statistical properties as the actual dynamical system (i.e., Eqs. (1)-(3)). Reproducing the actual dynamical system's long-term statistics (sometimes refer to as the system's *climate*) using a data-driven method can be significantly useful. In some problems, a surrogate model does not need to predict the evolution of a specific trajectory, but only the long-term statistics of a system. Furthermore, synthetic long datasets produced using an inexpensive data-driven method (trained on a shorter real dataset) can be used to examine the system's probability density functions (PDFs), including its tails, which are important for studying the statistics of rare/extreme events.

By examining return maps, Pathak *et al.* [**?**] have already shown that RC-ESNs can reproduce the long-term statistics of the Lorenz 63 and Kuramoto-Sivashinsky equations (Jaeger *et al.* [**?**] and Pathak *et al.* [**?**] have also shown that RC-ESNs can be used to accurately estimate a chaotic system's Lyapunov spectrum). Here, we focus on comparing the performance of RC-ESN, ANN, and RNN-LSTM in reproducing the system's long-term statistics by

- Examining the estimated PDFs and in particular their tails,

- Investigating whether the quality of the estimated PDFs degrades with time, which can negate the usefulness of long synthetic datasets.

Figure 9 compares the estimated PDFs obtained using the three deep learning methods. The data predicted using RC-ESN and RNN-LSTM are found to have PDFs closely matching the true PDF, even at the tails. Deviations at the tails of the PDF predicted by these methods from the true PDF are comparable to the deviations of the PDFs obtained from true data using the same number of samples. The ANN-predicted

Figure 8: For RC-ESN: scaling of average prediction error between $0 - 0.5$ MTU (E, Eq. (10)) with reservoir size D. $N = 500000$ is used.

data have reasonable PDFs between ± 2 standard deviations, but the tails have substantial deviations from those of the true PDFs. All predicted PDFs are robust and do not differ much (except near the end of the tails) among the quartiles.

The results show that RC-ESN and RNN-LSTM can accurately reproduce the system's long-term statistics, and can robustly produce long synthetic datasets with PDFs that are close to the PDF of the true data even near the tails. The ability of ANN to accomplish these tasks is limited.

4 Discussion

By examining the true and predicted trajectories (Figures 5-6) and the prediction errors and horizons (Figure 7), we have shown that RC-ESN substantially outperforms ANN and RNN-LSTM in predicting the short-term evolution of a multi-scale Lorenz 96 chaotic system (Eqs. (1)-(3)). Additionally, RC-ESN and RNN-LSTM both work well in reproducing the long-term statistics of this system (Figure 9). We emphasize that following the problem formulation of Duben and Bauer *et al.* [**?**], and unlike most other studies, only part of the multi-scale state-vector (the slow/large-scale variable X) has been available for training the data-driven model and has been of interest for testing. This problem design is more relevant to many practical problems but is more challenging as it requires the data-driven model to not only predict the evolution of X based on its past observations, but also to account for the effects of the intermediate and fast/small-scale variables Y and Z.

We have also found that the prediction horizon, and in particular the prediction accuracy, of RC-ESN to have a weak dependence on the size of the training set (Figure 7). This is an important property, as in many practical problems the data available for training are limited. Furthermore, the prediction error of RC-ESN is shown to have an asymptotic behavior for large reservoir sizes, which suggests that reasonable accuracy can be achieved with moderate reservoir sizes. Note that the skillful predictions with RC-ESN in Figures 5-6 have been obtained with a moderate-sized training set ($N = 500000$) and reservoir ($D = 5000$). Figures 7 and 8 suggest that slightly better results could have been obtained using larger N and D, although such improvements come with a higher computational cost during training.

The order we have found for the performance of the three deep learning methods (RC-ESN > RNN-LSTM > ANN) is different from the order of complexity of these methods in terms of the number of trainable parameters (RC-ESN < ANN < RNN-LSTM) but aligned with the order in terms of the learnable function space (ANN < RC-ESN and RNN-LSTM, as the latter methods are both RNNs and thus Turning complete, while ANN is a state-less feed-forward network and thus not Turing complete [**?**]). Whether the superior predictive performance of RC-ESN (especially over RNN-LSTM) is due to its explicit state representation

Figure 9: Estimated probability density functions (PDFs) of long datasets freely predicted using RC-ESN (red), ANN (blue), and RNN-LSTM (cyan) compared to the true PDF (dashed black). RC-ESN, ANN, and RNN-LSTM are first trained with $N = 500000$, and then predict, from an initial condition, for $4 \times 10^6 \Delta t$. The panels $Q_1 - Q_4$ correspond to the equally divided quartiles of 10^6 predicted samples. PDFs are approximated using kernel density estimation [?]. The green lines show the estimated PDFs from different sets of 10^6 samples that are obtained from numerical solution of Eqs. (1)-(3). Small differences between the green lines show the uncertainties in estimating the tails from 10^6 samples. The dashed black lines show the true PDF estimated from 10^7 samples from numerical solution. Note that the presented PDFs are for standardized data.

updated at every time step both during training and testing, or lower likelihood of overfitting due to the order-of-magnitude smaller number of trainable parameters, remains to be seen. Also whether we have not fully harnessed the power of RNN-LSTMs (see below) is unclear at this point, particularly because a complete theoretical understanding of how/why these methods work (or do not work) is currently lacking. That said, there have been some progress in understanding the RC-ESNs, in particular by modeling the network itself as a dynamical system [?, ?]. Such efforts, for example those aimed at understanding the echo states that are learned in the RC-ESN's reservoir, might benefit from recent advances in dynamical systems theory [?, ?, ?, ?, ?, ?].

An important next step in our work is to determine how generalizable our findings are. This investigation is important for the following two reasons. First, here we have only studied one system, a specially designed version of the Lorenz 96 system. The performance of these methods should be examined in a hierarchy of chaotic dynamical systems and high-dimensional turbulent flows. That said, our findings are overall consistent with the recently reported performance of these methods applied to chaotic toy models. [?, ?, ?] demonstrated that RC-ESN can predict, for several Lyapunov timescales, the spatio-temporal evolution, and can reproduce the climate of the Lorenz 63 and Kuramoto-Sivashinsky chaotic systems. Here, we have shown that RC-ESN performs similarly well even when only the slow/large-scale component of the multi-scale state-vector is known. Our results with ANN are consistent with those of Dueben and Bauer [?], who showed examples of trajectories predicted accurately up to 1 MTU with a large training set $N = 2 \times 10^6$

using ANN for the same Lorenz 96 system (see their Fig. 1). While RNN-LSTM is considered state-of-the-art for sequential data modeling, and has worked well for a number of applications involving time series, to the best of our knowledge, simple RNN-LSTMs, such as the one used here, have not been very successful when applied to chaotic dynamical systems. Vlachas *et al.* [?] found some prediction skills using RNN-LSTM applied to the original Lorenz 96 system (which does not have the multi-scale coupling); see their Fig. 5.

Second, we have only used a simple RNN-LSTM; there are other variants of this architecture as well as more advanced deep learning RNNs that might potentially yield better results. For our simple RNN-LSTM, we have extensively explored optimization of hyperparameters and tried variant formulations of the problem: predict $\Delta X(t)$ with or without updating the state of the LSTM. We have found that such variants do not lead to better results, consistent with the findings of Vlachas *et al.* [?]. Our preliminary explorations with more advanced variants of RNN-LSTM (seq2seq and encoder-decoder LSTM [?]) have not resulted in any improvement either. However, just as the skillful predictions of RC-ESN and ANN hinge on one key step (nonlinear transformation for the former and predicting ΔX for the latter), it is possible that changing/adding one step leads to major improvements in RNN-LSTM (it is worth mentioning that similar nonlinear transformation of the state in LSTMs is not straightforward due to their more complex architecture; see Appendix C). We have shared our codes publicly to help others explore such changes to our RNN-LSTM. Furthermore, there are other more sophisticated RNNs that we have not explored. For example [?] have introduced a tensor-train RNN that outperformed a simple RNN-LSTM in predicting the temporal evolution of the Lorenz 63 chaotic system. Mohan *et al.*. [?] showed that a compressed convolutional LSTM performs well in capturing the long-term statistics of a turbulent flow. The performance of more advanced deep learning RNNs should be examined and compared, side by side, with the performance of RC-ESNs and ANNs in future studies. The multi-scale Lorenz 96 system provides a suitable starting point for such comparisons.

The results of our work show the promise of deep learning methods such as RC-ESNs (and to some extent RNN-LSTM) for data-driven modeling of chaotic dynamical systems, which has broad applications in geosciences, e.g., in weather/climate modeling. Practical and fundamental issues such as interpretability, scalability to higher dimensional systems [?], presence of measurement noise in the training data and initial conditions [?], non-stationarity of the time series, and dealing with data that have two or three spatial dimensions (e.g., through integration with convolutional neural networks, CNN-LSTM [?] and CNN-ESN [?] should be studied in future work.

Here we have focused on a fully data-driven approach, as opposed to the conventional approach of direct numerical solutions (Figure 2). In practice, for example for large-scale, multi-physics, multi-scale dynamical systems such as weather and climate models, it is likely that a hybrid framework yields the best performance: depending on the application and the spatio-temporal scales of the physical processes involved [?, ?], some of the equations could be solved numerically with double precision, some could be solved numerically with lower precisions, and some could be approximated with a surrogate model learned via a data-driven approach, such as the ones studied in this paper.

A   More details on RC-ESN

Here we show a comparison of RC-ESN forecast skills with nonlinear transformation algorithms T_1 [?], T_2 and T_3, and without any transformation between $r(t)$ and $\tilde{r}(t)$. The three algorithms are (for $i = 1, 2, 3 \dots N$ and $j = 1, 2, 3 \dots D$)

Algorithm T_1
$\tilde{r}_{i,j} = r_{i,j} \times r_{i,j}$ if j is odd
$\tilde{r}_{i,j} = r_{i,j}$ if j is even

Algorithm T_2
$\tilde{r}_{i,j} = r_{i,j-1} \times r_{i,j-2}$ if $j > 1$ is odd
$\tilde{r}_{i,j} = r_{i,j}$ if j is 1 or even

Algorithm T_3

Figure 10: Comparison of RC-ESN forecast skills with nonlinear transformation algorithms T_1, T_2, and T_3 (cyan, green, and blue lines, respectively) and without any transformation (red line). Black line shows the truth. (A) Prediction from one initial condition, (B) Relative L_2-norm error averaged over 100 initial condition $e(t)$ (Eq. (9)). Time is in MTU where 1 MTU $= 200\Delta t \approx 4.5/\lambda_{max}$.

$$\tilde{r}_{i,j} = r_{i,j-1} \times r_{i,j+1} \qquad \text{if } j > 1 \text{ is odd}$$
$$\tilde{r}_{i,j} = r_{i,j} \qquad\qquad\qquad \text{if } j \text{ is 1 or even}$$

Figure 10(A) shows an example of short-term predictions from an initial condition using T_1, T_2, T_3, and no transformation, everything else kept the same. It is clear that the nonlinear transformation is essential for skillful predictions, as the prediction obtained without transformation diverges from the truth before 0.25 MTU. The three nonlinear transformation algorithms yield similar results, with accurate predictions for more than 2 MTU. Figure 10(B), which shows the relative prediction error averaged over 100 initial conditions, further confirms this point.

Why the nonlinear transformation is needed, and the best/optimal choice for the transformation (if it exists) should be studied in future work. We highlight that the nonlinear transformation resembles basis function expansion, which is commonly used to capture nonlinearity in linear regression models [?].

B   More details on ANN

The ANN used in this study (and the one that yields the best performance) is similar to that of Duben and Bauer *et al.* [?]. The ANN has four hidden layers each of which has 100 neurons with a tanh activation function. The input to the ANN has 8 neurons which takes $X(t)$ as the input and outputs $\Delta X(t)$ which is also a 8-neuron layer. The obtained $\Delta X(t)$ output is then used in an Adam-Bashforth integration scheme to obtain $X(t + \Delta t)$. We found that a simpler Euler scheme yields similar results. However, we found that training/testing on pairs of $X(t)$ and $X(t+\Delta t)$ (that was used for RC-ESN) leads to no prediction skill with ANN, and that following the procedure used in Dueben and Bauer [?] is essential for skillful predictions. We speculate that this might be due to the stateless nature of the ANN. By relating $X(t)$ to the change in $X(t)$ at the next time step, rather than the raw value $X(t + \Delta t)$, this ANN training architecture implicitly contains a reference to the previous time step. While the ANN itself is stateless, this particular training approach essentially encodes a first-order temporal dependence between successive states.

Note that our approach here is the same as the global ANN of Duben and Bauer *et al.* [?]. We also tried their local ANN approach, but consistent with their findings for the Lorenz system, found better performance with the global approach (results not reported for brevity).

The ANN is trained with a stochastic gradient descent algorithm with a learning rate of 0.001 with a batch size of 100 and mean absolute error as loss function (mean squared error also gives similar performance).

C   More details on RNN-LSTM

The governing equations for RNN-LSTM are:

$$g^f(t) = \sigma_f\left(\mathbf{W_f}\left[h(t-1), I(t)\right] + b_f\right), \tag{11}$$

$$g^i(t) = \sigma_f\left(\mathbf{W_i}\left[h(t-1), I(t)\right] + b_i\right), \tag{12}$$

$$\tilde{C}(t) = \tanh\left(\mathbf{W_c}\left[h(t-1), I(t)\right] + b_h\right), \tag{13}$$

$$C(t) = g^f(t)C(t-1) + g^i(t)\tilde{C}(t), \tag{14}$$

$$g^o(t) = \sigma_h\left(\mathbf{W_h}\left[h(t-1), I(t)\right] + b_h\right), \tag{15}$$

$$h(t) = g^o(t)\tanh\left(C(t)\right), \tag{16}$$

$$\Phi(t) = \mathbf{W_{oh}}h(t), \tag{17}$$

$$X(t+\Delta t) \approx \Phi(t). \tag{18}$$

σ_f is the softmax activation function; $g^f(t)$, $g^i(t)$, and $g^o(t) \in \Re^{d_h \times (d_h+d_i)}$ are the forget gate, input gate, and output gate respectively. d_h is the dimension of the hidden layers (chosen as 50 in our study) and d_i is the dimension of the input $(8 \times q)$. b_f, b_i, and b_h are the biases in the forget gate, input gate, and the hidden layers. $I(t) \in \Re^{d_i}$ is the input, which is a column of a time-delay-embedded matrix of $X(t)$. This matrix has the dimension of $(8 \times q) \times N$. $h(t) \in \Re^{d_h}$ is the hidden state and $C(t) \in \Re^{d_h}$ is the cell state (the states track the temporal evolution). The weights $\mathbf{W_o}$, $\mathbf{W_i}$, $\mathbf{W_f}$, $\mathbf{W_c}$, and $\mathbf{W_{oh}}$ are learned through the BPTT algorithm [?] using an ADAM optimizer [?]. $\Phi(t)$ is the output from the RNN-LSTM.

The LSTM used in this study is a stateless LSTM, wherein the two hidden states (C and h) are refreshed at the beginning of each batch during training (this is not the same as the stateless ANN, where there is no hidden state at all). Here, the training sequences in each batch are shuffled randomly, leading to an unbiased gradient estimator in the stochastic gradient descent algorithm [?].