

University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

Mechanical Engineering Faculty Publications
and Presentations

College of Engineering and Computer Science

4-6-2022

Model-assisted deep learning of rare extreme events from partial observations

Anna Asch

Ethan Brady

Hugo Gallardo

The University of Texas Rio Grande Valley

John Hood

Bryan Chu

See next page for additional authors

Follow this and additional works at: https://scholarworks.utrgv.edu/me_fac



Part of the [Mechanical Engineering Commons](#)

Recommended Citation

Chaos 32, 043112 (2022); <https://doi.org/10.1063/5.0077646>

This Article is brought to you for free and open access by the College of Engineering and Computer Science at ScholarWorks @ UTRGV. It has been accepted for inclusion in Mechanical Engineering Faculty Publications and Presentations by an authorized administrator of ScholarWorks @ UTRGV. For more information, please contact justin.white@utrgv.edu, william.flores01@utrgv.edu.

Authors

Anna Asch, Ethan Brady, Hugo Gallardo, John Hood, Bryan Chu, and Mohammad Farazmanda

Model-assisted deep learning of rare extreme events from partial observations

Cite as: Chaos **32**, 043112 (2022); <https://doi.org/10.1063/5.0077646>

Submitted: 04 November 2021 • Accepted: 21 March 2022 • Published Online: 06 April 2022

Anna Asch, Ethan J. Brady, Hugo Gallardo, et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[A tighter generalization bound for reservoir computing](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **32**, 043115 (2022); <https://doi.org/10.1063/5.0082258>

[Direct statistical simulation of the Lorenz63 system](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **32**, 043111 (2022); <https://doi.org/10.1063/5.0075580>

[Oscillation quenching in diffusively coupled dynamical networks with inertial effects](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **32**, 041102 (2022); <https://doi.org/10.1063/5.0087839>

APL Machine Learning

Open, quality research for the networking communities

MEET OUR NEW EDITOR-IN-CHIEF

LEARN MORE



Model-assisted deep learning of rare extreme events from partial observations

Cite as: Chaos 32, 043112 (2022); doi: 10.1063/5.0077646

Submitted: 4 November 2021 · Accepted: 21 March 2022 ·

Published Online: 6 April 2022



View Online



Export Citation



CrossMark

Anna Asch,¹ Ethan J. Brady,² Hugo Gallardo,³ John Hood,⁴  Bryan Chu,⁵  and Mohammad Farazmand^{5,a)} 

AFFILIATIONS

¹Department of Mathematics, Cornell University, 310 Malott Hall, Ithaca, New York 14853, USA

²Department of Mathematics, Purdue University, 150 N. University Street, West Lafayette, Indiana 47907, USA

³Department of Mechanical Engineering, The University of Texas Rio Grande Valley, 1201 W. University Drive, Edinburg, Texas 78539, USA

⁴Department of Mathematics, Bowdoin College, 8600 College Station Brunswick, Maine 04011, USA

⁵Department of Mathematics, North Carolina State University, Raleigh, North Carolina 27695-8205, USA

^{a)}Author to whom correspondence should be addressed: farazmand@ncsu.edu

ABSTRACT

To predict rare extreme events using deep neural networks, one encounters the so-called small data problem because even long-term observations often contain few extreme events. Here, we investigate a model-assisted framework where the training data are obtained from numerical simulations, as opposed to observations, with adequate samples from extreme events. However, to ensure the trained networks are applicable in practice, the training is not performed on the full simulation data; instead, we only use a small subset of observable quantities, which can be measured in practice. We investigate the feasibility of this model-assisted framework on three different dynamical systems (Rössler attractor, FitzHugh–Nagumo model, and a turbulent fluid flow) and three different deep neural network architectures (feedforward, long short-term memory, and reservoir computing). In each case, we study the prediction accuracy, robustness to noise, reproducibility under repeated training, and sensitivity to the type of input data. In particular, we find long short-term memory networks to be most robust to noise and to yield relatively accurate predictions, while requiring minimal fine-tuning of the hyperparameters.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0077646>

Deep learning has proven largely effective in predicting chaotic dynamical systems. However, to predict rare extreme events, one is confronted with the so-called small data problem: even long-term observations lack sufficient extreme events for training purposes. We introduce a model-assisted framework where the deep neural network is trained using long-term simulations with adequate sampling from the extreme event regime. However, to ensure the trained network can be used with observational data, input variables only contain partial information comprising the quantities that can be measured in experiments. We examine the feasibility of this model-assisted framework on three neural network architectures, each trained with partial observations from three different dynamical systems.

I. INTRODUCTION

Extreme events such as rogue waves, earthquakes, epileptic seizures, and stock market crashes are rare but have

devastating humanitarian, environmental, and financial consequences. Real-time prediction of these events is essential for optimal response management and mitigation of their most adverse consequences.^{1–3}

Although deep learning has proven largely effective in predicting chaotic systems, to predict extreme events, one is confronted with the *small data problem*: even long-term observations contain relatively few extreme events. As a result, the available observational data lack enough samples from the extreme event regime for training purposes. In addition, the full state of the system is often inaccessible in practice. This is especially the case for spatiotemporal systems where the state can only be measured at sparse sensor locations.^{4–6} Therefore, the training must be carried out using a limited set of system observables.

Here, we consider a model-assisted framework that addresses these sampling and partial observation issues. As depicted in Fig. 1, the proposed framework has two components:

1. Off-line: A computationally expensive component, which is carried out off-line and only once. This step involves

long-term numerical simulations and training a deep neural network.

2. Real-time: A computationally inexpensive component that uses observational data as input for the pre-trained neural network to make real-time predictions.

In the off-line step, we use a mathematical model of the system to generate long-term simulations, which contain enough extreme event samples to train a deep neural network. Although the simulation results contain the full state of the system $\mathbf{x}(t)$, we do not use this complete information for training the neural network. Instead, we use this information to generate time series of the observables $p_i(t)$, i.e., a set of quantities that can be actually measured in practice. The observable time series are then used to train a deep neural network. The trained network uses the observations $p_i(t)$ as input to predict the future values of a quantity of interest $q(t + \tau)$, where τ denotes the prediction time. The quantity of interest (QoI) refers to a scalar quantity, which is relevant to extreme events. For instance, in the case of ocean rogue waves, QoI is the maximum surface height,⁷ or in climatology, QoI may refer to the mean surface temperature.⁸

Once trained, the deep neural network is used in the real-time component to make predictions using the available observations $p_i(t)$. This component only relies on observational data and is quite fast since it uses a pre-trained network. We emphasize that the off-line component is designed to address the practical issues faced in predicting extreme events; mainly that in practice not all degrees of freedom can be measured (incomplete observations). Furthermore, the long-term numerical simulations provide enough sampling from extreme events to overcome the small data problem. Although we do not do so here, one can use rare event sampling techniques to reduce the amount of numerical simulations required.^{9–11}

We examine the feasibility of this framework (see Fig. 1) for predicting extreme events by considering three different dynamical systems and three different neural network architectures: feedforward (FF), long short-term memory (LSTM), and reservoir computing (RC). For different architectures, we study their accuracy, robustness to noise, sensitivity to hyperparameters, reproducibility under repeated training, and dependence on the input data.

A. Related work

As applications of machine learning have proliferated, they have been extensively used in predicting, analyzing, and discovering chaotic dynamical systems. We refer to Refs. ^{12–16} for recent reviews of applications of machine learning methods in dynamical systems. Here, we only review studies that use machine learning in the context of extreme events.

Ding *et al.*¹⁷ design a hybrid recurrent neural network structure for extreme event prediction that addresses the scarcity of extreme events by introducing a specialized loss function inspired by extreme value theory. Qi and Majda¹⁸ use a relative entropy loss function together with a convolutional neural network to predict extreme events in a truncated Korteweg–de Vries equation (see also Rudy and Sapsis¹⁹). Senthilvelan and co-workers²⁰ use standard deep learning methods to predict extreme events in a parametrically driven mechanical system (see also Refs. ^{21–23}).

The above studies rely on the complete system state for their prediction tasks. However, such complete information is often unavailable in practice. One remedy is to use reduced-order models that only rely on a subset of the state variables. For instance, Wan *et al.*²⁴ demonstrate that incorporating physical principles and reduced-order modeling in the neural network architecture can improve predictive accuracy and interpretability. Although reduced-order methods involve less information than the full system state, they still require solving a model which is not ideal for fast, real-time predictions.

There are relatively few studies that examine the capacity of deep learning for predicting extreme events based only on observational data. Chattopadhyay *et al.*²⁵ make model-free prediction of extreme-causing weather patterns using large-scale circulation data and surface temperature to train a capsule neural network (CapsNet). They find that supplementing the circulation data with surface temperature significantly enhances the prediction accuracy. Rudy and Sapsis²⁶ use an LSTM network to predict intermittent aerodynamic fluctuations from discrete pressure measurements on an airfoil. Chattopadhyay *et al.*²⁷ quantify the statistics of a Lorenz 96 model using only the slow variables of the system as input data. They find that LSTM and RC networks predict the heavy-tailed statistics (i.e., rare events) reasonably well, while FF networks fail to capture the tail. Pyragas and Pyragas²⁸ use RC networks to predict and mitigate extreme events in a FitzHugh–Nagumo system from a single quantity containing global information.

Here, to mimic real applications, we also assume that only a small number of observable time series are available as network inputs. Unlike previous studies that often consider one dynamical system or one network architecture, we conduct a comprehensive study by applying three different neural network architectures to three different dynamical systems that exhibit extreme events. In particular, we seek to address the following questions: (1) Accuracy: Is one neural network architecture more skillful in predicting extreme events across different dynamical systems? (2) Sensitivity to noise: Is a particular architecture more robust to observational noise? Furthermore, does training on clean simulation data outperform training on noisy data. (3) Sensitivity to hyperparameters: Which network requires less hyperparameter fine-tuning. (4) Reproducibility: Are the results reproducible under retraining? (5) Sensitivity to input data: Do the networks perform equally well when trained on different types of input data?

B. Outline of the paper

This paper is organized in the following manner. In Sec. II, we introduce the setup and notation and outline the dynamical systems studied here. Section III details the three deep learning structures we use. In Sec. IV, we present our numerical results. Section V contains our concluding remarks.

II. PRELIMINARIES AND SETUP

In this section, we introduce the dynamical systems setup for extreme events. In particular, we discuss three systems that are used later in this paper to demonstrate our results.

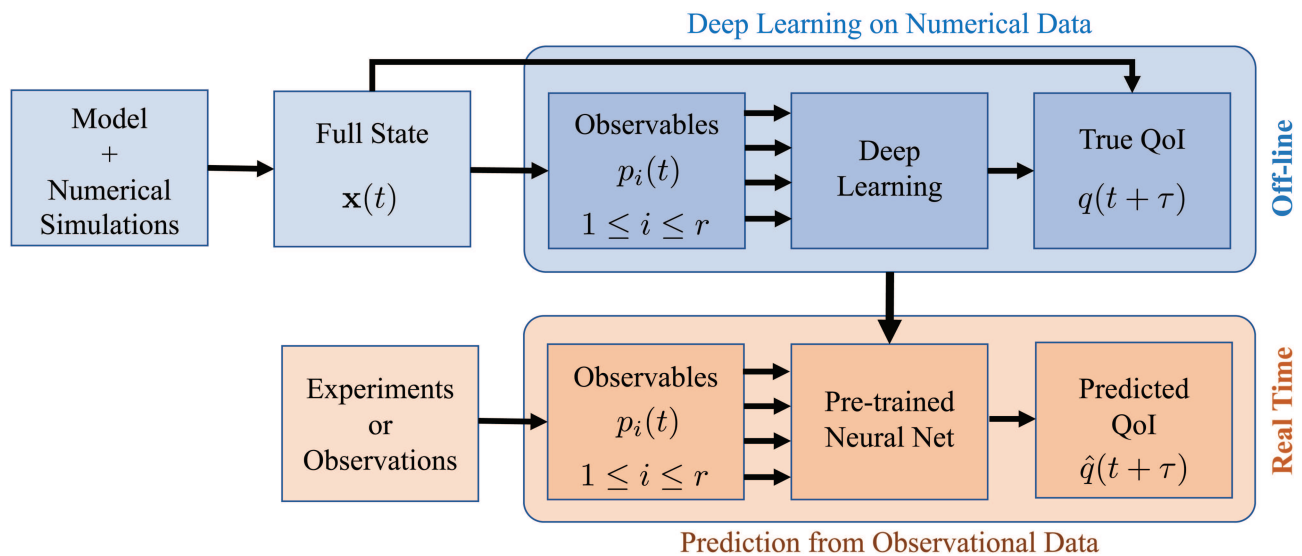


FIG. 1. Schematic diagram of the program followed in this paper. We use a two-phase procedure to make predictions. In the off-line phase, numerical simulations of a model are used to train a neural network. The trained neural network is then used in real-time to predict extreme events from observations.

A. Dynamical systems setup

We consider autonomous dynamical systems defined by a set of ordinary differential equations (ODEs),

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (1)$$

where the state $\mathbf{x}(t) = (x_1(t), x_2(t), \dots, x_n(t))$ belongs to \mathbb{R}^n for all $t \geq 0$. The vector field $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a potentially nonlinear map. ODE (1) may model a finite-dimensional system or arise from a finite-dimensional approximation of a partial differential equation (PDE), as is common in numerical discretization of PDEs. We denote the solution map of the system by $S^t : \mathbb{R}^n \rightarrow \mathbb{R}^n$, which maps an initial condition \mathbf{x}_0 to its time- t state $\mathbf{x}(t)$.

We consider a scalar quantity of interest $q : \mathbb{R}^n \rightarrow \mathbb{R}$ whose evolution is related to extreme events. This quantity is problem dependent; for instance, for rogue waves, the quantity of interest is the maximum wave height⁷ while, in turbulence, it may be energy dissipation.^{29,30} Evaluated along a trajectory $\mathbf{x}(t)$ of the system, the quantity of interest generates the time series $q(t) := q(\mathbf{x}(t))$. An extreme event refers to an episode where the quantity of interest is unusually large. More specifically, we say an extreme event has taken place if the quantity of interest $q(t)$ exceeds a prescribed threshold q_e . The case where extreme events correspond to unusually small values of QoI can be handled similarly by redefining the quantity as $q \mapsto -q$.

We seek to predict extreme events before they take place. For a prediction time τ , we use the available information up to the present time t to predict the future value of the QoI $q(t + \tau)$. Of course, if the full state $\mathbf{x}(t)$ is accessible, the ODE can be integrated numerically to estimate $q(t + \tau) = q(S^\tau(\mathbf{x}(t)))$.

However, in practice, we often lack complete information about the full system state $\mathbf{x}(t)$. Instead, we assume that only *partial*

observations $\mathbf{p} = (p_1, p_2, \dots, p_r) \in \mathbb{R}^r$, containing r unique measurements, are available where $r \ll n$. The i th component of the observations is given by the map $p_i : \mathbb{R}^n \rightarrow \mathbb{R}$, and we write $p_i(t) = p_i(\mathbf{x}(t))$ for notational simplicity. We emphasize that although the full state $\mathbf{x}(t)$ is not available, we assume the observables p_i are measurable quantities. These observations may contain certain coordinates of the system state \mathbf{x} or, more generally, may be linear or nonlinear functions of the state.

The partial observations $\mathcal{U} := \{\mathbf{p}(s) : s \leq t\}$ up to time t are used to predict the QoI $q(t + \tau)$ at the future time $t + \tau$. To this end, we seek to learn a map $N : \mathcal{U} \rightarrow \mathbb{R}$ that predicts the QoI in the future,

$$\hat{q}(t + \tau) = N[\mathbf{p}(s), s \leq t]. \quad (2)$$

We denote the predicted QoI by \hat{q} . In Sec. III, we detail the machine learning methods used to learn the map N .

In practice, we realize a solution to Eq. (1) as a discrete time series obtained by numerical integration or experimental measurements. For some small time step Δt , the discrete time series for the state $\mathbf{x}(t)$ is denoted by $\{\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)}, \dots\}$, where $\mathbf{x}^{(k)} = \mathbf{x}(k\Delta t)$. Similarly, for the observations \mathbf{p} and the quantity of interest q , we have the discrete time series $\{\mathbf{p}^{(0)}, \mathbf{p}^{(1)}, \dots, \mathbf{p}^{(k)}, \dots\}$ and $\{q^{(0)}, q^{(1)}, \dots, q^{(k)}, \dots\}$, where $\mathbf{p}^{(k)} = \mathbf{p}(\mathbf{x}^{(k)})$ and $q^{(k)} = q(\mathbf{x}^{(k)})$.

B. Assumptions and limitations

The proposed framework uses simulation data to pre-train a neural network, which is subsequently used for prediction from observational data (see Fig. 1). Some consistency assumptions have to be made in order to ensure that the pre-trained network is transferable to observational data. For instance, if the distribution of the simulation data is different from the observations, the trained network may fail to accurately predict the out-of-sample events in the

observations. To ensure consistency between the simulation data and the observational data, we require the following conditions: (i) Statistical stationarity: We assume that the system's attractor is statistically stationary so that the distribution of the simulation data and the observations are in agreement. (ii) Constant sampling rate: We assume that the time series $\{\mathbf{p}(s) : s \leq t\}$ are sampled at regular intervals Δt and that the observational data can be measured with the same frequency.

Assumption (i) ensures that the dynamics has reached a stationary distribution, which is common between the simulation data and the observations. Note that here we assume that the dynamical system (1) accurately models the system; therefore, the only discrepancy between the simulation and observation distributions can arise from observational noise, which does not significantly alter the distribution. As a result of assumption (i), the proposed framework is not applicable to non-autonomous systems where the vector field \mathbf{F} depends explicitly on time. We note that the system does not necessarily need to be ergodic; we only require that the components of the system attractor are adequately sampled.

Assumption (ii) ensures that the learned map N , which maps the input time series $\mathcal{U} = \{\mathbf{p}(s) : s \leq t\}$ to the QoI $q(t + \tau)$, is applicable to observations. The map N depends implicitly on the sampling time step Δt . In other words, a network N that maps input time series sampled at Δt intervals is different from the network N' that maps samples which are $\Delta t' (\neq \Delta t)$ apart. Therefore, if the simulation data are sampled at a different rate than the observational data, the pre-trained network cannot be applied to make predictions from observations. As a result of assumption (ii), when setting up the training data from simulations, one should ensure that the sampling rate Δt matches the frequency of available observations. We emphasize that the prediction time τ does not need to be equal to the sampling time step Δt .

Finally, we note that statistical stationarity is an inherent property of the dynamical system whereas constant sampling rate is related to the manner in which the system is observed.

C. Three systems studied

We consider three dynamical systems of increasing complexity: the Rössler system, the FitzHugh–Nagumo (FHN) system, and the Kolmogorov flow (KF). Figure 2 shows these systems, their observables, and the quantities of interest. Here, we briefly review each system and its significance; a detailed description of each system is provided in Appendix A.

1. Rössler: The Rössler system is a three-dimensional ODE routinely used as a prototypical model for extreme events. Denoting the state variable with $\mathbf{x} = (x_1, x_2, x_3)$, the extreme events in this system appear as intermittent chaotic bursts in its x_3 component (see Fig. 2). Here, we consider this system with x_1 and x_2 components of the state as our observables, and $q = x_3$ as the quantity of interest.
2. FitzHugh–Nagumo: The FitzHugh–Nagumo equations model excitable systems such as neural and cardiac activity. Here, we consider a discrete version of the FHN model consisting of $n = 101$ diffusively coupled units.³¹ Each unit i consists of two variables (v_i, w_i) leading to a 202 dimensional system. We

assume only the variables of the first unit (v_1, w_1) are observable. The QoI is the mean voltage $q = (\sum_i v_i)/n$.

3. Kolmogorov flow: The Kolmogorov flow refers to the Navier–Stokes equations with periodic boundary conditions and a sinusoidal shear forcing. At high enough Reynolds numbers, the system exhibits extreme events in the form of chaotic bursts of the energy dissipation rate.³² Here, we consider two sets of observables; one is a particular Fourier mode which was recently discovered to play a major role in the formation of extreme events.³³ The second set of observables is the vorticity field measured at a few discrete points. We investigate the performance of deep learning methods for each set of observables (i.e., network inputs). In both cases, the QoI is the energy dissipation rate.

III. MACHINE LEARNING ARCHITECTURES

We use artificial neural networks to predict upcoming extreme events given the available information about the system. More specifically, we train each neural network to learn the predictor N , which takes partial observations $\mathcal{U} = \{\mathbf{p}(s), s \leq t\}$ from the system as input and returns, as output, the predicted value of the quantity of interest, τ time units into the future, i.e., $\hat{q}(t + \tau)$.

We compare the performance of three deep neural network structures: feedforward (FF), long short-term memory (LSTM), and reservoir computing (RC) networks. Each neural network depends on many hyperparameters, which have to be prescribed before training the network. It is well-known that the choice of the hyperparameters can drastically alter the network performance. Here, we choose these hyperparameters after an extensive trial-and-error search in the hyperparameter space as listed in Table 1. The optimal set of hyperparameters depends on the system (Rössler, FitzHugh–Nagumo, or Kolmogorov flow) as detailed in Appendix B. Here, “optimal” refers to the best combination of hyperparameters among those we tested and not optimal over all possible hyperparameters. In the following, we briefly review each neural network architecture and its hyperparameters.

A. Feedforward neural networks

We use a fully connected feedforward neural network as shown in Fig. 3. Each node in the neural network takes the outputs from the previous layer and transforms them with a nonlinear mapping. Within a particular layer, the j th node takes the outputs of the previous layer x_i , for $1 \leq i \leq l$ and returns $\sigma(\sum_{i=1}^l w_{ij}x_i + b_j)$, where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a sigmoidal activation function. The weights w_{ij} and bias b_j are trained using backpropagation and stochastic gradient descent. This nonlinear composition continues until the network reaches the final (or output) layer. The output layer has its own bias b_o and activation function σ_o .

We use the `feedforwardnet` function from MATLAB's Deep Learning Toolbox to implement our FF neural networks for time-series prediction. We train our networks for 1000 epochs using the Levenberg–Marquardt algorithm, a variation of gradient descent, to minimize the loss.³⁴ The loss function is the mean

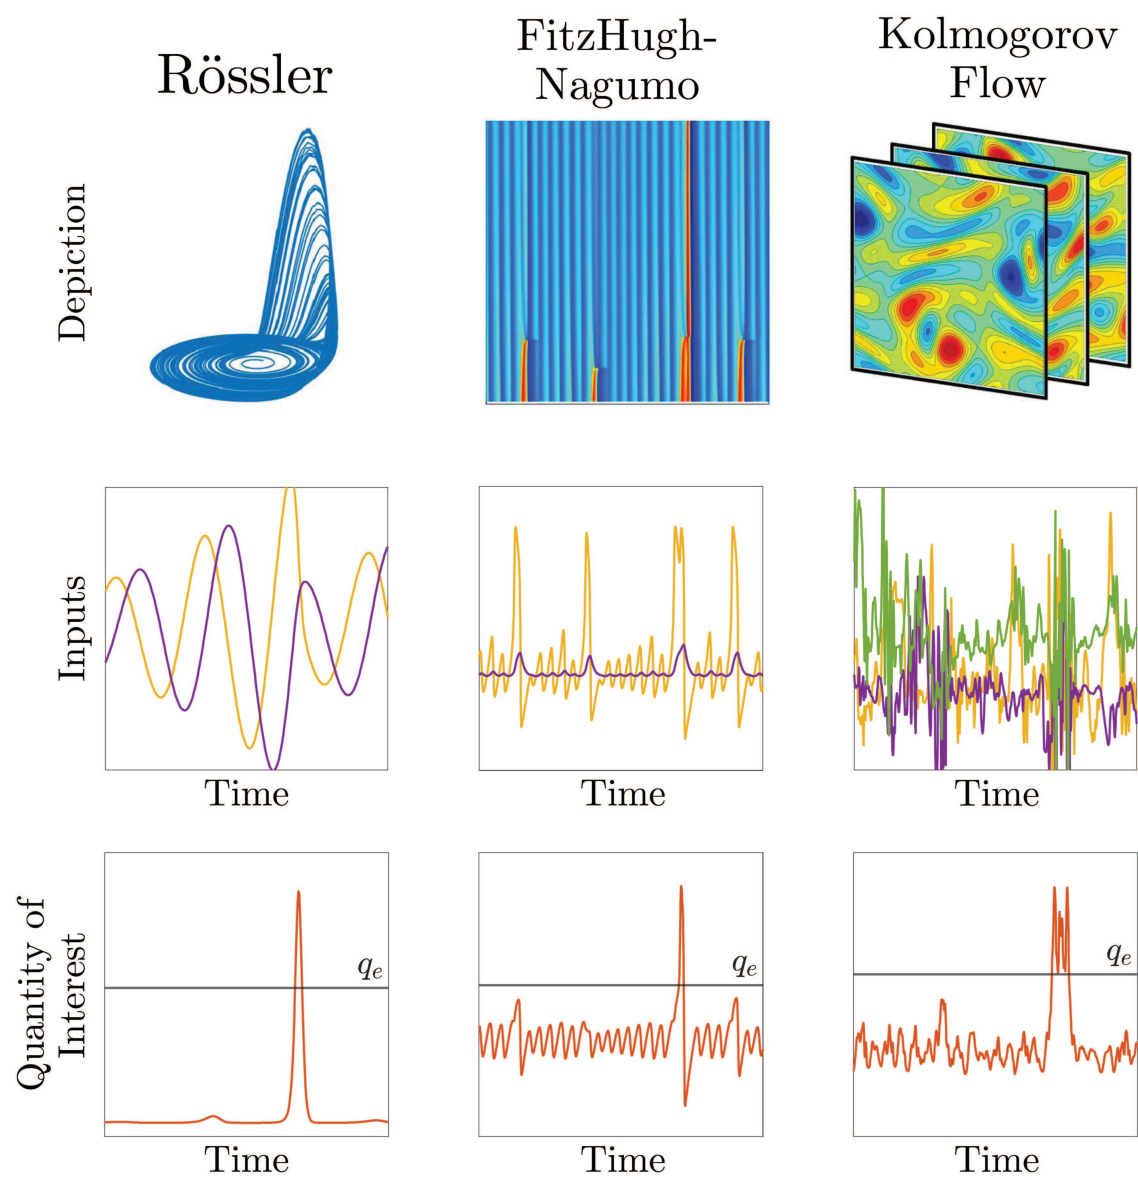


FIG. 2. The three dynamical systems studied. The complexity of the systems increases from left to right.

TABLE I. The list of fine-tuned hyperparameters for each neural network.

	Feedforward	LSTM	Reservoir computing
Hyperparameters	Number of layers Number of nodes Activation functions Number of delays Delay time	Hidden units Number of layers	Number of nodes Spectral radius Leaking rate Input density Reservoir density Regularization weight




FIG. 3. The FF network architecture with a single output and input data of size n . There are m layers and l nodes per layer.

squared error,

$$MSE = \frac{1}{K} \sum_{i=1}^K (q(t_k + \tau) - \hat{q}(t_k + \tau))^2. \quad (3)$$

where q (respectively, \hat{q}) denotes the true value of the quantity of interest (respectively, the predicted value of the quantity of interest). Here, K is the number of predictions made. We choose 15% of the original training data as validation data and terminate training after validation error increases over six consecutive epochs.

To determine the optimal hyperparameters, we consider between 2 and 6 layers with 4–64 nodes per layer. In addition, we consider log-sigmoidal, tanh-sigmoidal, and linear activation functions. Two additional hyperparameters correspond to time delays in our input data, as discussed in Sec. III A 1. The optimal combination of these hyperparameters for each dynamical system is discussed in Appendix B.

1. Time delay embedding

For FF networks, predicting the QoI $\hat{q}(t + \tau)$ from the instantaneous observations $\mathbf{p}(t)$ does not return accurate results. Motivated by Takens' embedding theorem,³⁵ we introduce time delays to take into account the history of the observations. The time delay embedding collects several previous observations from the time series $\{\mathbf{p}(s), s \leq t\}$. More precisely, for a prescribed number of delays $m \in \mathbb{N}$ and a delay time $s > 0$, the time delay embedding $\rho \in \mathbb{R}^{r \cdot m}$ is defined as

$$\rho(t; m, s) = [\mathbf{p}(t), \mathbf{p}(t-s), \dots, \mathbf{p}(t-(m-1)s)]. \quad (4)$$

The special case $m = 1$ corresponds to instantaneous observations, $\rho(t; m, s) = \mathbf{p}(t)$. The function learned using the FF neural network is, therefore, a map $N: \mathbb{R}^{r \cdot m} \rightarrow \mathbb{R}$ from past observations up to the current time t to the predicted quantity of interest $\hat{q}(t + \tau)$.

The values for m and s are chosen differently for each system as detailed in Appendix B. The number of delays m depends

on the dimension of the system attractor whereas the delay time s depends on the decorrelation time of the system. Arbitrarily increasing m is unfeasible due to training time constraints and s must be chosen carefully so that there is little redundancy between delayed observations.

LSTM and RC are recurrent neural networks, which implicitly take the history into account. As a result, time delay embedding is unnecessary for training LSTM and RC networks and is only used for training the FF neural network. In fact, we have observed that including the time delays in the input data explicitly tends to deteriorate the performance of the LSTM and RC networks.

B. Long short-term memory neural networks

Long short-term memory (LSTM) networks are a particular type of recurrent neural network (RNN), which use gates to control the effect of past history on the current state. This construction allows them to capture long-term dependencies in the time series.^{36,37}

Figure 4 illustrates our LSTM architecture consisting of LSTM layers followed by a fully connected network. The networks use an input layer of dimension n , followed by m LSTM layers, followed by a fully connected layer ending in an output layer of dimension one. The fully connected and output layers in Fig. 4 are identical to their FF neural network counterparts; the new LSTM layers enable the network to “remember” past history.

Each LSTM layer takes a time series and, after iterating through each time step, returns a reconstructed time series composed of significant details of the network remembered through training. At each time step t , the cell combines the time step information with the remembered information. The previous cell determines the remembered information and the modifications to make, known as the cell state c_{t-1} and the hidden state h_{t-1} , respectively. The number of modifications made, known as hidden units, is a hyperparameter.

Each LSTM node uses three gates to decide which information should be retained or discarded in producing the output

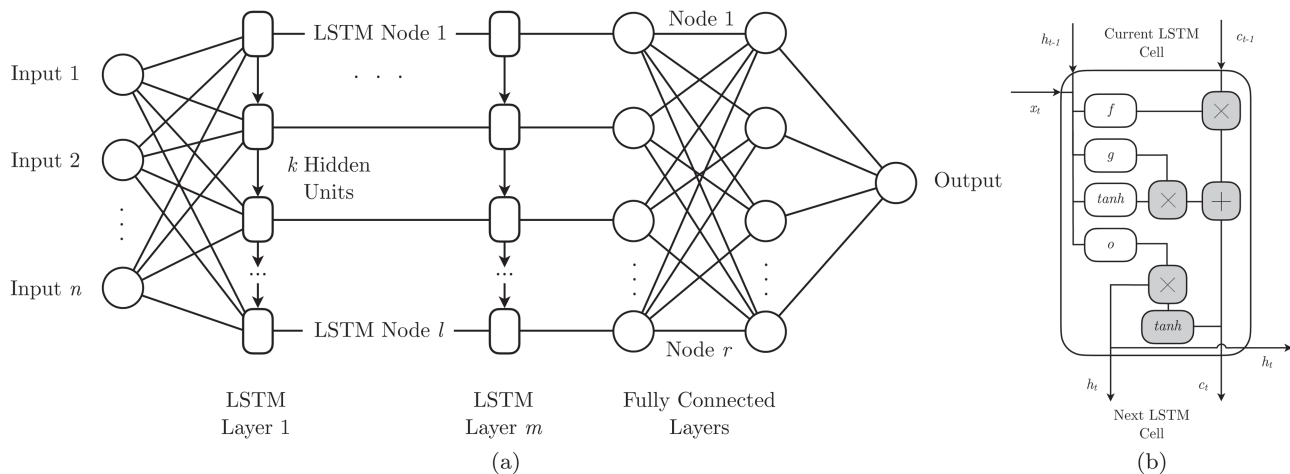


FIG. 4. (a) LSTM architecture with m layers followed by a fully connected network. (b) An LSTM cell.

[see Fig. 4(b)]. The forget gate f determines which information the LSTM block removes from the layer's memory. The input gate and cell candidate, denoted by \tanh and g , determine what information from the input data x_t the LSTM block incorporates to update the cell state in the LSTM node. The output gate o determines how the cell state c_{t-1} contributes to the output relative to h_{t-1} and x_t .

Each gate in the LSTM node has two weight matrices and a bias. One weight matrix is associated with the input x_t while the other weight matrix is associated with the hidden state h_{t-1} . In the training process, the weights and biases of both the LSTM network and the fully connected network are trained by backpropagation.

We implement LSTM networks for sequence-to-sequence regression using MATLAB's Deep Learning Toolbox. The two hyperparameters for LSTM networks we adjust are the number of layers and the number of hidden units per layer. We explore hyperparameter values between 1 and 3 layers and 8–200 hidden units per layer. Each network trains for 250 epochs using the Adam algorithm to minimize the loss function (3).

C. Reservoir computing

Reservoir computing networks are the second RNN we consider here. The key structure is a reservoir of nodes with prescribed connections capturing complex temporal dynamics.³⁸ Our RC network takes an input of dimension n and outputs a scalar.

As shown in Fig. 5, the reservoir architecture consists of two couplers and a high dimensional dynamical system. The input-to-reservoir (I/R) coupler maps the input data to the reservoir, while the reservoir-to-output (R/O) coupler maps the reservoir state to an output. Here, both couplers are chosen to be linear mappings. The I/R coupler uses $\mathbf{W}_{\text{in}} \in \mathbb{R}^{m \times n}$, whose entries are sampled from a uniform distribution $U(-a, a)$ for some $a > 0$, to map the observations $\mathbf{p}(t)$ into the reservoir,

$$\mathbf{u}(t) = \mathbf{W}_{\text{in}} \mathbf{p}(t), \quad (5)$$

where $\mathbf{u}(t)$ is used to update the reservoir state $\mathbf{r}(t)$.

The entry $r_j(t)$ of the reservoir state $\mathbf{r}(t) \in \mathbb{R}^m$ corresponds to the j th reservoir node. The reservoir is a directed Erdős–Rényi network with $m \times m$ adjacency matrix \mathbf{A} describing connections between nodes. We denote the spectral radius of \mathbf{A} by $\rho \in \mathbb{R}$. The reservoir density is the percentage of nonzero entries in \mathbf{A} . The initial reservoir state $\mathbf{r}(0)$ is chosen at random. Within the reservoir, the reservoir state is updated by $\mathbf{r}(t + \Delta t) = \tanh(\mathbf{A}\mathbf{r}(t) + \mathbf{u}(t))$. The reservoir then passes the updated reservoir state $\mathbf{r}(t + \Delta t)$ through the R/O coupler via $\hat{q}(t + \tau) = \mathbf{W}_{\text{out}} \mathbf{r}(t) \in \mathbb{R}$, where $\mathbf{W}_{\text{out}} \in \mathbb{R}^{1 \times m}$ denotes the output weights. The vector \mathbf{W}_{out} is optimized during training to minimize the least squares error, with an L_2 regularization term.

We create RC networks using `easyn` library in Python.⁴⁰ As in the case of the LSTM network, our inputs do not include time delay embedding since the recurrent nature of the RC network implicitly accounts for the history.

RC has several hyperparameters. We vary the input density or the portion of nonzero entries in \mathbf{W}_{in} between 0.1 and 1. Within the reservoir, we vary the number of nodes from 100 to 5000 and the spectral radius ρ of the adjacency matrix \mathbf{A} from 0.1 to 1. The penalization weight associated with the L_2 regularization is a hyperparameter, varied from 10^{-5} to 1. Finally, a reservoir computer's leaking rate ℓ determines how frequently the reservoir updates during training: ℓ^{-1} times per time instance t . We also vary the leaking rate from 0.1 to 1. The optimal value of the hyperparameters for each dynamical system is discussed in Appendix B.

IV. RESULTS AND DISCUSSION

In this section, we use deep learning to predict extreme events in three dynamical systems: Rössler (Sec. IV B), FitzHugh–Nagumo (Sec. IV C), and Kolmogorov flow (Sec. IV D). For each system, we generate a long-term numerical simulation, which returns the full state $\mathbf{x}(t)$. However, to mimic applications where the full state is not available, we do not use $\mathbf{x}(t)$ directly for training the neural networks. Instead, we use the full state information $\mathbf{x}(t)$ to generate




FIG. 5. Reservoir computing network with n inputs and 1 output.

the corresponding time series of the observables $\mathbf{p}(t)$, as depicted in Fig. 1.

The first 75% of the observable time series is used for training the neural network and the remaining 25% is used to test the trained network. For each system, we quantify the error between each system's QoI q and the corresponding neural network prediction \hat{q} using two measures of accuracy described in Sec. IV A. For each network structure outlined in Sec. III, we optimize the hyperparameters for each dynamical system separately. We present the optimal hyperparameter combinations that we use to train each neural network structure in Appendix B.

The produced observable time series $\mathbf{p}(t)$ only contain round-off and numerical discretization errors that are negligible. In practice, however, observations are always polluted with noise. To mimic the observational noise, we artificially add noise to our synthetic test data. Noise is added proportionally to each component of the observation $\mathbf{p}(t)$ at four intensities: 0%, 5%, 10%, and 20%. More precisely, let σ_i denotes the standard deviation of the time series of the i th observation $p_i(t)$. Then, the corresponding noisy data $\tilde{p}_i(\mathbf{x}(t))$ are given by

$$\tilde{p}_i(\mathbf{x}(t)) = p_i(\mathbf{x}(t)) + \alpha \xi_i(t), \quad \xi_i(t) \sim \mathcal{N}(0, \sigma_i), \quad (6)$$

where $\xi_i(t)$ is a realization of $\mathcal{N}(0, \sigma_i)$, the normal distribution with mean 0 and standard deviation σ_i . The parameter $\alpha \in \{0, 0.05, 0.10, 0.20\}$ controls the noise intensity with $\alpha = 0$ corresponding to the clean data obtained from numerical simulations.

Figures 6–10 contain our main results. All reported results correspond to the test data. Figure 6 displays samples of the predicted times series vs their true values. Figure 7 illustrates the predictive power of the deep learning architectures for each dynamical system. Figure 9 shows the effect of varying the prediction time τ , and Fig. 10 displays how noise affects prediction accuracy. In Secs. IV B–IV D, we describe these results in detail for each dynamical system. But

first, in Sec. IV A, we describe two quantities used to measure the prediction accuracy of the trained neural networks.

A. Quantifying prediction accuracy

In order to quantify the accuracy of our extreme event predictions, we use two measures: normalized root-mean square error and area under the precision–recall curve. The former simply quantifies the accuracy of time series predictions, while the latter is more suitable for quantifying the accuracy of extreme event predictions.

We define the normalized root-mean-squared error (NRMSE),

$$\text{NRMSE} = \frac{\left[\frac{1}{K} \sum_{k=1}^K (q(t_k + \tau) - \hat{q}(t_k + \tau))^2 \right]^{1/2}}{\sigma_q}, \quad (7)$$

where σ_q is the standard deviation of the true QoI, $\{q(t_k + \tau)\}_{k=1}^K$. Note that NRMSE is the ratio of the square root of loss function (3) and the standard deviation σ_q . The normalization by σ_q allows us to make base comparisons across different dynamical systems.

NRMSE measures the deviation in our predictions from the true value with no consideration given to extreme events. As reviewed in Ref. 41, there are measures of accuracy which are better suited for extreme event prediction. Here, we quantify the accuracy of extreme event prediction using area under the precision–recall curve, or area under curve (AUC) for short. Recall that an extreme event is registered if $q > q_e$, where q_e is the extreme event threshold. As a result, the predicted QoI \hat{q} may not coincide with q , but as long as $\hat{q} > \hat{q}_e$, the extreme event is correctly predicted. Here, \hat{q}_e is the extreme event threshold according to the predicted QoI; in practice, \hat{q}_e is close to the true extreme event threshold q_e , but the two do not necessarily coincide.

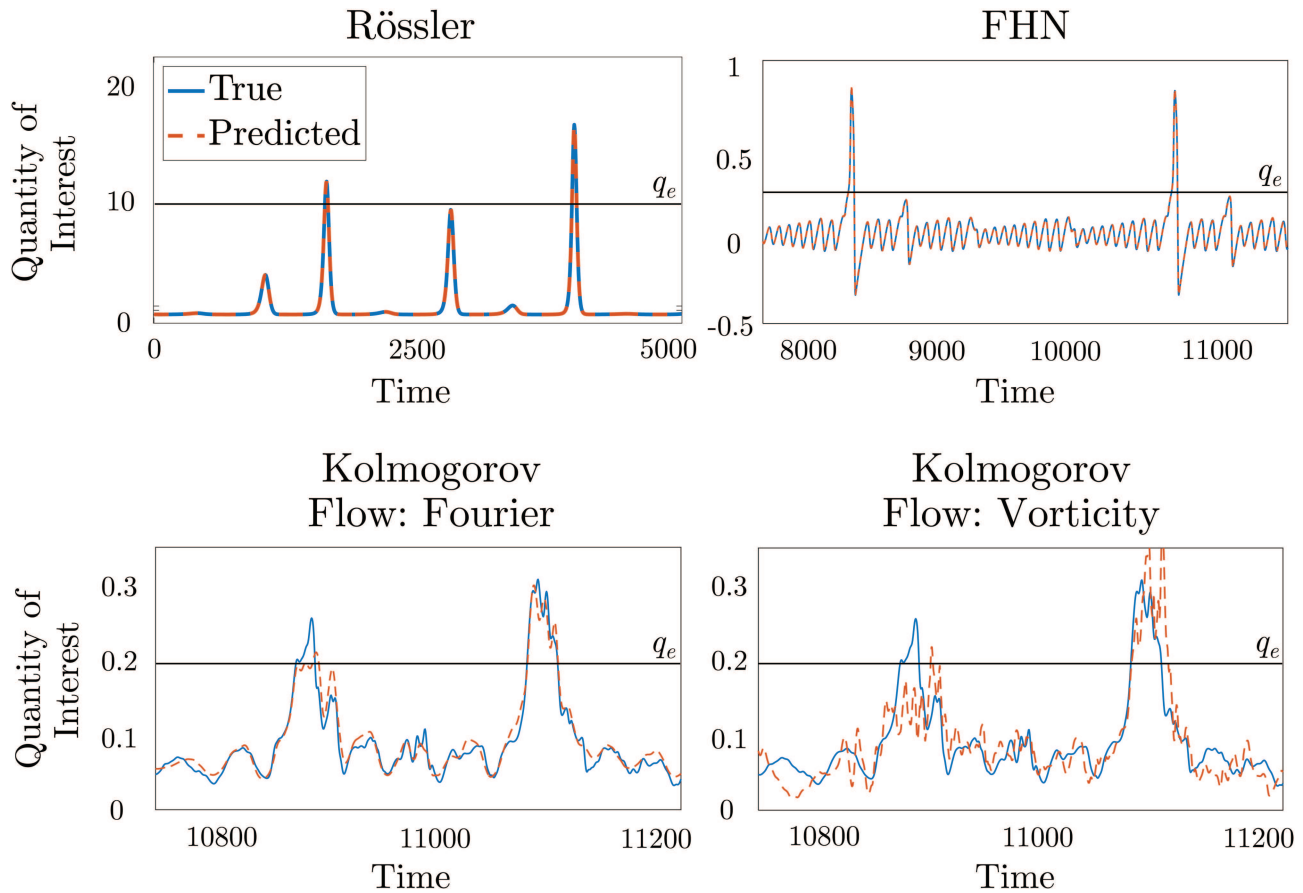


FIG. 6. Close-up views of the time series for each system's quantity of interest $q(t)$. The blue curves show the true value $q(t)$, while the dashed lines mark the best prediction $\hat{q}(t)$.

Each prediction \hat{q} is classified into one of the four ways: 1. True positive (TP): A true positive corresponds to a successful prediction of a true extreme event, i.e., $q > q_e$ and simultaneously $\hat{q} > \hat{q}_e$. 2. True negative (TN): A true negative corresponds to a correct prediction that no extreme event will take place, i.e., $q < q_e$ and simultaneously $\hat{q} < \hat{q}_e$. 3. False positive (FP): A false positive corresponds to the case where an extreme event was predicted but did not actually occur, i.e., $q < q_e$ but $\hat{q} > \hat{q}_e$. 4. False negative (FN): A false negative corresponds to the case where an extreme event occurred but the neural network failed to predict it, i.e., $q > q_e$ but $\hat{q} < \hat{q}_e$.

Precision and recall quantify extreme event prediction accuracy with regard to these four classifications. These quantities are defined by

$$\text{precision}(\hat{q}_e) = \frac{\text{TP}(\hat{q}_e)}{\text{TP}(\hat{q}_e) + \text{FP}(\hat{q}_e)}, \quad (8a)$$

$$\text{recall}(\hat{q}_e) = \frac{\text{TP}(\hat{q}_e)}{\text{TP}(\hat{q}_e) + \text{FN}(\hat{q}_e)}. \quad (8b)$$

For a given prediction threshold \hat{q}_e , precision measures the ratio of successfully predicted extreme events to the total number of predicted extreme events whereas recall measures the ratio of successfully predicted extreme events to the total number of true extreme events. Since extreme events are rare, precision and recall more accurately capture the prediction skill than TP and TN rates alone.^{41–43} In the best case scenario, where no false predictions are made, both precision and recall are 1.

We combine both these quantities to create the precision–recall curves, parameterized by the threshold \hat{q}_e . Varying the threshold \hat{q}_e allows us to examine how precision and recall change under a moving goalpost. Note that, if \hat{q}_e is very small, most extreme events will be correctly predicted leading to a large number of TPs, but at the same time there will be a large number of FPs resulting in low precision. Conversely, if \hat{q}_e is too large, then the number of FN predictions will be large resulting in low recall. A reliable predictor must have precision and recall close to 1 for a wide range of prediction thresholds \hat{q}_e . One way to quantify this range is to compute the area under the precision–recall curve, or AUC for short. The AUC is the area the precision–recall curve encompasses in the

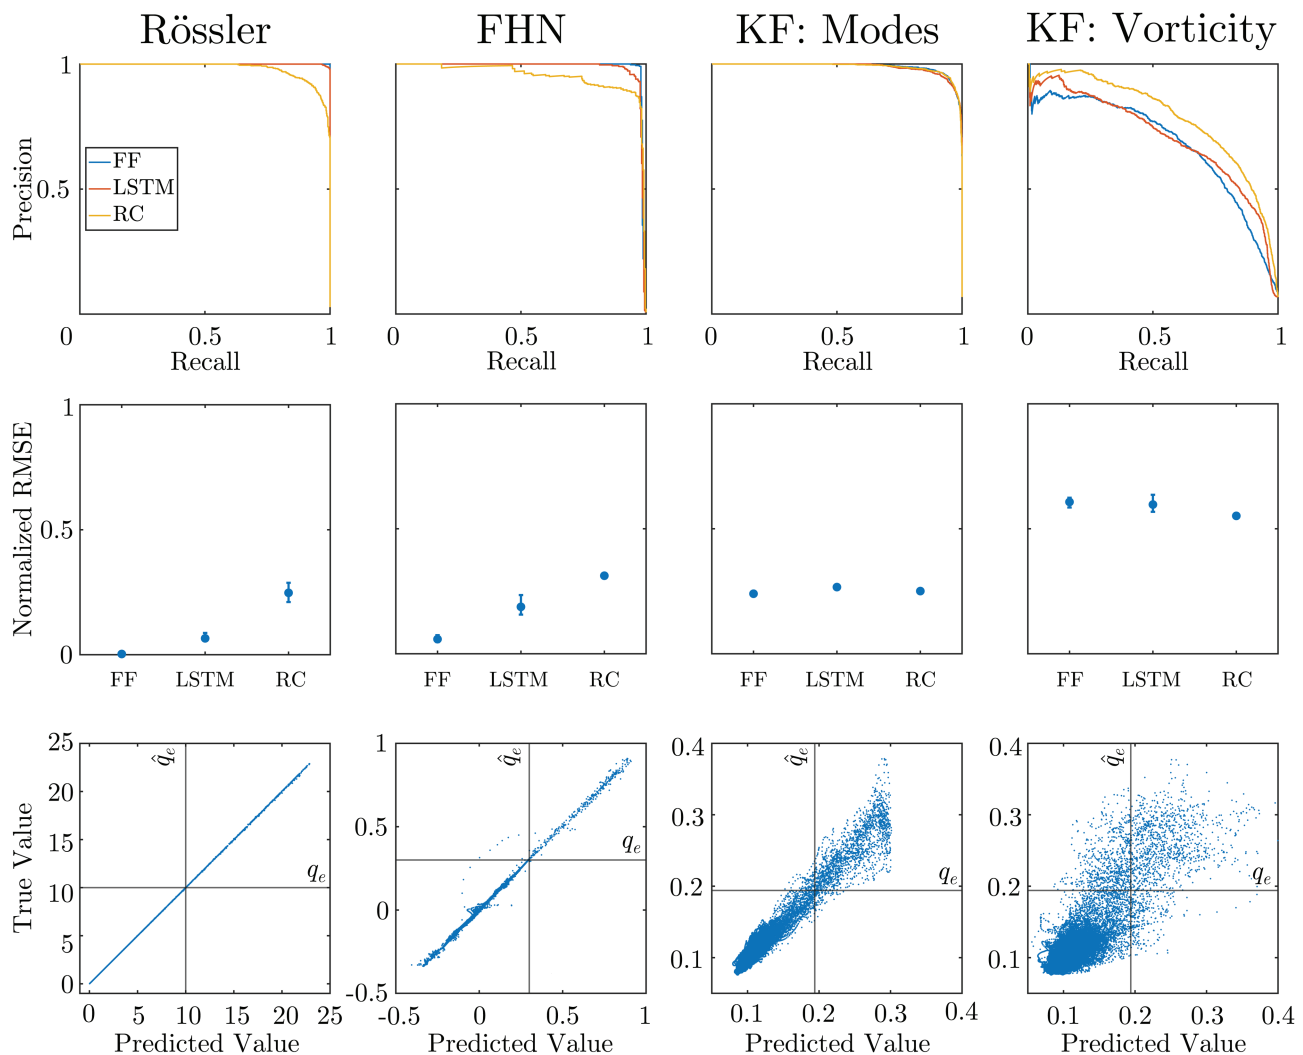


FIG. 7. Summary of neural network performance for each system with noiseless data. The precision–recall curves (first row) quantify the prediction skill of each network. The NRMSE plots (second row) quantify the fit of the predicted QoI, with error bars describing the maximum and minimum NRMSE for repeated training of the networks. The true vs predicted plots (third row) correspond to the best network architecture for each system.

unit square ranging between 0 and 1, with more skillful predictors having an AUC closer to 1. The first row of Fig. 7 provides examples of precision–recall curves.

B. Rössler system

The Rössler system is the least challenging system discussed here due to its lower dimensionality and relatively simpler dynamics. This system consists of three state variables $\mathbf{x} = (x_1, x_2, x_3)$, where the coordinate x_3 exhibits chaotic bursts (see Fig. 6). Therefore, we take $q = x_3$ as the quantity of interest with the extreme event threshold $q_e = 10$. The other two coordinates are used as observables, $p_1 = x_1$ and $p_2 = x_2$. We refer to

Subsection 1 of Appendix A for a detailed description of the Rössler system. In Subsection 1 of Appendix B, we discuss optimal hyperparameters used for each network when trained with Rössler data.

To generate training and testing data, we integrate the system for 500 time units with the results recorded every $\Delta t = 0.05$. Increasing the sampling time to $\Delta t = 0.5$ does not significantly alter the reported results, indicating robustness to moderate changes to the sampling rate. However, increasing the sampling rate further will eventually reduce the prediction skill of the networks across all three dynamical systems. In Rössler, for instance, increasing the sampling time beyond $\Delta t = 2$ led to a significant drop in the corresponding AUC.




FIG. 8. AUC for the Rössler system as the number of extreme events in the training data increases.

The amount of required training data is not *a priori* known. However, Fig. 8 gives an indication as to when enough training data from simulations are obtained. As the number of extreme events in the training data increases, the AUC initially increases rapidly. After the training data are long enough to contain about 100 extreme events, the AUC plateaus near 1, indicating that enough training data have been gathered, and therefore, we stop the numerical simulations. We also observed a similar plateauing behavior for the FitzHugh–Nagumo system and the Komogorov flow, although the required number of extreme events before reaching the plateau depends on the system.

We first discuss the neural network predictions in the absence of observational noise. As shown in the first column of Fig. 7, all three networks generate nearly exact predictions on testing data in the absence of noise, each with an AUC above 0.98. The RC

network performs slightly worse at an AUC of 0.98 whereas the AUC for Feedforward and LSTM networks is nearly equal to 1.00.

Training the network weights is not generally a convex optimization problem. Therefore, retraining the network may lead to different results. In order to quantify the reproducibility of the results, we retrain each network 10 times with different initial network weights. The error bars in the NRMSE plot within Fig. 7 correspond to the minimum and maximum value over 10 repeated training of each network. The error bars are relatively small across all systems and all networks, demonstrating robustness under retraining. For Rössler, the feedforward network has the best average NRMSE at 0.003 ± 0.0002 . LSTM networks perform slightly worse with an NRMSE within 0.098 ± 0.0109 . Finally, the RC network has the largest NRMSE within the range 0.248 ± 0.040 .

We also quantify the prediction skill as a function of the prediction time τ . Because of sensitivity to initial conditions, the predictability horizon of a chaotic system is always limited and often scales inversely with the leading Lyapunov exponent of the system. This limitation also applies to extreme event predictions.²⁹ However, as we explained in Sec. IV A, prediction of extreme events is a more forgiving task as compared to time series prediction. For the Rössler system, increasing the prediction time τ shows an expected decrease in accuracy for all networks, as shown in Fig. 9. Feedforward networks maintain the AUC at nearly 1 until $\tau = 20$, while the LSTM network maintains the AUC at nearly 1 until $\tau = 12$. The RC networks perform worse than FF and LSTM networks as τ increases.

Next, we examine the effect of observational noise on the predictions. Note that, although observations are always polluted with noise, the training data in our framework are gathered off-line using numerical simulations with no significant noise (see Fig. 1). Nonetheless, noise can be artificially added to the simulation data to mimic the observational noise. One may argue that the networks should be trained on clean simulation data (with no artificial noise) and hope that the network will perform well under small to moderate amounts of observational noise. In the following, we show that this approach is ill-advised since adding some artificial noise in the training phase often improves the network performance.

To this end, we first consider the case where no noise is added to the training data and only the testing data are polluted with




FIG. 9. AUC as a function of the prediction time τ .

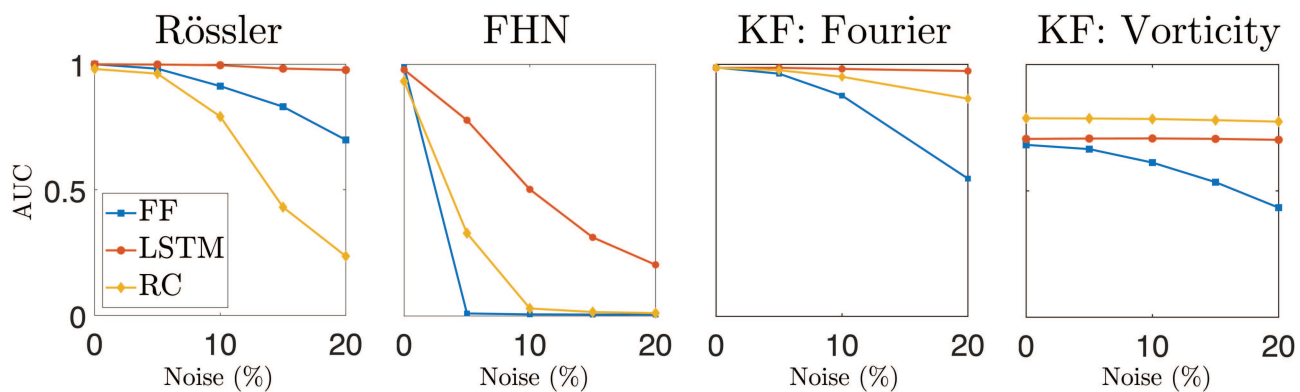


FIG. 10. Effect of noise on predictions. Each panel shows how the area under the precision–recall curve (AUC) changes as the noise intensity increases. Noise is added only to the testing data.

noise as described in Eq. (6). As shown in Fig. 10, we see notable differences in accuracy among the network architectures. LSTM networks consistently maintain their predictive power with the AUC decreasing from nearly 1.00 to 0.98 as the noise intensity increases from 0% to 20%. Feedforward networks perform similarly for 0% and 5% noise; but as the noise intensity increases to 20%, the AUC decreases to 0.7. Noise affects the RC network the most, with the AUC dropping to 0.79 for 10% noise. As the noise intensity reaches 20%, the AUC of RC decreases significantly to 0.24.

The detrimental effect of noise is more pronounced for the FitzHugh–Nagumo system. We show in Sec. IV C that adding artificial noise to the training data, as well as the testing data, significantly improves the prediction results (see Fig. 11).

C. FitzHugh–Nagumo system

The FitzHugh–Nagumo system is a network of chaotic oscillators composed of coupled units. The coordinates of the i th unit are denoted with (v_i, w_i) . Here, we consider a model with 101 completely coupled units resulting in a 202 dimensional system of ODEs. We assume that only the coordinates of the first unit are observable so that $p_1 = v_1$ and $p_2 = w_1$. The quantity of interest is the average of the v_i coordinates, $q = (\sum_{i=1}^n v_i) / n$. The FHN oscillators occasionally synchronize so that the v_i coordinates align, leading to intermittent bursts in the average q (see Fig. 6). The extreme event threshold is $q_e = 0.3$.

We integrate the FHN system for 2×10^5 time units with the results saved every $\Delta t = 1$ time unit. The FHN model is described in detail in Subsection 2 of Appendix A. The optimal hyperparameters for training the neural networks are detailed in Subsection 2 of Appendix B.

Examining the second column of Fig. 7, the precision–recall curves lie in the top right yielding large AUC values. The FF network performs best with an AUC of 0.99 while the LSTM network had a comparable AUC of 0.98. The reservoir network has the smallest AUC of 0.93. The scatterplot lies close to the diagonal line indicating the predicted quantity of interest closely matches its true value, although the predictions are not as accurate as the Rössler system.

This is to be expected since FHN is higher dimensional and only two out of its 202 coordinates are observed.

With multiple network trainings, the FF network performs the best with NRSME at 0.059 ± 0.014 . The LSTM network has the largest variability in NRMSE which lies in 0.188 ± 0.039 . Although the reservoir network has the largest NRMSE, the results have the most consistency, with NRMSE within 0.312 ± 0.09 . Recall that, for the Rössler system, the FF and LSTM networks were the most robust under retraining. In contrast, for the FHN system, the FF and RC networks are the most robust.

As shown in Fig. 9, the prediction accuracy decreases as prediction time increases. However, there is a significant loss in predictive power for the LSTM and RC networks, with the AUC declining to near 0 for the LSTM and RC networks at $\tau = 25$. In contrast, the FF network maintains accurate predictions, with an AUC of 0.97 for the prediction time $\tau = 30$.

For the FHN system, an interesting phenomenon takes place when adding noise to the testing data. As shown in Fig. 10, the results are extremely sensitive to noise. For instance, for the FF network, the AUC is 0.99 when using noise-free testing data; but even adding 5% noise to the testing data, the AUC drops to 0.01. While the LSTM network appears slightly more robust to noise, its predictions on data polluted with 20% noise are still inaccurate, with the AUC dropping to around 0.25. Contrast this with the LSTM results on the Rössler system where the AUC remains close to 1 even when 20% noise is added to the testing data.

Recall that the results reported in Fig. 10 correspond to a network trained on noise-free data and then tested on noisy observational inputs. Interestingly, if we add artificial noise to the training data, the prediction accuracy of all networks increases significantly. Figure 11 demonstrates the effect of training noise on predicting the quantity of interest q . With 0% training noise, the predicted quantity of interest $\hat{q}(t)$ deviates drastically from the true value $q(t)$ (the left panel of Fig. 11). In contrast, after adding 20% noise to the training data, the predicted quantity of interest closely resembles its true time series (see the middle panel of Fig. 11). This is also reflected in the AUC as shown in the right panel of Fig. 11 where AUC stays close to 1 as the noise intensity in the testing data increases. Contrast this

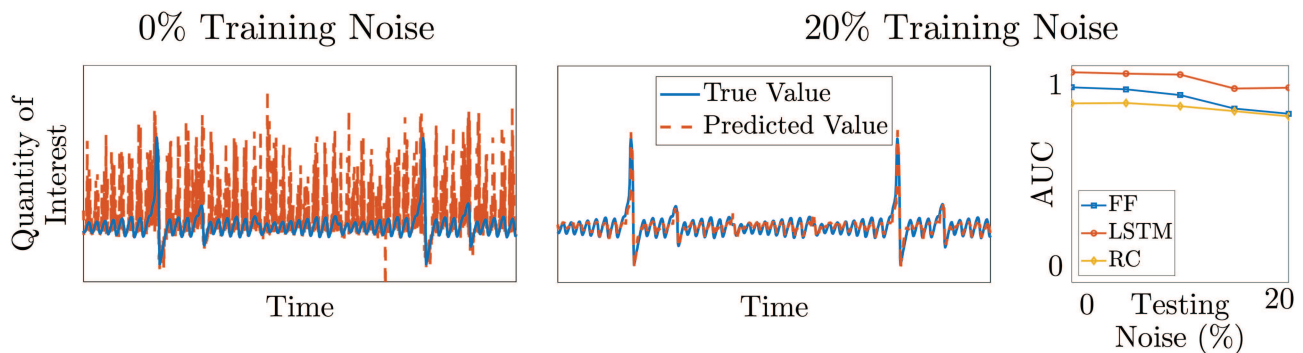


FIG. 11. Effect of training noise on predictions for the FitzHugh–Nagumo system. The feedforward network is trained with 0% training noise (the left panel) and 20% training noise (middle and right panels). The figures show the test data containing 5% noise. The networks trained with noisy data are more robust to noise than networks trained with noise-less data; contrast the AUC to the corresponding panel of Fig. 10.

with the corresponding Fig. 10 where no artificial noise was added to the training data.

In summary, adding artificial noise to the training data leads to more accurate extreme event predictions if the real-time observational measurements (i.e., testing data) are noisy. We have made a similar observation for the Rössler system and the Kolmogorov flow (not shown here for brevity). This is in line with previous studies which find that adding noise to the data prevents overfitting.^{44,45}

D. Kolmogorov flow

In this section, we finally consider the Kolmogorov flow, a particular type of turbulent fluid flow with periodic boundary conditions and sinusoidal forcing. Solving the corresponding Navier–Stokes equations, we generate the full simulation data that consists of the velocity field $\mathbf{u}(\mathbf{x}, t)$. The equations are numerically integrated for 10^5 time units and the velocity field is saved on a uniform 128×128 spatial grid at equispaced time intervals $\Delta t = 0.2$. A detailed description of the Kolmogorov flow is provided in Subsection 3 of Appendix A.

It is well-known that the energy dissipation rate $D(t)$, i.e., the spatial average of the vorticity field, exhibits extreme events in the form of intermittent bursts³² (see Fig. 6). Therefore, we take the quantity of interest to be the energy dissipation rate, $q = D$.

We consider two sets of neural network inputs for the Kolmogorov flow. First, we predict extreme events using the Fourier mode $a(1, 0) \in \mathbb{C}$, where $a(\mathbf{k})$ denotes the Fourier mode of the flow corresponding to the wave number \mathbf{k} . Using a variational method, Farazmand and Sapsis³³ discovered that the Fourier mode $a(1, 0)$ constitutes a precursor to extreme events in the Kolmogorov flow. More precisely, the magnitude $|a(1, 0)|$ of this mode decreases shortly before a burst in the energy dissipation D is observed. Therefore, we expect that using this mode as the observable will enable accurate prediction of extreme events in this flow. Consequently, we take the first set of observable to be the real and imaginary parts of this Fourier mode, i.e., $p_1 = \text{Re}[a(1, 0)]$ and $p_2 = \text{Im}[a(1, 0)]$. The results with Fourier observables are discussed in Sec. IV D 1.

As a second set of observables, we consider discrete measurements of the vorticity field $\omega(\mathbf{x}, t) = \nabla \times \mathbf{u}(\mathbf{x}, t)$. More precisely,

we consider the vorticity field evaluated on a uniform 3×3 grid at points $\mathbf{x}_i \in \{\frac{\pi}{3}, \pi, \frac{5\pi}{3}\} \times \{\frac{\pi}{3}, \pi, \frac{5\pi}{3}\}$. Then, the observables are given by $p_i(t) = \omega(\mathbf{x}_i, t)$, $i = 1, 2, \dots, 9$. The results with vorticity observables are discussed in Sec. IV D 2.

As we show below, the neural networks trained with Fourier observables vastly outperform those trained with vorticity measurements. This is in spite of the fact that Fourier mode data are of smaller dimension (two time series) than the vorticity measurements (nine time series), highlighting the importance of the input data for deep learning of extreme events. The optimal hyperparameters for training the neural networks are detailed in Subsection 3 of Appendix B.

1. Fourier mode

As shown in Fig. 7, all network structures perform well with the real and imaginary parts of the Fourier mode $a(1, 0)$ as input. The performance of the networks is comparable with their AUCs between 0.986 and 0.988. Nonetheless, the FF neural network slightly outperforms the LSTM and RC networks.

The FF network also marginally outperforms the other two networks when performance is measured in terms of NRMSE. The NRMSE for the FF network is 0.240, while the RC network had an error of 0.251 and the LSTM network had an error of 0.267. For all networks, the NRMSE varies insignificantly over repeated training, with variations of less than 0.01. This implies that the learned networks are robust with respect to retraining.

Keeping network structures fixed, we show the effect of increasing the prediction time τ in Fig. 9. All network architectures perform similarly as τ increases, retaining strong predictive power up to $\tau = 3$, where the AUC is approximately 0.9 for each network. As τ increases further above 5 time units, the predictive power drops off significantly, with the AUC reducing to around 0.5.

In Fig. 10, we analyze the effect of noisy input data on network performance. The LSTM network is remarkably robust to noise, returning AUC values for 20% noise which are comparable to those for noise-free data. The performance of RC network deteriorates slightly as the noise intensity increases. As in the case of Rössler

and FHN systems, the FF network performs significantly worse with noisy input data.

In summary, when supplied with the Fourier input data, all network structures produce accurate predictions for moderate prediction times ($\tau < 5$) and moderate noise intensity (less than 10%). This should not be surprising since the Fourier mode $a(1, 0)$ is known to act as a precursor to extreme events in the Kolmogorov flow.³³ In fact, Farazmand and Sapsis³³ were able to predict extreme events in this system by tracking the modulus $|a(1, 0)|$ without using a deep learning method.

However, measuring the Fourier mode $a(1, 0)$ is not straightforward in practice; to compute the Fourier transform, the flow must be measured on a dense spatial grid. A more realistic set of observables is velocity or vorticity measurements on a sparse spatial grid. As we see in Sec. IV D 2, the prediction accuracy deteriorates significantly if such realistic input data are used for training the neural networks.

2. Vorticity samples

Kolmogorov flow with vorticity sampling input data is the most difficult system-data combination we consider. As a result, our networks predict extreme events less accurately as compared to other systems.

As shown in Fig. 7, when supplied with the vorticity as the input data, the RC network yields the best predictions in terms of maximal AUC and minimal NRMSE. Reservoir computing gives an AUC of 0.787, whereas the LSTM network has an AUC of 0.706. The feedforward network performs the worst with an AUC of 0.682. Note that this is in contrast with previous systems where the FF network outperformed RC when tested on noiseless input data.

Similarly, conclusions are drawn when comparing performance based on NRMSE. RC performs best with a mean NRMSE of 0.552, relative to 0.597 for LSTM networks and 0.607 for FF networks. We see that LSTM network exhibits the most variation under multiple training, differing from the mean NRMSE by 0.040 between different trainings. NRMSE varies significantly less for FF and RC networks, on the order of 10^{-4} .

Figure 9 shows the prediction performance, measured by AUC, as the prediction time τ increases from $\tau = 0$ to $\tau = 10$. As expected, larger prediction times lead to lower accuracy. The prediction accuracies decline at a similar rate across networks.

As shown in Fig. 10, when noise is added to the test data, LSTM and RC networks perform significantly better than the FF network. As in the previous systems, the LSTM network is quite robust to noise, with the AUC barely dropping even when the noise intensity increases to 20%. In the Kolmogorov flow with the vorticity input, the RC network is also equally robust to noise. In contrast, the AUC drops sharply for the FF network as the noise intensity increases. We find that adding noise to the training data slightly improves the performance of the FF network on noisy test data (not shown here for brevity). Recall that a similar improvement was observed for the FHN system (see Fig. 11); however, the improvement for the Kolmogorov flow is not as drastic as the FHN system.

V. SUMMARY AND CONCLUSIONS

Using deep learning to predict extreme events remains challenging due to the small data problem. Even large amounts of observational data contain few extreme events to sufficiently train a neural network. Here, we investigated whether training on numerically generated data, with sufficient samples from extreme events, can overcome this obstacle. In addition, to mimic the practical situation where observable quantities are limited, we did not use the entire simulation data for training and instead used a small subset of the available time series data as network inputs.

To examine this framework, we conducted a thorough study on three dynamical systems (Rössler system, FitzHugh–Nagumo model, and Kolmogorov flow) and three standard machine learning architectures (Feedforward, long short-term memory, and reservoir computing networks), resulting in nine total combinations.

As summarized in Table II, there is no universal answer: no network architecture consistently outperforms others across different dynamical systems. Nonetheless, some broad conclusions can be drawn from our analysis. For instance, LSTM networks are most robust to observational noise, i.e., they largely maintain their accuracy of extreme event prediction when a moderate amount of noise is added to the observational data even if the training data (obtained from simulations) have no significant noise. In contrast, FF neural networks are most sensitive to observational noise, with their prediction accuracy deteriorating rapidly as the noise intensity increases. Note that although RC works best for the Kolmogorov flow with vorticity inputs, its performance is not significantly better than the corresponding LSTM network (see Fig. 10).

Furthermore, LSTM networks performed well with minimal fine-tuning of hyperparameters. For LSTM networks, we only had to fine-tune the number of layers and units per layer while for FF and RC networks several hyperparameters had to be fine-tuned to achieve comparable prediction accuracy.

We also find that adding artificial noise to the training data consistently improves the predictions. Recall that our training data are provided by numerical simulations, which contain no significant noise (other than round-off and numerical discretization errors). However, one can add some artificial noise to the simulation data to mimic the observational noise. Adding noise to the training data improved the performance of all networks across all systems. This improvement was specifically significant for the FF neural network trained on the FHN data. This observation is in line with previous studies, which find that adding noise to the data prevents overfitting.^{44,45}

Another important observation is the sensitivity of the results to the type of observations, i.e., network inputs. This point is clearly

TABLE II. The best performing machine learning architecture as measured by the area under the precision–recall curve (or AUC). FHN stands for the FitzHugh–Nagumo model and KF stands for the Kolmogorov flow.

	Rössler	FHN	KF: Fourier	KF: Vorticity
Noiseless data	FF	FF	FF, LSTM, RC	RC
Noisy data	LSTM	LSTM	LSTM	RC

demonstrated on the Kolmogorov flow where extreme events are predicted with high accuracy if the input data are obtained from the Fourier mode $a(1, 0)$. In contrast, a large number of false positive and false negative predictions are made if the input data consists of sparse vorticity measurements. Using a variational method, Farazmand and Sapsis³³ had previously found the Fourier mode $a(1, 0)$ to play an important role in the formation of extreme events in the Kolmogorov flow. Therefore, when possible, it is advised to discover precursors to extreme events and train the neural networks using these precursors. Discovering precursors to extreme events is itself challenging and remains the subject of ongoing research.²

On a related note, assume that a large number of observable time series of the system are available. To maintain a reasonable training cost, one needs to select a subset of the observables as input data. Currently, a concrete method to select the optimal subset of the observed time series, ensuring maximal extreme event prediction skill, is missing. More theoretical work in this direction is highly desirable.

The performance of a deep neural network is sensitive to its hyperparameters, such as the number of layers, nodes, and learning rate. Here, we determined the hyperparameters manually using extensive trial-and-error searches. More systematic methods, such as grid search or random search,⁴⁶ can be used to determine the optimal hyperparameters.

Finally, the length of training data was here determined in an *ad hoc* manner. Our future theoretical work will focus on deriving a lower bound on the number of required extreme events in the training data, which guarantees accurate prediction of upcoming extreme events in the testing data.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (NSF) under Grant No. DMS-2051010 and National Security Agency under Grant Nos. H98230-20-1-0259 and H98230-21-1-0014.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data and code that support the findings of this study are openly available in GitHub at <https://github.com/mfarazmand/DeepLearningExtremeEvents>, Ref. 39.

APPENDIX A: GOVERNING EQUATIONS

1. Rössler system

The Rössler system is defined by the ODEs,

$$\begin{aligned}\dot{x}_1 &= -x_2 - x_3, \\ \dot{x}_2 &= x_1 + ax_2, \\ \dot{x}_3 &= b + x_3(x_1 - c),\end{aligned}\quad (\text{A1})$$

with $a = 0.2$, $b = 0.2$, and $c = 5.7$. This system exhibits a chaotic attractor and intermittent extreme event *bursts* in the x_3 component (see Fig. 6). Therefore, we consider the quantity of interest,

$$q(\mathbf{x}(t)) = x_3(t) \quad (\text{A2})$$

and the partial observations $p_1(\mathbf{x}(t)) = x_1(t)$, $p_2(\mathbf{x}(t)) = x_2(t)$. The extreme event threshold is $q_e = 10$.

To generate trajectory data, we perform numerical integration using the Runge–Kutta method, as implemented in MATLAB's `ode45`. We set the initial condition as $(0, 1, 0.1)$ and integrate the system for 500 time units. The results are saved every $\Delta t = 0.05$ time units.

2. FitzHugh–Nagumo system

The FitzHugh–Nagumo system consists of n excitable units (v_i, w_i) , $i = 1, \dots, n$. These pairs can be interpreted as neurons, where v_i is the voltage and w_i is a recovery term. The dynamics are described by a system of $2n$ differential equations,

$$\begin{aligned}\dot{v}_i &= v_i(a_i - v_i)(v_i - 1) - w_i + k \sum_{j=1}^n A_{ij}(v_j - v_i), \\ \dot{w}_i &= b_i v_i - c_i w_i.\end{aligned}\quad (\text{A3})$$

The full state is described by $\mathbf{x}(t) = [v_1(t), w_1(t), \dots, v_n(t), w_n(t)]$. Parameters a_i , b_i , and c_i describe the dynamics of each unit internally, k is the coupling strength among units, and A_{ij} describes the coupling connections, where $A_{ij} = A_{ji} = 1$ indicates units i and j are coupled. We choose the parameters in case B of Ref. 31, which shows chaotic behavior with extreme events. This system has $n = 101$ completely coupled units, meaning $A_{ij} = 1, \forall i, j$. The parameters are $a_i = -0.02651$ and $c_i = 0.02$ for all units while $b_i = 0.006 + 0.008(i - 1)/(n - 1)$. These choices, especially the coupling strength $k = 0.00128$, are further discussed by Feudel and co-workers.⁴⁷

The quantity of interest is the average voltage of all the nodes,

$$q(\mathbf{x}(t)) = \frac{1}{n} \sum_{i=1}^n v_i(t). \quad (\text{A4})$$

The units of FHN oscillator occasionally synchronize, leading to intermittent bursts in the average q . We choose an extreme event threshold $q_e = 0.3$. We use $p_1(\mathbf{x}(t)) = v_1(t)$, $p_2(\mathbf{x}(t)) = w_1(t)$ as observables, but other units give comparable performance.

We use a Runge–Kutta method in MATLAB's `ode45` to numerically integrate the FHN system for 2×10^5 time units and save the results every $\Delta t = 1$ time units. We discard the first 100 time units as transient data from the initial condition $(v_i, w_i) = (0.1, 0.1)$, $i = 1, \dots, 101$.

3. Kolmogorov flow

We consider the Navier–Stokes equations for incompressible fluid flow,

$$\begin{aligned}\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \nu \Delta \mathbf{u} + \mathbf{F}, \\ \nabla \cdot \mathbf{u} &= 0,\end{aligned}\quad (\text{A5})$$

where $\mathbf{u}(\mathbf{x}, t)$ is the velocity field, $p(\mathbf{x}, t)$ is the pressure field, ν is the kinematic viscosity, and \mathbf{F} is the external forcing function. We specify the Kolmogorov flow as a specific two-dimensional realization of the Navier–Stokes equations. The domain is a doubly-periodic box of size $2\pi \times 2\pi$. A turbulent regime is chosen by using the Reynolds number $Re = \nu^{-1} = 40$. The external forcing function is the sinusoidal shearing function, $\mathbf{F}(x, y) = \sin(4y)\mathbf{e}_1$, where \mathbf{e}_1 is the standard basis vector in the x direction in \mathbb{R}^2 . We numerically integrate the Kolmogorov flow using a Fourier pseudospectral method in space with 2^7 modes in each dimension and a fourth-order Runge–Kutta scheme for time stepping. The flow is evolved for 10^5 time units and results are saved every $\Delta t = 0.2$ time units. We start the simulation from a random initial condition and discard the first 20 time units of data as transients.

Here, the quantity of interest is the energy dissipation rate,

$$q(\mathbf{u}) = \frac{\nu}{|\Omega|} \int_{\Omega} |\nabla \mathbf{u}|^2 d\mathbf{x}, \quad (\text{A6})$$

where $|\Omega| = (2\pi)^2$ is the area of the spatial domain $\Omega = [0, 2\pi] \times [0, 2\pi]$. It is known that the Kolmogorov flow exhibits intermittent bursts where the energy dissipation rate q increases to several standard deviations above its mean value³² (see Fig. 6). We use the extreme event threshold $q_e = 0.194$, which coincides with the mean plus twice the standard deviation of energy dissipation rate.

It has been recently discovered that these extreme events are instigated by the energy transfer between particular Fourier modes.³³ Consider the Fourier series expansion of the velocity field,

$$\mathbf{u}(\mathbf{x}, t) = \sum_{\mathbf{k} \in \mathbb{Z}^2} \frac{a(\mathbf{k}, t)}{|\mathbf{k}|^2} \begin{pmatrix} k_y \\ -k_x \end{pmatrix} e^{i\mathbf{k} \cdot \mathbf{x}}, \quad (\text{A7})$$

where $\mathbf{k} = (k_x, k_y)$ denotes the wave number and $a(\mathbf{k}, t) \in \mathbb{C}$ denotes the corresponding Fourier coefficient. Note that the incompressibility of the velocity field ($\nabla \cdot \mathbf{u} = 0$) is used to reduce the Fourier coefficients to scalars. Since the velocity is real-valued, we have $a(-\mathbf{k}, t) = -a(\mathbf{k}, t)^*$.

Using a variational method, Farazmand and Sapsis³³ showed that the extreme energy dissipation events are preceeded by an energy transfer from mode $a(1, 0, t)$ to the mode $a(0, 4, t)$. Motivated by this work, we use the real and imaginary parts of $a(1, 0, t)$ as the first set of observables for the Kolmogorov flow. More precisely, the Fourier observables are $p_1(t) = \text{Re}[a(1, 0, t)]$ and $p_2(t) = \text{Im}[a(1, 0, t)]$.

A second set of observables are obtained through the vorticity field $\omega = (\nabla \times \mathbf{u}) \cdot \mathbf{e}_3$, where \mathbf{e}_3 is the unit vector normal to the planar fluid domain. We evaluate the vorticity field at nine distinct spatial locations $\mathbf{x}_i \in \{\frac{\pi}{3}, \pi, \frac{5\pi}{3}\} \times \{\frac{\pi}{3}, \pi, \frac{5\pi}{3}\}$, leading to nine observables $p_i(t) = \omega(\mathbf{x}_i, t)$.

APPENDIX B: NETWORK ARCHITECTURES

In this section, we discuss the optimal parameters found for each of our example systems and the neural networks. The hyperparameters are found manually by trial-and-error. Therefore, the word “optimal” is used loosely; it only refers to the optimal combination of hyperparameters we investigated, not optimal among all possible combinations. All networks were optimized on noise-free data

before exploring network resilience to noise in training and testing data. All hyperparameters not mentioned here were set to the default values (Matlab default values for FF and LSTM networks, and `easyn` library for RC).

1. Rössler system

The best performing FF network architecture for the Rössler system consists of three layers with six nodes each with a tanh activation function. The architecture used $m = 3$ time delays, $s = 1$ time unit apart. These time delays improve the AUC and NRMSE of the network.

The optimal LSTM architecture consists of 2 layers with 55 nodes per layer. The optimal RC network contains 850 nodes in the reservoir. However, with 5% testing noise, the network AUC quickly decreased to 0.0359, suggesting that the network is overfitted. Re-optimizing the RC network with noisy data yields a drastically different optimal network: 50 nodes in the reservoir. Retraining the network on noiseless data with 50 nodes shows that at 0% testing noise the AUC is 0.9811, but the predictive power remains reasonable for higher levels of noise. At 5%, the AUC decreases to 0.9623. The optimal spectral radius is $\rho = 0.3$, and the leaking rate is $\ell = 1$. The optimal input density is 1, the reservoir density is 0.2, and the regularization weight is 10^{-4} .

2. FitzHugh–Nagumo system

The optimal FF network architecture is three layers with eight nodes, each with the tanh activation function. The architecture used $m = 2$ time delays with $s = 2$ time unit. However, this architecture is susceptible to noise in the testing data. We address the overfitting by adding noise to the training data. When varying the magnitude of noise in the training data, we find that increasing noise in the training data notably increases performance with high noise but slightly decreases performance with low noise. We choose to favor robustness and add relatively high 20% noise to the training data. Other alterations to neural network training such as reducing the number of epochs trained and adding a regularization term that penalized large weights did not produce significant improvements.

The LSTM network architecture of 2 layers with 64 hidden units each was found to be optimal. The time it takes to train the LSTM networks is much longer than FF and RC networks, as expected.

We optimized a wider range of parameters for the RC networks due to poor performance with default parameters. The optimal architecture used 500 nodes with spectral radius $\rho = 0.9$ and leaking rate $\ell = 0.3$. The input density is 1 with reservoir density 0.9 and regularization weight 10^{-4} . The spectral radius had the most significant effect on performance in the parameter ranges we examined.

3. Kolmogorov flow

a. Fourier mode

The best performing FF network architecture uses three layers, each with four nodes and the tanh activation function. We augment input data with $m = 8$ time delays, taken at time steps of $s = 0.2$

time units. The best performing LSTM network has 3 layers, each with 32 nodes.

The best performing RC network has 600 nodes and a spectral radius of $\rho = 0.9$ with leaking rate $\ell = 1$. The network also uses an input density of 1, a reservoir density of $\rho = 0.2$, and a regularization weight of 0.1.

b. Vorticity samples

The best performing FF neural network has four layers with eight nodes per layer, the tanh activation function, and $m = 12$ time delays each $s = 0.2$ time units apart. The optimal FF network taking in vorticity samples requires more layers and nodes because the vorticity data are higher dimensional than the Fourier modes.

The LSTM network with 2 layers and 16 nodes per layer performed best. Using more than 16 nodes per layer increases computational cost without improving the predictions. In fact, networks with 2 LSTM layers of 64 or more nodes per layer exhibit signs of overfitting. We added dropout layers to try to fix this issue, but prediction accuracy, as measured by NRMSE and AUC, did not improve when measured on testing data.

The best RC network has 1000 nodes, spectral radius $\rho = 0.2$, and leaking rate $\ell = 0.3$. The optimal input density is 0.3 with a reservoir density of 0.2 and a regularization weight of 10^{-4} .

REFERENCES

- ¹L. K. Comfort, A. Boin, and C. C. Demchak, *Designing Resilience: Preparing for Extreme Events* (University of Pittsburgh, 2010).
- ²M. Farazmand and T. P. Sapsis, "Extreme events: Mechanisms and prediction," *Appl. Mech. Rev.* **71**, 050801 (2019).
- ³L. E. McPhillips, H. Chang, M. V. Chester, Y. Depietri, E. Friedman, N. B. Grimm, J. S. Kominoski, T. McPhearson, P. Méndez-Lázaro, E. J. Rosi, and J. Shafiei Shiva, "Defining extreme events: A cross-disciplinary review," *Earth's Future* **6**, 441–455 (2018).
- ⁴J. L. Callahan, K. Maeda, and S. L. Brunton, "Robust flow reconstruction from limited measurements via sparse representation," *Phys. Rev. Fluids* **4**, 103907 (2019).
- ⁵B. Chu and M. Farazmand, "Data-driven prediction of multistable systems from sparse measurements," *Chaos* **31**, 063118 (2021).
- ⁶B. Kramer, P. Grover, P. Boufounos, S. Nabi, and M. Benosman, "Sparse sensing and DMD-based identification of flow regimes and bifurcations in complex flows," *SIAM J. Appl. Dyn. Syst.* **16**, 1164–1196 (2017).
- ⁷M. Farazmand and T. P. Sapsis, "Reduced-order prediction of rogue waves in two-dimensional deep-water waves," *J. Comput. Phys.* **340**, 418–434 (2017).
- ⁸A. Mendez and M. Farazmand, "Investigating climate tipping points under various emission reduction and carbon capture scenarios with a stochastic climate model," *Proc. R. Soc. London A*, **477**, 20210697 (2021).
- ⁹J. Bucklew, *Introduction to Rare Event Simulation* (Springer Science & Business Media, 2004).
- ¹⁰G. Dematteis, T. Grafke, and E. Vanden-Eijnden, "Rogue waves and large deviations in deep sea," *Proc. Natl. Acad. Sci. U.S.A.* **115**, 855–860 (2018).
- ¹¹M. A. Mohamad and T. P. Sapsis, "Sequential sampling strategy for extreme event statistics in nonlinear dynamical systems," *Proc. Natl. Acad. Sci. U.S.A.* **115**, 11138–11143 (2018).
- ¹²S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Machine learning for fluid mechanics," *Annu. Rev. Fluid Mech.* **52**, 477–508 (2020).
- ¹³H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P. Muller, "Deep learning for time series classification: A review," *Data Min. Knowl. Discov.* **33**, 917–963 (2019).
- ¹⁴G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nat. Rev. Phys.* **3**, 422–440 (2021).
- ¹⁵D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer, "Machine learning-accelerated computational fluid dynamics," *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2101784118 (2021).
- ¹⁶J. Pathak, Z. Lu, B. R. Hunt, M. Girvan, and E. Ott, "Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data," *Chaos* **27**, 121102 (2017).
- ¹⁷D. Ding, M. Zhang, X. Pan, M. Yang, and X. He, "Modeling extreme events in time series prediction," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19* (Association for Computing Machinery, New York, NY, 2019), pp. 1114–1122.
- ¹⁸D. Qi and A. J. Majda, "Using machine learning to predict extreme events in complex systems," *Proc. Natl. Acad. Sci. U.S.A.* **117**, 52–59 (2020).
- ¹⁹S. H. Rudy and T. P. Sapsis, "Output-weighted and relative entropy loss functions for deep learning precursors of extreme events," *arXiv:abs/2112.00825* (2021).
- ²⁰J. Meiyazhagan, S. Sudharsan, and M. Senthilvelan, "Model-free prediction of emergence of extreme events in a parametrically driven nonlinear dynamical system by deep learning," *Eur. Phys. J. B* **94**, 1–13 (2021).
- ²¹M. Lellep, J. Prexl, M. Linkmann, and B. Eckhardt, "Using machine learning to predict extreme events in the Hénon map," *Chaos* **30**, 013113 (2020).
- ²²M. Närhi, L. Salmela, J. Toivonen, C. Billet, J. M. Dudley, and G. Genty, "Machine learning analysis of extreme events in optical fibre modulation instability," *Nat. Commun.* **9**, 1–11 (2018).
- ²³P. K. Yeditha, V. Kasi, M. Rathinasamy, and A. Agarwal, "Forecasting of extreme flood events using different satellite precipitation products and wavelet-based machine learning methods," *Chaos* **30**, 063115 (2020).
- ²⁴Z. Y. Wan, P. Vlachas, P. Koumoutsakos, and T. Sapsis, "Data-assisted reduced-order modeling of extreme events in complex dynamical systems," *PLoS One* **13**, e0197704 (2018).
- ²⁵A. Chattopadhyay, E. Nabizadeh, and P. Hassanzadeh, "Analog forecasting of extreme-weather patterns using deep learning," *J. Adv. Model. Earth Syst.* **12**, e2019MS001958 (2020).
- ²⁶S. H. Rudy and T. P. Sapsis, "Prediction of intermittent fluctuations from surface pressure measurements on a turbulent airfoil," *AIAA* (online 2022).
- ²⁷A. Chattopadhyay, P. Hassanzadeh, and D. Subramanian, "Data-driven predictions of a multiscale Lorenz 96 chaotic system using machine-learning methods: Reservoir computing, artificial neural network, and long short-term memory network," *Nonlinear Process. Geophys.* **27**, 373–389 (2020).
- ²⁸V. Pyragas and K. Pyragas, "Using reservoir computer to predict and prevent extreme events," *Phys. Lett. A* **384**, 126591 (2020).
- ²⁹M. Farazmand and T. P. Sapsis, "Dynamical indicators for the prediction of bursting phenomena in high-dimensional systems," *Phys. Rev. E* **94**, 032212 (2016).
- ³⁰P. J. Blonigan, M. Farazmand, and T. P. Sapsis, "Are extreme dissipation events predictable in turbulent fluid flows?," *Phys. Rev. Fluids* **4**, 044606 (2019).
- ³¹G. Ansmann, R. Karnatak, K. Lehnertz, and U. Feudel, "Extreme events in excitable systems and mechanisms of their generation," *Phys. Rev. E* **88**, 052911 (2013).
- ³²M. Farazmand, "An adjoint-based approach for finding invariant solutions of Navier-Stokes equations," *J. Fluid Mech.* **795**, 278–312 (2016).
- ³³M. Farazmand and T. P. Sapsis, "A variational approach to probing extreme events in turbulent dynamical systems," *Sci. Adv.* **3**, e1701533 (2017).
- ³⁴J. J. Moré, "The Levenberg-Marquardt algorithm: Implementation and theory," in *Numerical Analysis*, edited by G. A. Watson (Springer, Berlin, 1978), pp. 105–116.
- ³⁵F. Takens, "Detecting strange attractors in turbulence," in *Dynamical Systems and Turbulence, Warwick 1980*, edited by D. Rand and L. Young (Springer, Berlin, 1981), pp. 366–381.
- ³⁶S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.* **9**, 1735–1780 (1997).
- ³⁷A. Sherstinsky, "Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network," *Physica D* **404**, 132306 (2020).
- ³⁸B. Schrauwen, D. Verstraeten, and J. Van Campenhout, "An overview of reservoir computing: Theory, applications and implementations," in *Proceedings of*

the 15th European Symposium on Artificial Neural Networks (i6doc.com, 2007), pp. 471–482.

³⁹A. Asch, E. J. Brady, H. Gallardo, J. Hood, B. Chu, and M. Farazmand, “GitHub repository: Deep learning for extreme events,” (2021); available at <https://github.com/mfarazmand/DeepLearningExtremeEvents>.

⁴⁰R. Zimmermann, see <https://github.com/kalekiu/easyesn/> for “Easyesn: A Library for Recurrent Neural Networks Using Echo State Networks,” version 0.1.6.1 (2021).

⁴¹S. Guth and T. P. Sapsis, “Machine learning predictors of extreme events occurring in complex dynamical systems,” *Entropy* **21**, 925 (2019).

⁴²J. Davis and M. Goadrich, “The relationship between precision-recall and ROC curves,” in *Proceedings of the 23rd International Conference on Machine Learning, ICML '06* (Association for Computing Machinery, New York, NY, 2006), pp. 233–240.

⁴³T. Saito and M. Rehmsmeier, “The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets,” *PLoS One* **10**, e0118432 (2015).

⁴⁴G. An, “The effects of adding noise during backpropagation training on a generalization performance,” *Neural Comput.* **8**, 643–674 (1996).

⁴⁵R. M. Zur, Y. Jiang, L. L. Pesce, and K. Drukker, “Noise injection for training artificial neural networks: A comparison with weight decay and early stopping,” *Med. Phys.* **36**, 4810–4818 (2009).

⁴⁶S. R. Young, D. C. Rose, T. P. Karnowski, S.-H. Lim, and R. M. Patton, “Optimizing deep learning hyper-parameters through an evolutionary algorithm,” in *Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments, MLHPC '15* (Association for Computing Machinery, 2015).

⁴⁷R. Karnatak, G. Ansmann, U. Feudel, and K. Lehnertz, “Route to extreme events in excitable systems,” *Phys. Rev. E* **90**, 022917 (2014).