

Data-driven prediction and control of extreme events in a chaotic flow

Alberto Racca 

*Department of Engineering, University of Cambridge, Trumpington Street,
CB2 1PZ Cambridge, United Kingdom*

Luca Magri *

*Aeronautics Department, Imperial College London, London SW7 2AZ, United Kingdom;
Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom;
The Alan Turing Institute, London NW1 2DB, United Kingdom;
and Institute of Advanced Study, TU Munich, 85748 Munich Garching, Germany (visiting)*



(Received 21 April 2022; accepted 13 September 2022; published 31 October 2022)

An extreme event is a sudden and violent change in the state of a nonlinear system. In fluid dynamics, extreme events can have adverse effects on the system's optimal design and operability, which calls for accurate methods for their prediction and control. In this paper, we propose a data-driven methodology for the prediction and control of extreme events in a chaotic shear flow. The approach is based on echo state networks, which are a type of reservoir computing that learn temporal correlations within a time-dependent data set. The objective is fivefold. First, we exploit *ad hoc* metrics from binary classification to analyze (1) how many of the extreme events predicted by the network actually occur in the test set (precision) and (2) how many extreme events are missed by the network (recall). We apply a principled strategy for optimal hyperparameter selection, which is key to the networks' performance. Second, we focus on the time-accurate prediction of extreme events. We show that echo state networks are able to predict extreme events well beyond the predictability time, i.e., up to more than five Lyapunov times. Third, we focus on the long-term prediction of extreme events from a statistical point of view. By training the networks with data sets that contain nonconverged statistics, we show that the networks are able to learn and extrapolate the flow's long-term statistics. In other words, the networks are able to extrapolate in time from relatively short time series. Fourth, we design a simple and effective control strategy to prevent extreme events from occurring. The control strategy decreases the occurrence of extreme events up to one order of magnitude with respect to the uncontrolled system. Finally, we analyze the robustness of the results for a range of Reynolds numbers. We show that the networks perform well across a wide range of regimes. This work opens up new possibilities for the data-driven prediction and control of extreme events in chaotic systems.

DOI: [10.1103/PhysRevFluids.7.104402](https://doi.org/10.1103/PhysRevFluids.7.104402)

I. INTRODUCTION

Extreme events are sudden and violent changes of an observable. These events arise in a variety of fluid dynamic systems in the form of oceanic rogue waves [1,2], dissipation events in turbulence [3,4], thermoacoustic bursts in aeronautical engines [5], and atmospheric blocking events in weather forecasting [6], to name a few. Because of their large amplitudes, extreme events

*lm547@cam.ac.uk

can become problematic in many engineering applications because they may lead to component failure [7]. Because of the potentially negative consequences of extreme events, their prediction and suppression is an active area of research [8–11].

There are two ways to study extreme events: (1) computing their statistics and (2) time accurately predicting the incoming events. With a statistical approach, the goal is to predict the probability of the occurrence of an extreme event. To do so, methods such as extreme value theory [12], large deviation theory [13], and tailored Monte Carlo algorithms [14] are used to compute the heavy tail of the probability density functions that describe the events. With a time-accurate approach, the goal is to identify precursors, which are quantities that indicate at the current time the occurrence of an extreme event in the near future [1,15]. Although the time-accurate prediction is mathematically possible in deterministic systems, it is hindered by the “butterfly effect”: in chaotic dynamics the system is predictable only up to a time scale, known as the predictability time [16]. One simple scale for the predictability time is the inverse of the dominant Lyapunov exponent, also known as the Lyapunov time, which is typically a small characteristic scale of the flow [17].

In turbulence and chaotic flows, extreme events have been characterized with precursors [10]. The authors of [18] developed an optimal time-dependent reduced-order model, which consists of a basis that at any point of the trajectory spans the most unstable directions characterized by the largest finite-time growth of perturbations. The eigenvalues of the linear operator describing the evolution of perturbations in the reduced-order model performed well in predicting dissipation events in forced periodic turbulence [19]. However, computing the elements of the basis may be computationally prohibitive for systems with a large number of degrees of freedom [20]. The authors of [1] developed a variational method to find the optimal condition along the attractor that maximizes the instantaneous growth of the observable exhibiting extreme events, which was applied to predict rogue waves and forced periodic turbulence. These methods are *model-based*, because they rely on the perfect knowledge of the governing equations.

A different approach to the prediction of extreme events is *data-driven*. Given some time series (data), we wish to predict the future evolution of the system, usually without any knowledge of the governing equations. In the realm of data-driven techniques for predicting time series, Recurrent Neural Networks (RNNs) [21], and their variants, are the state-of-the-art machine learning architectures. They are designed to have memory of the past history of the system through an internal state, which enables the network to learn temporal correlations within the data. Because of their design, RNNs typically outperform other common machine learning architectures, such as feed-forward neural networks, in time series prediction [22]. In fluid dynamics, RNNs have been successfully applied for different goals such as closure for reduced-order modeling [23], prediction of flows past bluff bodies [24], replication of chaotic flow statistics [25,26] and nonlinear control for gliding [27], to name a few.

Among recurrent neural networks, Echo State Networks (ESNs) [28,29], which are a class of reservoir computing, are a versatile architecture for the prediction of chaotic dynamics. They have been proven to be universal approximators [30], which can infer the invariant measure of dynamical systems [31] under nonstringent assumptions of ergodicity [32]. ESNs perform as well as other architectures such as Long Short-Term Memory (LSTM) networks [33], while requiring less computational resources for training [22,34]. Although training echo state networks is a straightforward task as compared to other networks, such as LSTMs, their performance is highly sensitive to the selection of hyperparameters [35,36]. To improve the performance of the networks, prior knowledge of the governing equations has been embedded in the hybrid [37] and physics-informed [38] architectures. In fluid dynamics, echo state networks have been employed to (1) learn [31] and (2) optimize ergodic averages in thermoacoustic oscillations [39], and (3) time accurately [40] and (4) statistically predict extreme events in chaotic flows [40,41]. In particular, Doan *et al.* [40] studied the effect of including prior knowledge of the governing equations in a chaotic shear flow model to improve the prediction of extreme events when only a small training set is available. Although leveraging physics knowledge improves performance, information about the governing equations is not always available, and the networks remain sensitive to the selection of hyperparameters [36].

Once an upcoming extreme event is accurately predicted, the next objective is to suppress the event, i.e., to prevent the event from occurring. This is usually achieved via a control strategy that is based on the prediction of the event. The authors of [15] used information from an empirical precursor to suppress extreme events in a specifically designed six-dimensional system of electronic circuits, whereas the authors of [42] apply ESNs to suppress local extreme events in a system of coupled oscillators and those of [43] computed the most unstable initial conditions by solving an optimization problem to successfully suppress extreme events through closed-loop control in two-dimensional turbulence.

The objective of the paper is fivefold. First, we define a set of metrics to evaluate the performance of the networks for the time-accurate and statistical prediction of extreme events. The importance of a reliable hyperparameter selection is emphasized. Second, we investigate the time-accurate prediction of extreme events using echo state networks for times larger than the Lyapunov time. Third, we analyze the networks' capability to extrapolate statistical knowledge of the system with respect to the available data used for training. Fourth, we exploit the time-accurate prediction of extreme events to actively suppress the events, with a simple and effective control strategy. Fifth, we investigate the robustness of the networks in a wide range of Reynolds numbers.

The paper is organized as follows. Section II discusses the qualitative model of chaotic shear flow and the generation of the data set. Section III describes the echo state network architecture. Section IV defines the metrics and analyzes the time-accurate prediction of extreme events. Section V analyzes the statistical prediction of extreme events with respect to the available data sets. Section VI proposes an open-loop control strategy based on the time-accurate prediction of extreme events. Finally, Sec. VII summarizes the results of this study and discusses current and future work.

II. A QUALITATIVE MODEL OF CHAOTIC SHEAR FLOW

We consider a qualitative model of a shear flow between infinite plates subjected to a sinusoidal body forcing [44,45]. For brevity, this model is referred to as ‘‘Moehlis, Faisst, and Eckhardt (MFE)’’ from the author names of [44]. Because the MFE captures qualitatively nonlinear phenomena, such as relaminarization and turbulent bursts, it has been employed to qualitatively study turbulence transition [46,47] and predictability of chaotic flows [25,40]. The dynamics are governed by the nondimensional Navier-Stokes equations for forced incompressible flows

$$\frac{d\mathbf{v}}{dt} = -(\mathbf{v} \cdot \nabla)\mathbf{v} - \nabla p + \frac{1}{\text{Re}} \Delta \mathbf{v} + \mathbf{F}(y), \quad (1)$$

where $\mathbf{v} = (u, v, w)$ is the velocity, p is the pressure, Re is the Reynolds number, ∇ is the gradient, Δ is the Laplacian, $\mathbf{F}(y) = \sqrt{2}\pi^2/(4\text{Re}) \sin(\pi y/2)\mathbf{e}_x$ is the body forcing along x , y is the direction of the shear between the plates, and z is the spanwise direction (Fig. 1). We solve the flow in the domain $L_x \times L_y \times L_z$, where the boundary conditions are free slip at $y \pm L_y/2$ and periodic at $x = [0; L_x]$ and $z = [0; L_z]$ (Fig. 1). The body forcing, $\mathbf{F}(y)$, drives the flow for it to sustain turbulence with homogeneous boundary conditions [45]. To reduce the partial differential equations (1) to a set of ordinary differential equations, the velocity is projected onto compositions of Fourier modes [44], $\hat{\mathbf{v}}_i(\mathbf{x})$ as

$$\mathbf{v}(\mathbf{x}, t) = \sum_{i=1}^9 a_i(t) \hat{\mathbf{v}}_i(\mathbf{x}). \quad (2)$$

The ansatz (2) is substituted in the Navier-Stokes equations (1), which are projected onto the same modes, $\hat{\mathbf{v}}_i(\mathbf{x})$. This spawns nine nonlinear ordinary differential equations for the amplitudes, $a_i(t)$, which, in turn, become the unknowns of the system (see the Supplemental Material [48]). The system displays a chaotic transient, which ultimately converges to the stable laminar solution $a_1 = 1, a_2 = \dots = a_9 = 0$. During the transient, the system shows intermittent large bursts in the

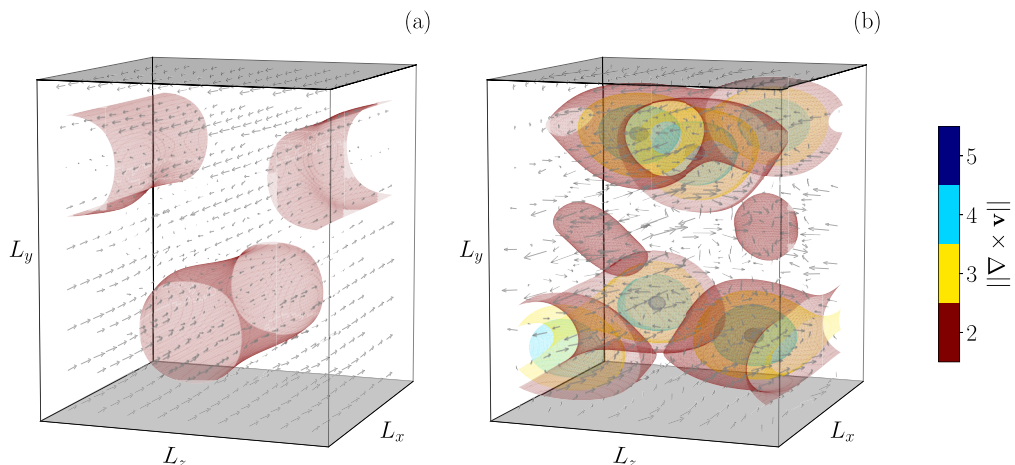


FIG. 1. Vorticity isosurfaces and velocity flowfield (a) before and (b) after an extreme event. The laminar structure (a) breaks down into vortices (b).

kinetic energy, which, in the MFE model, can be computed as

$$k(t) = \frac{1}{2} \sum_{i=1}^9 a_i^2(t). \quad (3)$$

The bursts in the kinetic energy are the extreme events we wish to predict [Figs. 1 and 2(a)]. The extreme events are characterized by a heavy tail of the distribution [11] [Fig. 2(b)]. In this work, we define an event as extreme extreme when the kinetic energy, $k(t)$, exceeds a threshold, k_e ,

$$k(t) \geq k_e, \quad (4)$$

where $k_e = 0.1$ is the user-defined extreme event's threshold [40]. Physically, an extreme event occurs because the stable laminar solution intermittently attracts the chaotic dynamics up to full laminarization. In reference to Fig. 3, during an extreme event, the flow slowly laminarizes (t_1 to t_3), but the laminar structure suddenly breaks down into vortices (t_4).

To generate extreme events, we set $L_x = 4\pi$, $L_y = 2$, $L_z = 2\pi$, and $\text{Re} = 400$ and follow the procedure of Srinivasan *et al.* [25]. We generate an ensemble of 2000 time series by integrating

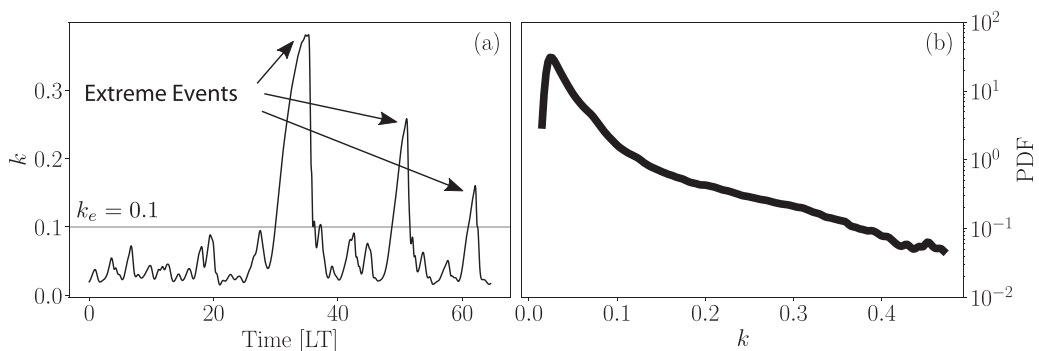


FIG. 2. (a) Typical evolution of the kinetic energy. Time is expressed in Lyapunov times (LT) (Sec. II A). (b) Probability density function of the kinetic energy.

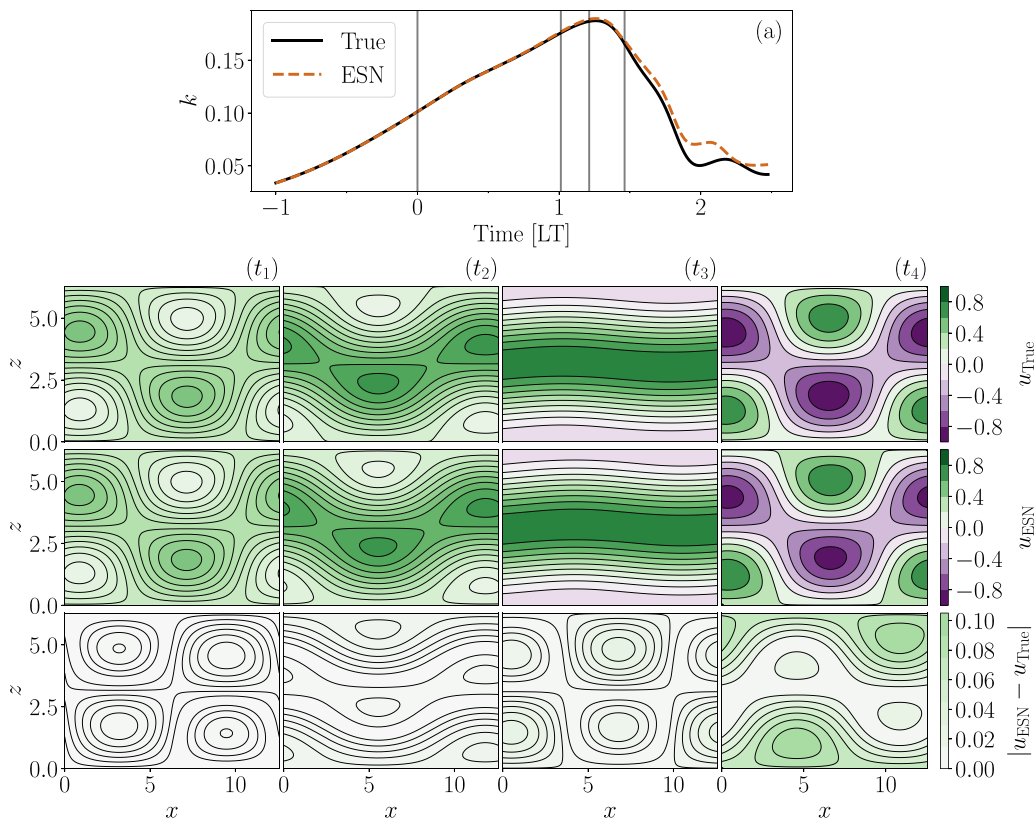


FIG. 3. Velocity component along x in the midplane during an event in the test set for the data, u_{True} , and an ESN, u_{ESN} . The times, (t_i) , of the flowfields are indicated by vertical lines in (a). Times t_3 and t_4 coincide with (a) and (b) of Fig. 1. The maximum of the absolute error, $|u_{\text{ESN}} - u_{\text{True}}|$, is approximately 5% the range of the flowfield.

the governing equations using a fourth-order Runge-Kutta scheme with time step $dt = 0.25$, where each time series consists of 4000 time units [≈ 65 LT (Lyapunov times)] (more details are given in the Supplemental Material [48]). Figure 2(a) shows a typical time series. We generate an ensemble of time series because we focus on the chaotic transient dynamics only. The ensemble of time series allows us to generate large amount of data, whereas a single long simulation may quickly laminarize, which would provide limited information about the chaotic transient. To create a training data set with rich dynamics, we discard all the time series that laminarize, i.e., the time series whose maximum kinetic energy is larger than the laminarization threshold, $k_l = 0.48$, where k_l is selected to be close to the (asymptotic) laminarization value, $k = 0.5$. Approximately $\simeq 28\%$ of all the time series are discarded in the ensemble due to laminarization. Hence, the data set consists of 1436 time series, which are used to train, validate, and test the echo state networks (Sec. III). Specifically, we use ten time series for training and validating an ensemble of echo state networks. Each network is trained on the same ten times series and receives the entire state of the system with downsampled $\delta t = 4dt$ [41] (more details are in the Appendix). We use an ensemble of ten networks due to the pseudorandom component of the initialization of the matrices in echo state networks (Sec. III). A summary of the parameters used in the generation of the data set is shown in Table I.

TABLE I. Summary of parameters used for data generation.

Parameter	Explanation	Value
L_x	Length of the domain in the x direction	4π
L_y	Length of the domain in the y direction	2
L_z	Length of the domain in the z direction	2π
Re	Reynolds number	400
dt	Integration time step	0.25
k_l	Laminarization threshold	0.48
k_e	Extreme event's threshold	0.1
	Number of time series	1436
	Length of each time series	4000
	Number of time series in the training set	10

A. Lyapunov exponent

In chaotic systems, almost all nearby trajectories¹ diverge at an average exponential rate given by the leading Lyapunov exponent of the system, Λ [17]. This means that any error in the prediction of the system grows exponentially in time at a rate given by Λ . Therefore, the leading Lyapunov exponent provides a timescale for assessing the time-accurate prediction of chaotic systems. We use Λ to normalize the time units in Lyapunov times, which is defined as $1\text{LT} = \Lambda^{-1}$. For the calculation of the Lyapunov exponent, we follow the procedure used by the authors of [49] in linear flow analysis. We consider a nonlinear autonomous dynamical system in the form of $\dot{\mathbf{q}} = \mathbf{f}(\mathbf{q})$, where \mathbf{q} is the system's state and \mathbf{f} is a nonlinear operator. The state is decomposed as $\mathbf{q}(t) = \bar{\mathbf{q}}(t) + \mathbf{q}'(t)$, where \mathbf{q}' is a first-order perturbation that develops on the unperturbed solution, $\bar{\mathbf{q}}(t)$. The time evolution of the perturbation is governed by

$$\mathbf{q}'(t) = \mathbf{M}(t, t_0)\mathbf{q}'(t_0), \quad \dot{\mathbf{M}}(t, t_0) = \mathbf{J}(t)\mathbf{M}(t, t_0), \quad (5)$$

where $\mathbf{M}(t, t_0)$ is the tangent propagator from t_0 to t , $\mathbf{J}(t) \equiv d\mathbf{f}(\mathbf{q})/d\mathbf{q}|_{\bar{\mathbf{q}}(t)}$ is the Jacobian of the system, and $\mathbf{M}(t_0, t_0) = \mathbf{I}$. In chaotic systems, perturbations grow in time at an average exponential rate given by the (positive) leading Lyapunov exponent, Λ , of the system [17]

$$\Lambda = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \left(\frac{\|\mathbf{q}'(t)\|}{\|\mathbf{q}'(t_0)\|} \right). \quad (6)$$

This means that tangent dynamics are unstable and that the norm, $\|\mathbf{q}'(t)\| = \|\mathbf{M}(t, t_0)\mathbf{q}'(t_0)\|$, becomes unbounded and causes numerical overflow. To overcome this problem, we use the QR algorithm proposed by Ginelli *et al.* [50] to compute the leading Lyapunov exponent. The QR algorithm overcomes the instability of tangent space by performing periodic orthonormalizations of the tangent propagator, $\mathbf{M}(t, t_0)$, thereby avoiding numerical overflow. Through the QR algorithm, we find the Lyapunov exponent to be $\Lambda = 0.0163$,² which implies that the Lyapunov time of the system is $1\text{LT} = 0.0163^{-1}$.

¹Specifically, ‘‘almost’’ means that the set of nearby trajectories that will not diverge in time has a Lebesgue measure of zero.

²The value $\Lambda = 0.0163$ is different from the Lyapunov exponent, $\Lambda = 0.0296$, computed by Srinivasan *et al.* [25]. This difference is because the latter is not the leading Lyapunov exponent, but a local Lyapunov exponent. The local Lyapunov exponent describes the average evolution of perturbations only from one specific initial condition along the trajectory, before nonlinearities saturate the norm of the perturbation. On the other hand, the leading Lyapunov exponent describes the average evolution of perturbations along the entirety of the trajectory. We further verify our results by computing the leading Lyapunov exponent as the average of the exponential growth of perturbations up to nonlinear saturation from different initial conditions along the trajectory and find, consistently with the QR method, $\Lambda = 0.0163$.

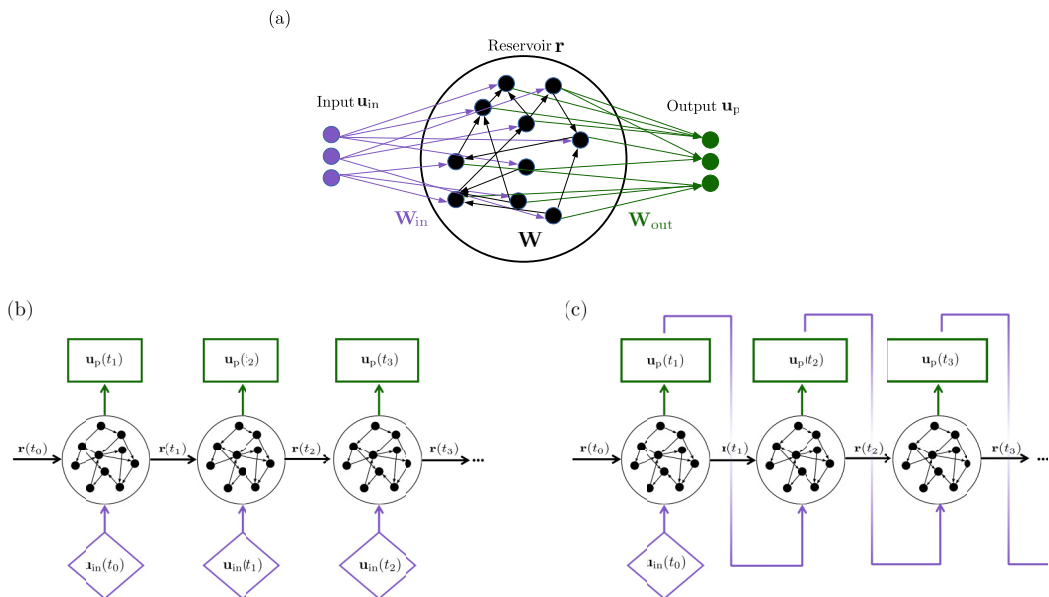


FIG. 4. (a) Schematic representation of the echo state network. (b) Open-loop and (c) closed-loop configurations.

III. ECHO STATE NETWORKS

In time-series prediction, the data are sequentially ordered in time. Recurrent neural networks are designed to infer the correlation within sequential data through an internal hidden state, which is updated at each time step. Because of the long-lasting time dependencies of the hidden state, however, training RNNs with backpropagation through time is notoriously difficult [51]. Echo state networks overcome this problem by nonlinearly expanding the inputs into a high-dimensional system, the reservoir, which also acts as the memory of the system [35]. Hence, the output is computed as a linear combination of the reservoir's dynamics, whose weights are the only trainable parameters of the system. Thanks to this specific architecture, training echo state networks consists of a straightforward linear regression problem, which bypasses backpropagation through time.

As shown in Fig. 4(a), in an echo state network, at any time t_i the input vector, $\mathbf{u}_{in}(t_i) \in \mathbb{R}^{N_u}$, is mapped into the reservoir state, by the input matrix, $\mathbf{W}_{in} \in \mathbb{R}^{N_r \times (N_u+1)}$, where $N_r \gg N_u$. The reservoir state, $\mathbf{r} \in \mathbb{R}^{N_r}$, which is updated at each time iteration as a function of the current input and its previous value, is used to compute the predicted output, $\mathbf{u}_p(t_{i+1}) \in \mathbb{R}^{N_u}$. The reservoir state and output are governed by the discrete dynamical equations, respectively,

$$\mathbf{r}(t_{i+1}) = \tanh[\mathbf{W}_{in}[\hat{\mathbf{u}}_{in}(t_i); b_{in}] + \mathbf{W}\mathbf{r}(t_i)], \quad (7)$$

$$\mathbf{u}_p(t_{i+1}) = [\mathbf{r}(t_{i+1}); 1]^T \mathbf{W}_{out}, \quad (8)$$

where $(\hat{\cdot})$ indicates normalization by the range component-wise, $\mathbf{W} \in \mathbb{R}^{N_r \times N_r}$ is the state matrix, b_{in} is the input bias and $\mathbf{W}_{out} \in \mathbb{R}^{(N_r+1) \times N_u}$ is the output matrix. Matrices \mathbf{W}_{in} and \mathbf{W} are (pseudo)randomly generated and fixed, while the weights of the output matrix, \mathbf{W}_{out} , are computed by training the network. The input matrix, \mathbf{W}_{in} , has only one element different from zero per row, which is sampled from a uniform distribution in $[-\sigma_{in}, \sigma_{in}]$, where σ_{in} is the input scaling. The state matrix, \mathbf{W} , is an Erdős-Rényi matrix with average connectivity d , in which each neuron (each row of \mathbf{W}) has on average only d connections (nonzero elements), which are obtained by sampling from a uniform distribution in $[-1, 1]$; the entire matrix is then rescaled by a multiplication factor

to set the spectral radius, $\rho \leq 1$ to enforce the echo state property [35]. The bias in the inputs and outputs layers are added to break the inherent symmetry of the basic ESN architecture [31,52]. The input bias, b_{in} is a hyperparameter, selected in order to have the same order of magnitude of the normalized inputs, $\hat{\mathbf{u}}_{\text{in}}$, while the output bias is determined by training the weights of the output matrix, \mathbf{W}_{out} .

The ESN can be run either in open loop or closed loop. In the open-loop configuration [Fig. 4(b)], we feed the data as the input at each time step to compute and store the reservoir dynamics, $\mathbf{r}(t_i)$. In the initial transient of this process, which is the washout interval, we do not compute the output, $\mathbf{u}_p(t_i)$. The purpose of the washout interval is for the reservoir state to satisfy the echo state property. In doing so the reservoir state becomes (1) up to date with respect to the current state of the system and (2) independent of the arbitrarily chosen initial condition, $\mathbf{r}(t_0) = 0$. After washout, we train the output matrix, \mathbf{W}_{out} , by minimizing the mean square error between the outputs and the data over the training set. Training the network on $N_{\text{tr}} + 1$ points needs neither backpropagation nor gradient descent, but it needs only a solution of a linear system (ridge regression)

$$(\mathbf{R}\mathbf{R}^T + \beta\mathbf{I})\mathbf{W}_{\text{out}} = \mathbf{R}\mathbf{U}_d^T, \quad (9)$$

where $\mathbf{R} \in \mathbb{R}^{(N_r+1) \times N_r}$ and $\mathbf{U}_d \in \mathbb{R}^{N_r \times N_r}$ are the horizontal concatenation of the reservoir states with bias, $[\mathbf{r}; 1]$, and of the output data, respectively; \mathbf{I} is the identity matrix and β is the Tikhonov regularization parameter [53].

In the closed-loop configuration [Fig. 4(c)], starting from an initial data point as an input and an initial reservoir state obtained after the washout interval, the output, \mathbf{u}_p , is fed back to the network as an input for the next time step prediction. In doing so, the network is able to autonomously evolve in the future. The closed-loop configuration is used during validation and test, but not during training.

A. Validation

During validation, we use part of the data to select the hyperparameters of the network by minimizing the logarithm of the mean square error of the kinetic energy with respect to the data (Sec. II). We use the logarithm because the error varies by multiple orders of magnitude with different hyperparameters. In this work, we tune the input scaling, σ_{in} , spectral radius, ρ , and Tikhonov parameter, β , which are the key hyperparameters for the performance of the network [35,36]. We use Bayesian optimization to select σ_{in} and ρ , and perform a grid search within $[\sigma_{\text{in}}, \rho]$ to select β [36]. We explore the hyperparameter space $[0.1, 10] \times [0.1, 1]$ for $[\sigma_{\text{in}}, \rho]$ in logarithmic scale and the grid $[10^{-6}, 10^{-9}, 10^{-12}]$ for β . The Bayesian optimization starts from a grid of 4×4 points in the $[\sigma_{\text{in}}, \rho]$ domain, and then it selects nine additional points through the gp-hedge algorithm [54]. We set $b_{\text{in}} = 0.1$, $d = 20$ and add Gaussian noise with zero mean and standard deviation, $\sigma_n = 0.01\sigma_u$, where σ_u is the standard deviation of the data component-wise, to the training and validation data [34]. Adding noise to the data improves the performance of ESNs in chaotic dynamics because the network explores the region around the attractor and becomes more robust to errors when forecasting the future evolution of the trajectory [34,35]. A summary of the hyperparameters is shown in Table II.

To select the hyperparameters, we use single shot validation (SSV) and recycle validation (RV) [36]. Single shot validation is the most commonly used validation strategy for RNNs in the literature [55]. It consists of splitting the data in a training set and a single small validation set subsequent in time to the training set [Fig. 5(a)]. The size of the validation set is limited by the chaotic nature of the signal, which causes the predicted trajectory to diverge exponentially in time from the data (Sec. II A). Because the hyperparameters are tuned with respect to a small interval, which is not representative of the available data set, the validated networks may perform poorly in the test set [36]. To improve the selection of hyperparameters, the K-Fold Cross-Validation (KCV) [56], which validates the network in multiple intervals of the data set, is sometimes used. The KCV is computationally expensive, because the network needs to be trained multiple times for each set of hyperparameters. To overcome the high computational cost of the KCV, while keeping the improved

TABLE II. Echo state networks' hyperparameters. Multiple values indicate that the parameter is optimized in the shown range.

Parameter		Value
ρ	Spectral radius	[0.1,1]
σ_{in}	Input scaling	[0.1,10]
β	Tikhonov parameter	$[10^{-6}, 10^{-9}, 10^{-12}]$
d	Connectivity	20
b_{in}	Input bias	0.1
σ_n	Standard deviation of the noise in the training inputs	$0.01\sigma_u$

performance, the authors of [36] proposed the recycle validation (RV). In the RV, the network is trained only once on the entire data set, and validation is performed on multiple intervals already used for training. This is possible because recurrent neural networks operate in two configurations (open loop and closed loop), which means that the networks can be validated in closed loop on data used for training in open loop [Fig. 5(b)].

IV. TIME-ACCURATE PREDICTION OF EXTREME EVENTS

The time-accurate prediction of extreme events is the prediction in the near future of the occurrence of an extreme event given data at the current time. Forecasting an incoming extreme event can be viewed as a *binary classification* problem, in which we evaluate whether (or not) the event will happen. Because extreme events are also rare, the classification operates on an imbalanced data set, in which the occurrence of an extreme event is significantly less likely than the nonoccurrence of an event. To assess the capability of echo state networks to time accurately predict extreme events, therefore, we tailor figures of merit that take into account the imbalanced nature of the data set [57]. In this section, we (1) assess the performance of the network to accurately predict extreme events from multiple points along the test set through precision, recall, and the F-score; (2) compute how far in advance we can foretell extreme events through the prediction horizon; and (3) extend the analysis to multiple Reynolds number to verify the robustness of the networks and the results.

A. Precision and recall

We assess the capability of the networks to predict extreme events. To do so, we monitor the system in real time by letting the network evolve autonomously (closed loop) to forecast the future behavior of the flow in the test set. We define the Prediction Time (PT) to evaluate if the network is




FIG. 5. Schematic representation of the validation strategies, where \mathbf{u} represents the number of degrees of freedom of the data at a single time step. (a) Single shot validation and (b) recycle validation. Three validation intervals are shown in the recycle validation [36].

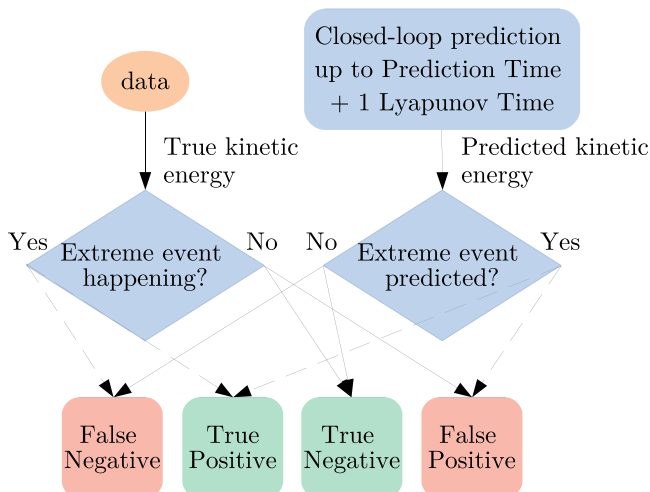


FIG. 6. Schematic representation of the prediction of extreme events with echo state networks. We evaluate whether the predicted kinetic energy and the true kinetic energy display an extreme event in the interval lasting 1 LT after the prediction time. A more detailed block diagram can be found in the Supplemental Material [48].

predicting an extreme event in a specific window in the future. For a user-defined prediction time, we evaluate if the prediction generated by the network crosses the extreme event's threshold, k_e , in the 1 LT interval subsequent to the prediction time. The prediction time determines how far ahead into the future we are assessing the forecast of the network. For example, with a $PT = 2$ LTs, we evaluate if the network is forecasting an event between 2 LTs and 3 LTs in the future. There are three cases we consider: (1) if the network's prediction crosses k_e and the true signal crosses k_e , we have a true positive; (2) if the network's signal crosses k_e and the true signal does not, we have a false positive; and (3) if the network's signal does not cross k_e and the true signal crosses k_e , we have a false negative (Figs. 6 and 7).

To evaluate the performance of the networks to forecast extreme events, we use precision, P , recall, R , and the F-score, F , which are defined as [57]

$$P = \frac{TP}{TP + FP}, \quad R = \frac{TP}{TP + FN}, \quad F = \frac{2}{P^{-1} + R^{-1}}, \quad (10)$$

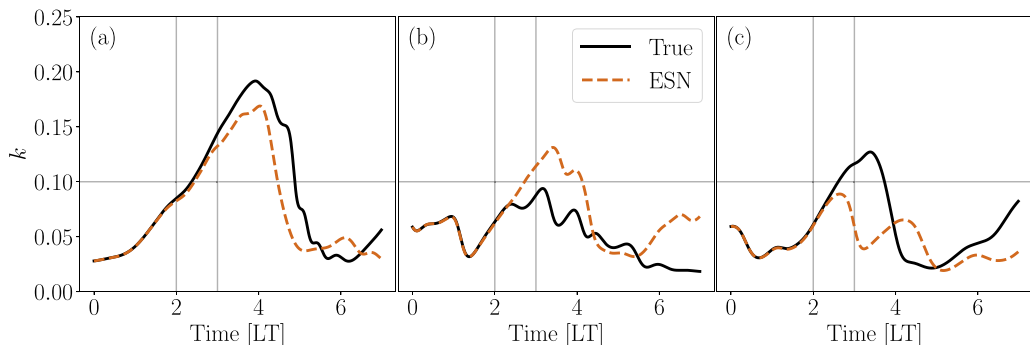


FIG. 7. (a) True positive, (b) false positive, and (c) false negative predictions of an extreme event with prediction time equal to 2 LTs. The horizontal line indicates the extreme event threshold, k_e . The vertical lines indicate the time window of the prediction.

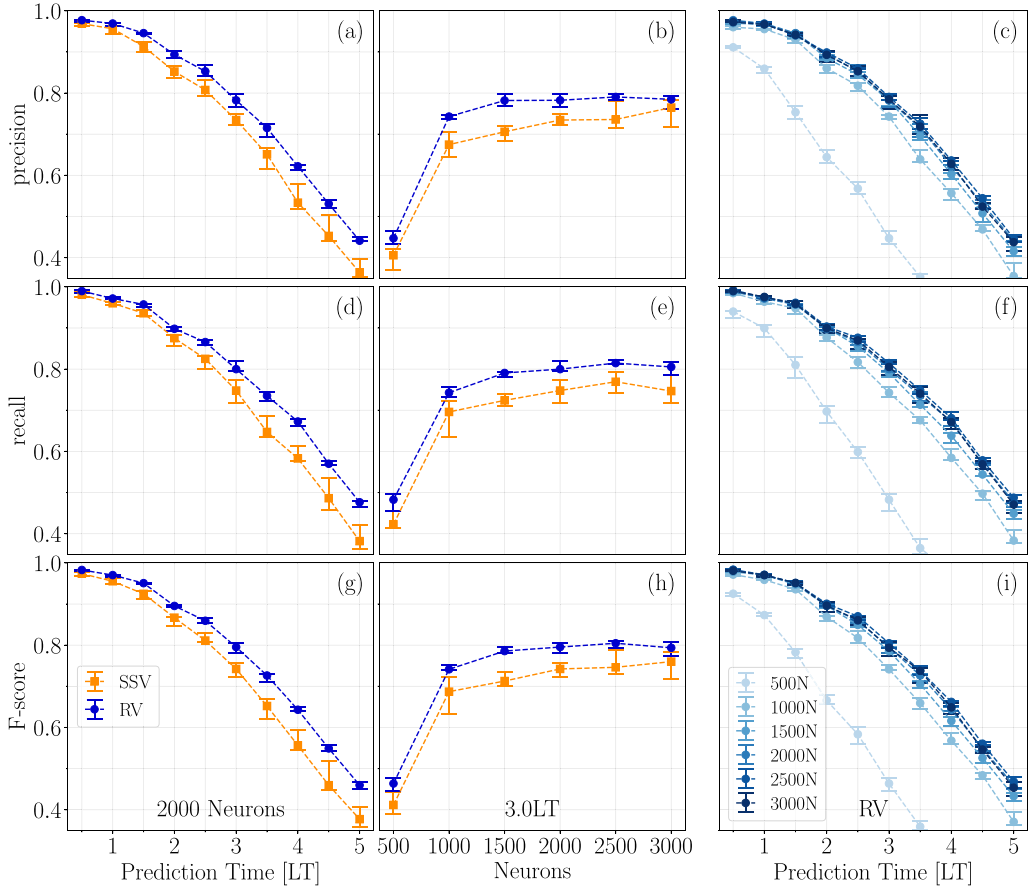


FIG. 8. 25th, 50th, and 75th percentiles of precision, recall, and F-score, as a function of the prediction time and number of neurons for the single shot validation (SSV) and recycle validation (RV). Panels (a), (d), and (g) have fixed number of neurons, (b), (e), and (h) have a fixed prediction time of 3 LTs, and (c), (f), and (i) are for the RV only.

where TP, FP, and FN are true positives, false positives, and true negatives, respectively. Precision measures how many of the predicted events actually happen; recall measures how many events are predicted out of the total number of events, which gives an indication on how many events are missed out by the prediction; and the F-score is the geometric average of the two. Therefore, $P = 1$ means that all the predicted events happened, $R = 1$ means that all the events have been predicted, and $F = 1$ means that all the events were predicted and all predictions were correct.

Figure 8 shows the results obtained by predicting the occurrence of extreme events for different prediction times in the test set. The results are obtained from 9000 starting points, which span 150 time series and consist of 360 extreme events. We select the starting points to be spaced 1 LT from each other. First, the recycle validation outperforms the single shot validation in all the figures of merit, which implies that networks optimized with RV correctly predict more events than networks optimized with SSV. In large reservoirs (neurons ≥ 1500), the recycle validation improves the F-score with respect to the single shot validation by an amount comparable to the improvement obtained by increasing the size of the reservoir by hundreds of neurons. For example, the 1500-neuron networks trained with the recycle validation have a median of the F-score of $F_{RV} = 0.79$, which is greater than the median of the F-score of the 3000-neuron networks trained




FIG. 9. 25th, 50th, and 75th percentiles for the F-score as a function of the extreme event threshold and prediction time for 2000-neuron networks. Blue round markers indicate recycle validation, and orange square markers indicate single shot validation.

with single shot validation ($F_{SSV} = 0.76$) [Fig. 8(h)]. Second, the networks are able to accurately predict extreme events multiple Lyapunov times in advance, with $F \geq 0.9$ for 2 LT prediction time and up to $F = 0.8$ for 3 LT prediction time [Fig. 8(g)]. This means that we predict almost all the events ($\simeq 90\%$) [Fig. 8(d)] with only a few false positives ($\simeq 10\%$) [Fig. 8(a)] 2LTs before the events happen. Third, Figs. 8(c), 8(f), and 8(i) display the performance for the recycle validation as the reservoir size increases. After a significant improvement in the performance between 500 and 1000 neurons, the performance of the networks remains approximately constant for networks with more than 1500 neurons. This indicates that 1500 neurons are able to learn the chaotic dynamics needed to forecast the evolution of the MFE system. Crucially, the marked difference in performance between different validation strategies [Figs. 8(b), 8(e), and 8(h)] means that for large reservoirs tuning the hyperparameters is more effective than increasing the number of neurons.

Figure 9 shows the F-score as a function of the extreme event threshold, k_e , for the ensemble of 2000-neuron networks. We analyze different values of k_e to verify the robustness of echo state networks with respect to the threshold. We observe qualitatively similar results to the case $k_e = 0.1$. The recycle validation outperforms the single shot validation in all figures of merits. The networks are able to predict extreme events at different k_e . These results indicate that the networks' performance does not change significantly with k_e and, therefore, the networks are robust with respect to the threshold that defines an extreme event.

B. Prediction horizon of extreme events

To assess how far in advance the networks predict extreme events, we use the prediction horizon of extreme events, PH_{ee} . The prediction horizon is a metric to evaluate the time-accurate prediction of chaotic dynamics with data-driven methods [34,36,37,40]. In this work, it is defined as the time interval during which the absolute error between the kinetic energy predicted by the ESN, k_{ESN} , and the true kinetic energy, k_{True} is bounded as

$$\frac{|k_{\text{ESN}}(t) - k_{\text{True}}(t)|}{k_e - \bar{k}} < 0.2, \quad (11)$$

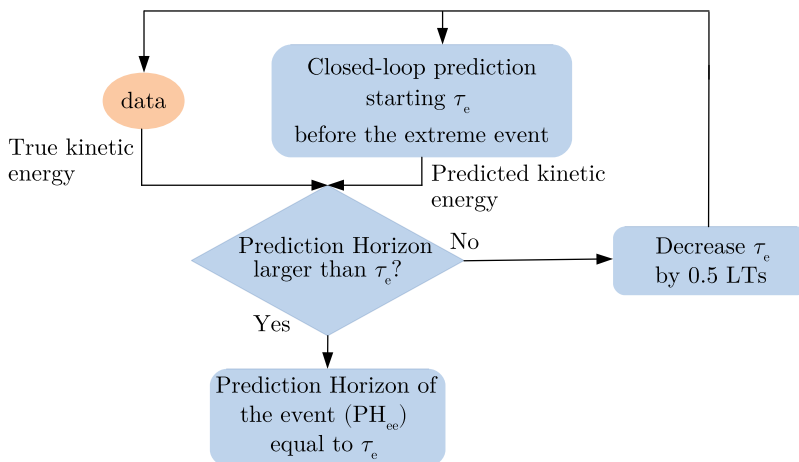


FIG. 10. Block diagram for the computation of the prediction horizon of extreme events, PH_{ee} . We evaluate the prediction horizon (11) of the closed-loop prediction starting from τ_e before the event. If the prediction horizon is larger than τ_e , $\text{PH}_{\text{ee}} = \tau_e$; if not, we decrease τ_e and evaluate the prediction horizon starting from the new τ_e before the event. The time τ_e is initialized to 10 LTs. A more detailed diagram is shown in the Supplemental Material [48].

where \bar{k} is the time average of the kinetic energy, k_e is the extreme event threshold, $k_e - \bar{k}$ acts as a normalization factor, and 0.2 is a user-defined threshold.

We define the prediction horizon of extreme events, PH_{ee} , by identifying the instants $t^{(e)}$ at which extreme events start, such that $k(t) = k_e$ and $dk(t)/dt > 0$ at $t = t^{(e)}$. The value of $t^{(e)}$ is obtained by preprocessing the test set data. For the i th event, we start a closed-loop prediction of the ESN at time $t_i^{(e)} - \tau_e$, $\tau_e = 10\text{LTs}$, and evaluate whether the prediction horizon is larger than τ_e . If the prediction horizon is larger than τ_e , then PH_{ee} is equal to τ_e , if the prediction horizon is smaller than τ_e , we decrease τ_e by 0.5 LTs, evaluate whether the prediction horizon is larger than τ_e , and so on (Figs. 10 and 11).

Figure 12 shows the average PH_{ee} over 500 extreme events for different sizes of the reservoirs. The networks are able to predict the extreme events up to almost 5 LTs in advance on average. The performance increases sharply in smaller networks, and levels off for larger networks. The recycle validation outperforms the single shot validation in all the figures of merit, and in most




FIG. 11. Last four evaluations before the event is predicted with $\text{PH}_{\text{ee}} = 6.5$ LTs. In each panel, the prediction starts from 0.5 LT closer to the extreme event with respect to the previous panel.

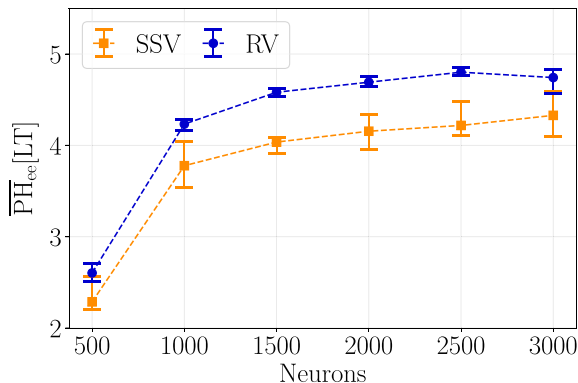


FIG. 12. 25th, 50th, and 75th percentiles for the average prediction horizon of extreme events as a function of the number of neurons for the single shot validation (SSV) and the recycle validation (RV).

cases generates an improvement in performance comparable to the increase in size of the reservoir of hundred, and even thousands, of neurons. For example, the 1000-neuron networks trained with recycle validation predict 50% of the events (median of $\overline{\text{PH}}_{\text{ee-RVC}} = 4.23$ LTs) with approximately the same prediction horizon of the 2500-neuron networks trained with single shot validation (median of $\overline{\text{PH}}_{\text{ee-SSV}} = 4.22$ LTs).

C. Analysis at different Reynolds numbers

In order to further corroborate the results of Secs. IV A and IV B, we analyze the time-accurate prediction of the MFE system's extreme events for different Reynolds numbers, $\text{Re} = [600, 800, 1000]$, for which the Lyapunov times are $\text{LT} = [0.0135^{-1}, 0.0126^{-1}, 0.0121^{-1}]$, respectively. Due to the reduced-order nature of the model, the MFE captures only the near wall cycle in turbulence, in which the flow is independent of the Reynolds number. However, increasing the Reynolds number affects the behavior of the low-order model. The flow still exhibits a transient that eventually converges to the laminar solution, but extreme events (and laminarization) happen less frequently [58]. Therefore, we study this enlarged range of Reynolds numbers to verify the robustness of echo state networks in time accurately predict chaotic dynamics with respect to changes of the physical parameter in the governing equations. For each Reynolds number, the networks are retrained and validated (see Sec. II and the Appendix).

Figure 13 shows the results for the prediction horizon of extreme events and the F-score. On the one hand, increasing the Reynolds number increases the prediction horizon of extreme events

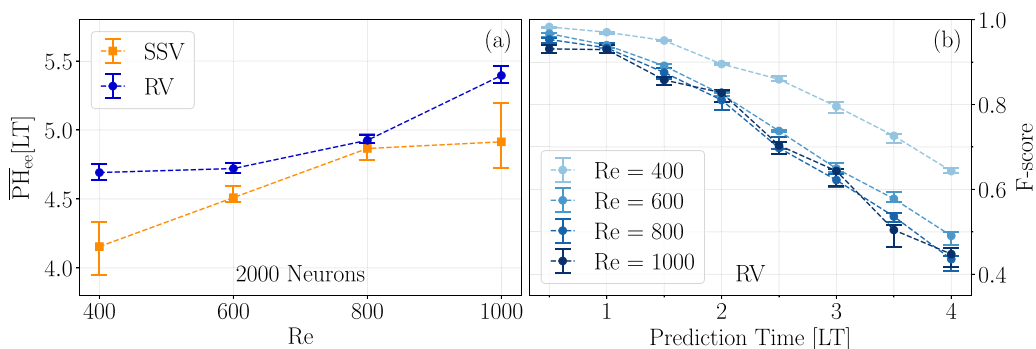


FIG. 13. 25th, 50th, and 75th percentiles for the (a) average prediction horizon of extreme events and (b) F-score as a function of the Reynolds number for 2000-neuron networks.

[Fig. 13(a)]. On the other hand, the F-score deteriorates from $Re = 400$ to the other cases, which perform similarly one to the other [Fig. 13(b)]. In agreement with the observed behavior in the $Re = 400$ case, the recycle validation outperforms the single shot validation in all figures of merit in the new range of Reynolds values. These results indicate that the analysis for the $Re = 400$ case extends to a broader range of Reynolds numbers, which means that ESNs are robust machines with respect to changes in the physical parameters of the system.

V. STATISTICAL PREDICTION OF EXTREME EVENTS

In Sec. IV, we analyze time-accurate predictions. In this section, we move our focus on to the flow long-term statistics. In ergodic systems, the flow statistics are robust quantities, which are not affected by the butterfly effect, in contrast to the time-accurate prediction.³ Physically, nearby trajectories evolve within the same bounded region of the phase space, which means that they share the same long-term statistics.

We analyze the capability of the networks to predict the statistical behavior of the flow through long-term predictions. Long-term predictions are closed-loop predictions that last several tens of Lyapunov times and diverge from the true data. The long-term predictions are generated from 500 different starting points in the training set, from which we generate a total of 500 different time series per network (one per starting point). Each time series is obtained by letting the ESN evolve in closed loop for 4000 time units ($\simeq 65$ LTs). We discard all the time series that laminarize and compute the statistics produced by the ESN using the remaining time series. This procedure is analogous to the generation of the data through the numerical integration of the governing equations discussed in Sec. II. We use the long-term prediction to (1) assess the prediction of the Probability Density Function (PDF) of the kinetic energy for different sizes of the reservoir and (2) predict the statistics of flow variables in the physical domain $L_x \times L_y \times L_z$. In all cases, we compare the statistical prediction of the networks with the statistics of the training set. We do so because the objective of predicting the statistics is to improve our knowledge via the networks from the already available knowledge in the training data. This is a markedly different approach from previous works such as Srinivasan *et al.* [25], where the large data sets available during training already have converged statistics, and the networks are trained to replicate the statistics, but not to extrapolate them. In this study, we use small training data sets with nonconverged statistics to train the network and extrapolate the statistics. This improves the statistical knowledge of the flow with respect to the available data.

To assess the ESN performance on the prediction of the statistics of the kinetic energy, we use the Kantorovich metric, the mean logarithmic error and relative errors of the high-order moments of the PDF, which are computed with respect to the true PDF of the kinetic energy, $\text{PDF}_{\text{True}}(k)$. The first-order Kantorovich metric [59], also known as the Wasserstein distance or Earth mover's distance, measures the distance between two PDFs. Intuitively, it evaluates the minimum work to transform one PDF to another. It is computed as

$$\mathcal{K} = \int_{-\infty}^{\infty} |\text{CDF}_{\text{True}}(k) - \text{CDF}_j(k)| dk, \quad (12)$$

where CDF is the cumulative distribution function and j indicates the PDF that we are comparing with the true data. We use the Kantorovich metric to assess the accuracy on the prediction of the PDF of the kinetic energy. If \mathcal{K} is small, the two PDFs are similar to each other, while if \mathcal{K} is large,

³In chaotic flows, the trajectories cannot be time accurately predicted indefinitely because the tangent space of chaotic flows is exponentially unstable (butterfly effect). This is an intrinsic limit to which the time-accurate prediction of chaos is subject (see Sec. II A).

the two PDFs significantly differ one from another. The mean logarithmic error is defined as

$$\text{MLE} = \sum_{i=1}^{n_b} n_b^{-1} |\log_{10}(\text{PDF}_{\text{True}}(k)_i) - \log_{10}(\text{PDF}_j(k)_i)|, \quad (13)$$

where the sum is performed over the total number of bins, n_b , used in the PDF. In case a bin has a value equal to zero and the logarithm is undefined, we saturate the logarithmic error in the bin to 1. The MLE is able to assess the accuracy on the prediction of the tail of the PDF because the logarithm amplifies the errors in the small values of the tail. The i th standardized moment of the kinetic energy, $\tilde{\mu}_i$, is

$$\tilde{\mu}_i = \mathbb{E} \left[\left(\frac{k - \bar{k}}{\sigma(k)} \right)^i \right], \quad (14)$$

where $\mathbb{E}[\]$ is the expectation operator, \bar{k} is the mean of the kinetic energy, and $\sigma(k)$ is the standard deviation. High-order moments of the distribution are significantly affected by values of the kinetic energy far from the mean, thus, they describe the tail of the PDFs. The third and fourth standardized moments, which are the skewness and kurtosis, respectively, are commonly used metrics to characterize the tail of the distributions [60]. In this work, we assess the prediction of the tail of the distribution with the relative error of these two high-order moments, $\text{RE}(\tilde{\mu}_i) = |(\tilde{\mu}_i - \tilde{\mu}_{i_{\text{True}}})/\tilde{\mu}_{i_{\text{True}}}|$.

Figure 14 shows the statistics of the kinetic energy for the ESNs optimized with the recycle validation and the training set. Figure 14(a) shows the true PDF, the PDF predicted by the ESN and the training set PDF. Although the statistics of the training set is not close to the true statistics, the network learns the tail of the distribution by extrapolation. In other words, the echo state networks have learned a good approximation of the converged statistics from a relatively short training set. Figures 14(b), 14(c), 14(d), and 14(e) show the Kantorovich metric, the mean logarithmic error, and the relative error for skewness and kurtosis for an ensemble of ten networks. The networks (1) improve significantly the statistics with respect to the training set in all cases analyzed and (2) perform slightly better when increasing their size. This means that the networks are able to both learn from an imperfect data set and improve the prediction of the overall dynamics of the flow, even in small reservoirs. These results extend the findings of Racca and Magri [41] to smaller sizes of the reservoir.

Finally, we study the statistics of the flow variables, following Srinivasan *et al.* [25] and Doan *et al.* [40]. We assess the prediction for the generic variable α through the Normalized Root-Mean-Square Error (NRMSE) with respect to the true data,

$$\text{NRMSE}(\alpha) = \frac{\sqrt{\frac{1}{N} \sum_i (\alpha_i - \alpha_{\text{True}_i})^2}}{\max(\alpha_{\text{True}}) - \min(\alpha_{\text{True}})}, \quad (15)$$

where $[\max(\alpha_{\text{True}}) - \min(\alpha_{\text{True}})]$ is the range of α_{True} . Figures 15(d)–15(f) show the NRMSE for the time, $(\bar{\ })$, and spatial (along x and z), $\langle \langle \rangle \rangle$, average of the streamwise velocity, u , its square, u^2 , and Reynolds stress, uv , whose profiles are plotted in Fig. 15(a)–15(c). As seen for the PDF of the kinetic energy, (1) the ESN improves the statistics with respect to the available data and (2) networks of different sizes perform similarly. Overall, there is a smaller improvement in the statistics of the flowfield than in the statistics of the kinetic energy of Fig. 14. This is because the networks are validated with a focus on the kinetic energy (Sec. III A). The time average of the square of the velocity along x in the midplane is shown in Fig. 16(a) for the entire data. As seen for the spatial averages of Fig. 15, the error for a network with 2000 neurons [Fig. 16(b)] is smaller than the error of the training set [Fig. 16(c)]. These results show that echo state networks improve the prediction of the statistics of the flowfield in the physical space with respect to the training data.

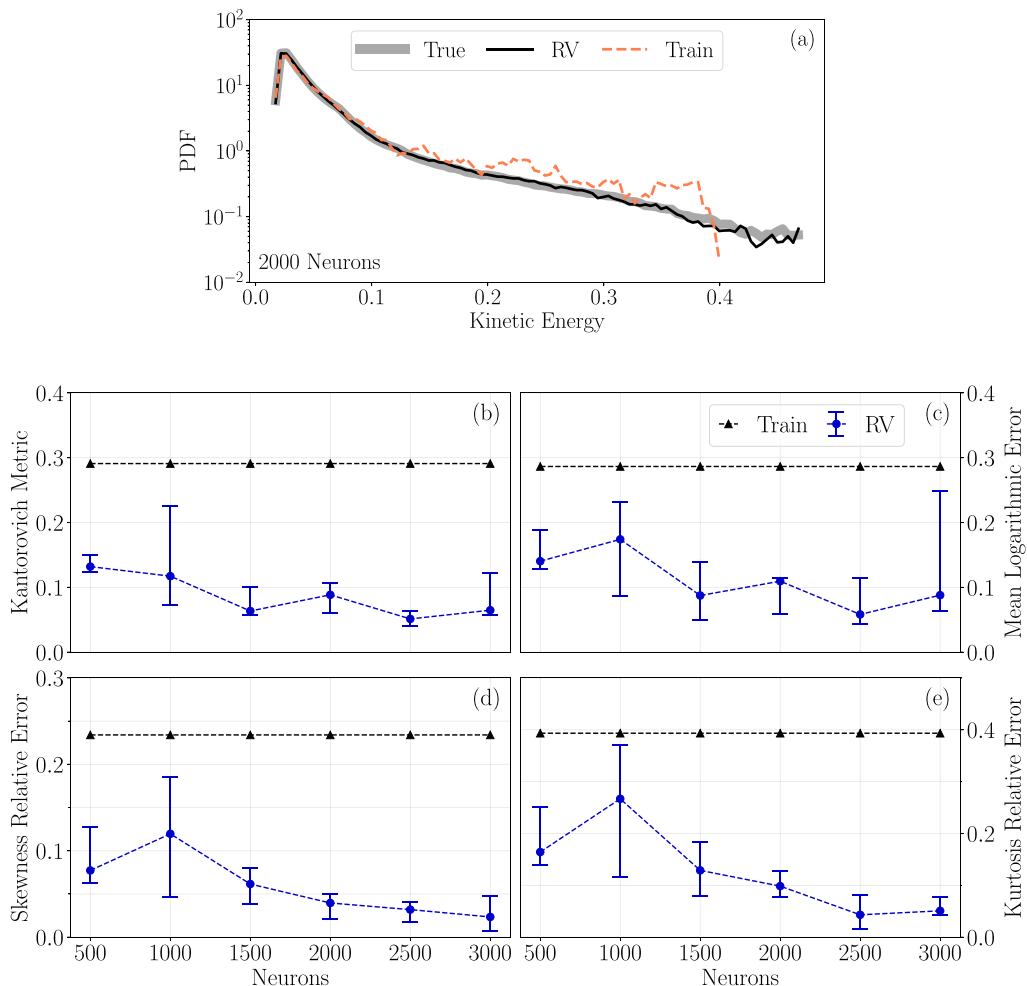


FIG. 14. (a) Probability density function (PDF) of the kinetic energy for the entire time series (True), training set (Train), and best performing 2000-neuron ESN for the recycle validation (RV). 25th, 50th, and 75th percentiles of (b) the Kantorovich metric, (c) the mean logarithmic error and (d), (e) relative errors of the high-order moments for the ESNs ensemble and training set.

VI. CONTROL OF EXTREME EVENTS

In this section, we use the time-accurate prediction of echo state networks (Sec. IV A) to control and suppress extreme events before they occur. We propose an active control strategy as a proof of concept. We (1) develop a simple and effective control strategy based on the predictions of the networks, (2) analyze the effects of the control strategy on decreasing the probability of the occurrence of extreme events at different Reynolds numbers, and (3) compare the ESN-based control strategy with other possible strategies.

We actuate the control strategy when the ESN forecasts an extreme event in the future time window lasting 1 Lyapunov time (LT) after the prediction time (PT) (see Sec. IV A). The control strategy consists of increasing the Reynolds number to 2000 for 1.5 LTs, starting from $PT - 1LT$ (Fig. 17). We choose the value 1.5 LTs because it is the average time between the start of the control strategy and the event. Physically, increasing the Reynolds corresponds to decreasing the forcing

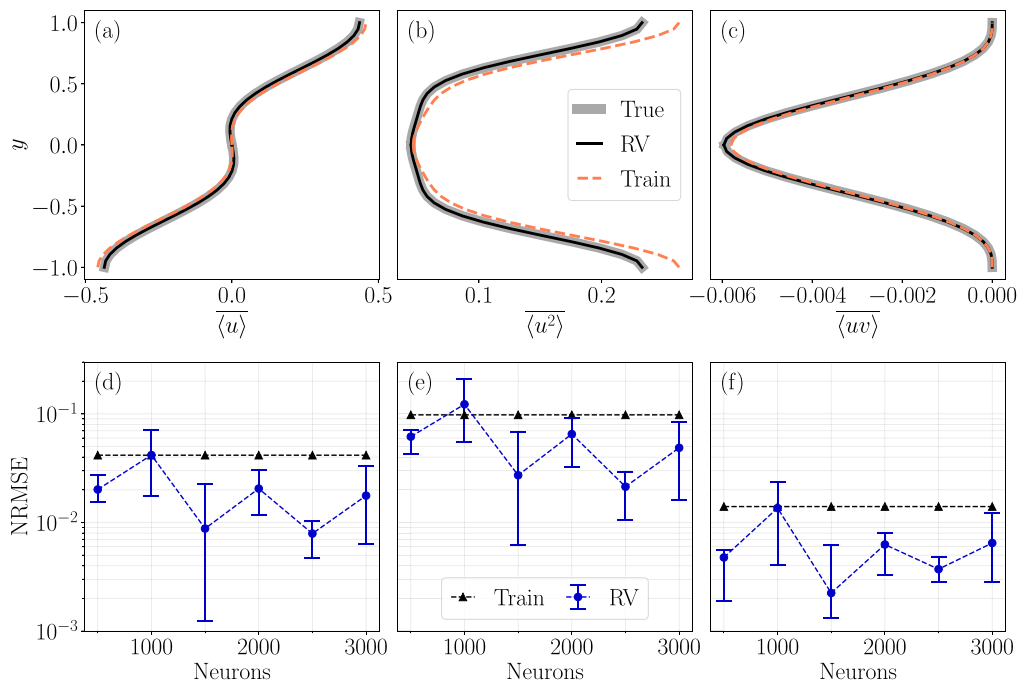


FIG. 15. (a) Profile of the spatial and time average of the streamwise velocity, (b) its square, and (c) Reynolds stress. ESNs profiles are the best performing networks with 2000 neurons. (d)–(f) 25th, 50th, and 75th percentiles of the NRMSE in the variables from (a)–(c) for the network ensemble and training set.

term and the dissipation, which in turn results in a slower change in time of the total energy of the flow [45]. Because we act on the flow before the event happens, reducing the increase of the energy of the flow prevents the kinetic energy from increasing and passing the extreme event threshold. An example of a controlled extreme event is shown in Fig. 18. The control strategy is effective and the event is suppressed. This results in significantly lower values of the velocity flowfield in the suppressed case with respect to the uncontrolled case (times t_3 – t_4). Practically, the suppressed time series in panel (a) is obtained by reintegrating the governing equations from 1 LT, with $\text{Re} = 2000$

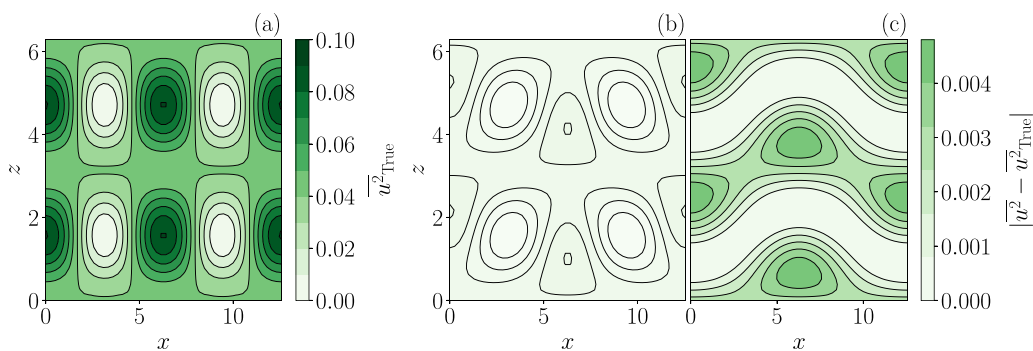


FIG. 16. (a) Time average of the square of the midplane velocity along x . Error with respect to the true data for (b) a 2000-neuron network, and for (c) the training data.

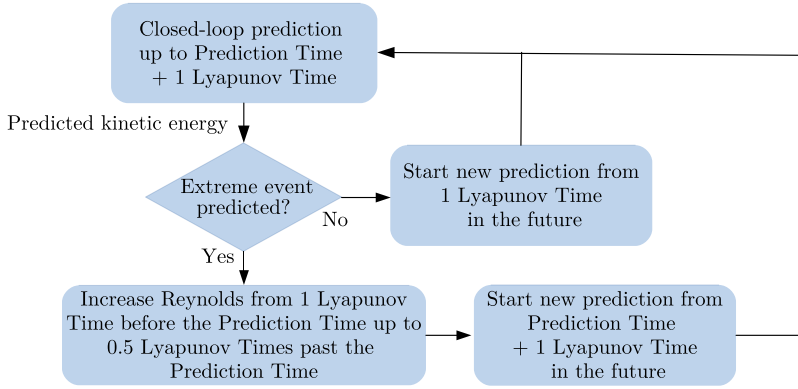


FIG. 17. Schematic representation of the control strategy. If the network predicts an event in the 1 LT interval after the prediction time, we increase the Reynolds number before the event happens to suppress it. A more detailed schematic is shown in the Supplemental Material [48].

from 1 to 2.5 LTs and $Re = 400$ from 2.5 LTs onwards. The new time series becomes the true data, which is used to monitor the flow to predict future extreme events.

Figure 19 shows the effect of the control strategy on the average PDF of the kinetic energy for the 2000-neuron network ensemble trained with recycle validation at different Reynolds numbers. The tail of the PDF decreases in all cases analyzed by up to one order of magnitude for values

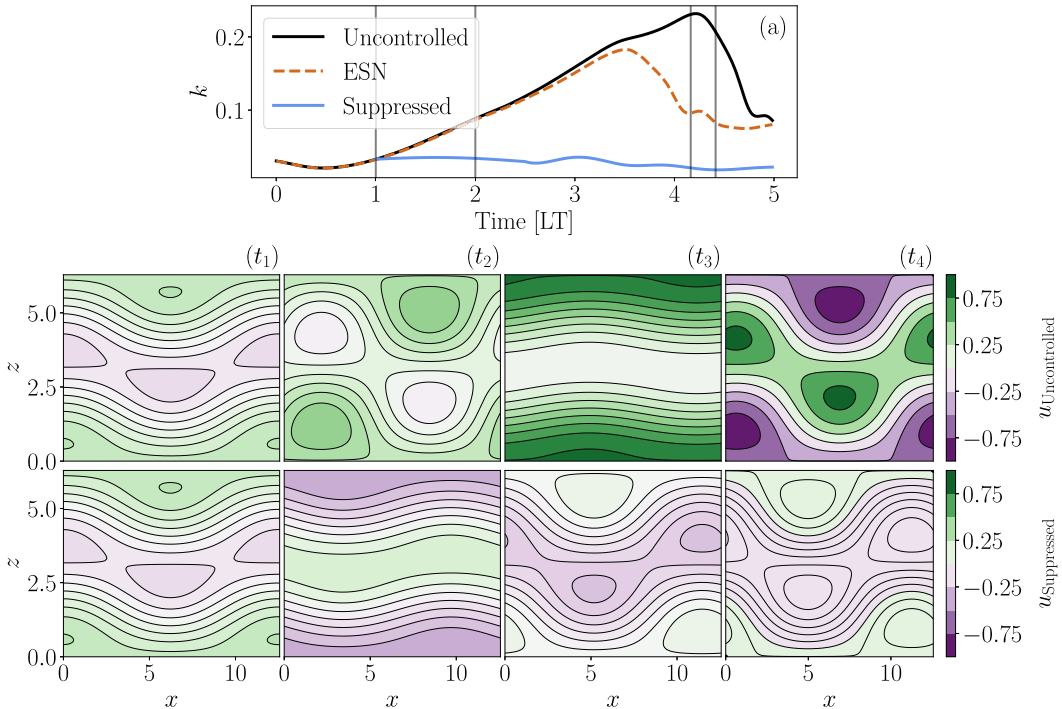


FIG. 18. Velocity component along x in the midplane with and without control during an extreme event. The times, (t_i) , of the flowfields are indicated by the vertical lines in (a). The event is predicted in the interval [2 LTs, 3 LTs] (PT = 2 LTs), and control is active from 1 LT to 2.5 LTs.




FIG. 19. Probability density function of the kinetic energy with and without control for different prediction times (PTs) and Reynolds numbers.

larger than the extreme event threshold, $k_e = 0.1$. This physically means that we are able to predict the vast majority of the events and suppress them. However, the tail does not go to zero for two reasons: (1) not all events are predicted by the networks and (2) the control strategy is not perfect, i.e., some predicted events are not suppressed. A predicted event is not suppressed if the kinetic energy crosses the extreme event threshold in the interval before the next prediction is made, which is for $t < 2PT + 1LT$ since the previous prediction. For example, in the event of Fig. 18, the next ESN prediction starts at 3 LTs and predicts whether an event is occurring between 5 and 6 LTs. If the extreme event occurs in the window 1–5 LTs, the event is correctly predicted by the ESN, but not suppressed by the control strategy. Overall, there is an improvement in the suppression of the tail as the prediction time decreases because the number of false negatives increases with the prediction time [Fig. 8(d)]. This is due to the butterfly effect that hinders the time-accurate prediction for large times. This implies that fewer events are predicted and controlled when the networks try to predict the flow for larger times.

To assess the efficacy of controlling extreme events through the prediction of the ESNs, we compare the control strategy actuated on the ESN predictions with ideal and random actuation. In ESN-based control, the ESN prediction activates the control strategy. In ideal control, by preprocessing the test data, we identify the instants at which extreme events starts, and we actuate the control strategy based on these instants. In random control, the control strategy has a fixed probability to be activated at each Lyapunov time, which we select in order for the average time of actuation to be similar to the ESN-based and ideal actuation. In all cases, the control strategy is the same: the Reynolds number is increased to 2000 for 1.5 LTs; the difference is the condition that triggers the activation of the control strategy.

Figure 20(a) shows the PDF of the kinetic energy for the different cases. The ideal actuation performs slightly better than the ESN-based strategy, while the random actuation performs poorly. Figure 20(b) shows the average period between extreme events and between actuations of the control strategy. The objective is to compare the results with the uncontrolled system, in which extreme events occur on average every 19 LTs (more details are in the Supplemental Material [48]). On the one hand, ESN-based and ideal control greatly improve the average period between extreme events, with average periods of 91 LTs and 99 LTs, respectively. On the other hand, random control only marginally suppresses extreme events, with an average period of 21 LTs. The computational cost of

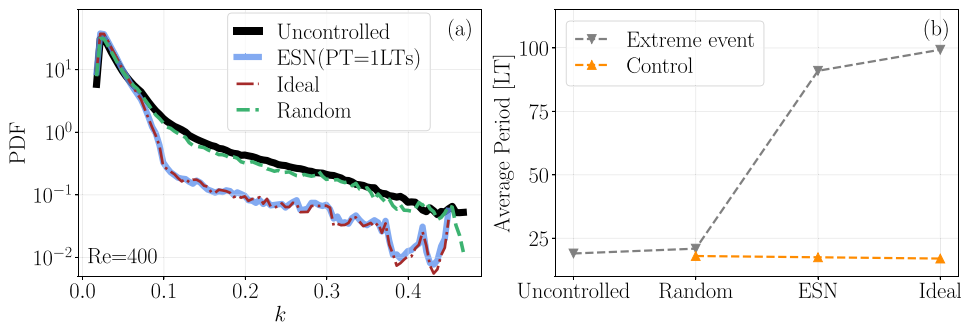


FIG. 20. (a) Probability density function of the kinetic energy for different control scenarios and (b) average period between extreme events and between activations of the control strategy.

actuating the three strategies is qualitatively the same: ideal control is on average actuated every 17.0 LTs, ESN-based control every 17.5 LTs, and random control every 18.0 LTs. These results show that ESN-based control (1) outperforms control methods that are unaware of incoming extreme events (random control) and (2) is close to the optimal suppression of extreme events (ideal control).

VII. CONCLUSIONS

We propose a data-driven framework and methodology to predict and control extreme events in nonlinear flows. The data-driven framework is based on echo state networks, which are a type of reservoir computing. We analyze a prototypical chaotic shear flow, in which extreme events manifest themselves as sudden bursts of the kinetic energy, which are caused by an intermittent competition between laminar and chaotic dynamics.

First, we show that a careful hyperparameter optimization is key to the networks' precision, recall, and robustness. We deploy the recycle validation for hyperparameter selection, which is a principled strategy based on a chaotic timescale. The recycle validation exploits Bayesian sampling, it is computationally cheap, and it allows the machine to accurately predict more extreme events than other hyperparameter-selection techniques (such as the commonly used single shot validation). Second, we analyze the networks' capability of time accurately predict extreme events before they occur. We show that the networks correctly predict extreme events up to five Lyapunov times in advance with large precision (up to 98%) and recall (up to 99%). Physically, the networks accurately forecast the incoming attempts of the flow to laminarize and readjust itself through chaotic bursts. Third, we analyze the networks' capability of inferring the long-term statistics of the kinetic energy and the velocity from training data with nonconverged statistics. Fourth, we leverage on the echo state networks' accurate predictions to control the flow and prevent extreme events from occurring. The control strategy is simple and effective. The network provides a precursor of the event, which activates a control strategy that modifies the Reynolds number to decrease the energy growth rate. We reduce the occurrence of extreme events by up to one order of magnitude with respect to the uncontrolled flow, which physically reduces the frequency of the flow laminarization attempts. Fifth, by investigating the chaotic shear flow at different Reynolds numbers, we find that the networks are able to predict and control extreme events for a wide range of regimes. This means that the networks performance is robust with respect to changes in the system's physical parameter.

This paper provides a data-driven framework to predict extreme events in chaotic systems, which can be potentially applied to other nonlinear systems within and beyond fluid mechanics. Current and future work is focused on extending the framework to higher-dimensional turbulent flows with a data-compression strategy.

The code is available on the github repository [61].

ACKNOWLEDGMENTS

A.R. is supported by the EPSRC-DTP and the Cambridge Commonwealth, European & International Trust under a Cambridge European Scholarship. L.M. gratefully acknowledges financial support from the ERC Starting Grant PhyCo 949388 and the visiting fellowship at the Technical University of Munich, Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under Grant Agreement No. 291763. The authors are grateful to Yongyun Hwang for helpful discussions on the MFE model. The authors report no conflict of interest.

APPENDIX: TRAINING AND VALIDATION SETS

In this section, we report additional details regarding the training and validation of echo state networks. In all the cases analyzed, the washout interval consists of 0.5 LTs. In the single shot validation, the single validation interval consists of 2 LTs at the end of the first time series in the ten time series ensemble we use for training (Sec. II). In the recycle validation, in the $Re = 400$ case, we use 90 validation intervals, each of length 2 LTs. Subsequent intervals are generated by moving the validation interval forward in time by its own length. Therefore, 90 intervals span the first three time series of the ten time series ensemble we use for training (each time series lasts $\simeq 65$ LTs (see Sec. II) so that 30 intervals per time series span 60 LTs). In the $Re = [600, 800, 1000]$ cases, we use 88 validation intervals, each of length 2 LTs, spanning the first four time series of the ten time series ensemble we use for training. We do so because, due to smaller Lyapunov exponents, 4000 time units, which constitute one time series, consist of less than 60 LTs.

-
- [1] M. Farazmand and T. P. Sapsis, A variational approach to probing extreme events in turbulent dynamical systems, *Sci. Adv.* **3**, e1701533 (2017).
 - [2] A. J. Majda, M. Moore, and D. Qi, Statistical dynamical model to predict extreme events and anomalous features in shallow water waves with abrupt depth change, *Proc. Natl. Acad. Sci. USA* **116**, 3982 (2019).
 - [3] P. Yeung, X. Zhai, and K. R. Sreenivasan, Extreme events in computational turbulence, *Proc. Natl. Acad. Sci. USA* **112**, 12633 (2015).
 - [4] P. J. Blonigan, M. Farazmand, and T. P. Sapsis, Are extreme dissipation events predictable in turbulent fluid flows? *Phys. Rev. Fluids* **4**, 044606 (2019).
 - [5] T. C. Lieuwen and V. Yang, *Combustion Instabilities in Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modeling* (American Institute of Aeronautics and Astronautics, Reston, VA, 2005).
 - [6] S. Tibaldi and F. Molteni, On the operational predictability of blocking, *Tellus A* **42**, 343 (1990).
 - [7] M. Hassanaly and V. Raman, Classification and computation of extreme events in turbulent combustion, *Prog. Energy Combust. Sci.* **87**, 100955 (2021).
 - [8] M. Ghil, P. Yiou, S. Hallegatte, B. Malamud, P. Naveau, A. Soloviev, P. Friederichs, V. Keilis-Borok, D. Kondrashov, V. Kossobokov *et al.*, Extreme events: Dynamics, statistics and prediction, *Nonlinear Proc. Geophys.* **18**, 295 (2011).
 - [9] S. Rahmstorf and D. Coumou, Increase of extreme events in a warming world, *Proc. Natl. Acad. Sci. USA* **108**, 17905 (2011).
 - [10] M. Farazmand and T. P. Sapsis, Extreme events: Mechanisms and prediction, *Appl. Mech. Rev.* **71**, 050801 (2019),.
 - [11] T. P. Sapsis, Statistics of extreme events in fluid flows and waves, *Annu. Rev. Fluid Mech.* **53**, 85 (2021).
 - [12] V. Lucarini, D. Faranda, J. M. M. de Freitas, M. Holland, T. Kuna, M. Nicol, M. Todd, S. Vaienti *et al.*, *Extremes and Recurrence in Dynamical Systems* (John Wiley & Sons, 2016).
 - [13] S. S. Varadhan, Large deviations, in *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures* (World Scientific, 2010), pp. 622–639.

- [14] G. Margazoglou, L. Biferale, R. Grauer, K. Jansen, D. Mesterházy, T. Rosenow, and R. Tripiccone, Hybrid Monte Carlo algorithm for sampling rare events in space-time histories of stochastic fields, *Phys. Rev. E* **99**, 053303 (2019).
- [15] H. L. de S. Cavalcante, M. Oriá, D. Sornette, E. Ott, and D. J. Gauthier, Predictability and Suppression of Extreme Events in a Chaotic System, *Phys. Rev. Lett.* **111**, 198701 (2013).
- [16] E. N. Lorenz, Deterministic nonperiodic flow, *J. Atmos. Sci.* **20**, 130 (1963).
- [17] G. Boffetta, M. Cencini, M. Falcioni, and A. Vulpiani, Predictability: A way to characterize complexity, *Phys. Rep.* **356**, 367 (2002).
- [18] H. Babae and T. Sapsis, A minimization principle for the description of modes associated with finite-time instabilities, *Proc. R. Soc. A* **472**, 20150779 (2016).
- [19] M. Farazmand and T. P. Sapsis, Dynamical indicators for the prediction of bursting phenomena in high-dimensional systems, *Phys. Rev. E* **94**, 032212 (2016).
- [20] T. P. Sapsis, New perspectives for the prediction and statistical quantification of extreme events in high-dimensional dynamical systems, *Philos. Trans. R. Soc. A* **376**, 20170133 (2018).
- [21] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [22] A. 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 Proc. Geophys.* **27**, 373 (2020).
- [23] Z. 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).
- [24] K. Hasegawa, K. Fukami, T. Murata, and K. Fukagata, Machine-learning-based reduced-order modeling for unsteady flows around bluff bodies of various shapes, *Theor. Comput. Fluid Dyn.* **34**, 367 (2020).
- [25] P. A. Srinivasan, L. Guastoni, H. Azizpour, P. Schlatter, and R. Vinuesa, Predictions of turbulent shear flows using deep neural networks, *Phys. Rev. Fluids* **4**, 054603 (2019).
- [26] H. Eivazi, L. Guastoni, P. Schlatter, H. Azizpour, and R. Vinuesa, Recurrent neural networks and Koopman-based frameworks for temporal predictions in a low-order model of turbulence, *Int. J. Heat Fluid Flow* **90**, 108816 (2021).
- [27] G. Novati, L. Mahadevan, and P. Koumoutsakos, Controlled gliding and perching through deep-reinforcement-learning, *Phys. Rev. Fluids* **4**, 093902 (2019).
- [28] H. Jaeger and H. Haas, Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication, *Science* **304**, 78 (2004).
- [29] W. Maass, T. Natschläger, and H. Markram, Real-time computing without stable states: A new framework for neural computation based on perturbations, *Neural Comput.* **14**, 2531 (2002).
- [30] L. Grigoryeva and J.-P. Ortega, Echo state networks are universal, *Neural Netw.* **108**, 495 (2018).
- [31] F. Huhn and L. Magri, Learning ergodic averages in chaotic systems, in *Computational Science—ICCS 2020*, Lecture Notes in Computer Science (Springer International Publishing, Cham, 2020), pp. 124–132, Vol. 12142.
- [32] A. G. Hart, J. L. Hook, and J. H. Dawes, Echo state networks trained by Tikhonov least squares are $12 (\mu)$ approximators of ergodic dynamical systems, *Physica D* **421**, 132882 (2021).
- [33] S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**, 1735 (1997).
- [34] P. Vlachas, J. Pathak, B. Hunt, T. Sapsis, M. Girvan, E. Ott, and P. Koumoutsakos, Backpropagation algorithms and reservoir computing in recurrent neural networks for the forecasting of complex spatiotemporal dynamics, *Neural Netw.* **126**, 191 (2020).
- [35] M. Lukoševičius, A practical guide to applying echo state networks, in *Neural Networks: Tricks of the Trade*, Lecture Notes in Computer Science, Vol. 7700, edited by G. Montavon, G. B. Orr, and K. R. Müller (Springer, Berlin, Heidelberg, 2012), pp. 659–686.
- [36] A. Racca and L. Magri, Robust optimization and validation of echo state networks for learning chaotic dynamics, *Neural Netw.* **142**, 252 (2021).

- [37] J. Pathak, A. Wikner, R. Fussell, S. Chandra, B. R. Hunt, M. Girvan, and E. Ott, Hybrid forecasting of chaotic processes: Using machine learning in conjunction with a knowledge-based model, *Chaos* **28**, 041101 (2018).
- [38] N. Doan, W. Polifke, and L. Magri, Physics-informed echo state networks, *J. Comput. Sci.* **47**, 101237 (2020).
- [39] F. Huhn and L. Magri, Gradient-free optimization of chaotic acoustics with reservoir computing, *Phys. Rev. Fluids* **7**, 014402 (2022).
- [40] N. Doan, W. Polifke, and L. Magri, Short-and long-term predictions of chaotic flows and extreme events: physics-constrained reservoir computing approach, *Proc. R. Soc. A* **477**, 20210135 (2021).
- [41] A. Racca and L. Magri, Statistical prediction of extreme events from small datasets, in *Computational Science – ICCS 2022*, Lecture Notes in Computer Science, edited by D. Groen, C. de Mulatier, M. Paszynski, V. V. Krzhizhanovskaya, J. J. Dongarra, and P. M. A. Sloot (Springer, Cham, 2022), pp. 707–713, Vol. 13352.
- [42] V. Pyragas and K. Pyragas, Using reservoir computer to predict and prevent extreme events, *Phys. Lett. A* **384**, 126591 (2020).
- [43] M. Farazmand and T. P. Sapsis, Closed-loop adaptive control of extreme events in a turbulent flow, *Phys. Rev. E* **100**, 033110 (2019).
- [44] J. Moehlis, H. Faisst, and B. Eckhardt, A low-dimensional model for turbulent shear flows, *New J. Phys.* **6**, 56 (2004).
- [45] F. Waleffe, On a self-sustaining process in shear flows, *Phys. Fluids* **9**, 883 (1997).
- [46] J. D. Skufca, J. A. Yorke, and B. Eckhardt, Edge of Chaos in a Parallel Shear Flow, *Phys. Rev. Lett.* **96**, 174101 (2006).
- [47] B. Eckhardt, T. M. Schneider, B. Hof, and J. Westerweel, Turbulence transition in pipe flow, *Annu. Rev. Fluid Mech.* **39**, 447 (2007).
- [48] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevFluids.7.104402> for more details about the governing equations and the control strategies.
- [49] F. Huhn and L. Magri, Stability, sensitivity and optimisation of chaotic acoustic oscillations, *J. Fluid Mech.* **882**, A24 (2020).
- [50] F. Ginelli, P. Poggi, A. Turchi, H. Chaté, R. Livi, and A. Politi, Characterizing Dynamics with Covariant Lyapunov Vectors, *Phys. Rev. Lett.* **99**, 130601 (2007).
- [51] P. J. Werbos, Backpropagation through time: What it does and how to do it, *Proc. IEEE* **78**, 1550 (1990).
- [52] Z. Lu, J. Pathak, B. Hunt, M. Girvan, R. Brockett, and E. Ott, Reservoir observers: Model-free inference of unmeasured variables in chaotic systems, *Chaos* **27**, 041102 (2017).
- [53] A. N. Tikhonov, A. V. Goncharky, V. V. Stepanov, and A. G. Yagola, *Numerical Methods for the Solution of Ill-Posed Problems* (Springer Science & Business Media, 2013), Vol. 328.
- [54] M. D. Hoffman, E. Brochu, and N. de Freitas, Portfolio allocation for Bayesian optimization, <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.406.7117&rep=rep1&type=pdf> (2011).
- [55] M. Lukoševičius and A. Uselis, Efficient cross-validation of echo state networks, in *Artificial Neural Networks and Machine Learning – ICANN 2019: Workshop and Special Sessions*, Lecture Notes in Computer Science, edited by I. Tetko, V. Kůrková, P. Karpov, and F. Theis (Springer, Cham, 2019), pp. 121–133, Vol. 11731.
- [56] M. Stone, Cross-validators choice and assessment of statistical predictions, *J. R. Stat. Soc. B* **36**, 111 (1974).
- [57] H. He and E. A. Garcia, Learning from imbalanced data, *IEEE Trans. Knowl. Data Eng.* **21**, 1263 (2009).
- [58] J. Moehlis, H. Faisst, and B. Eckhardt, Periodic orbits and chaotic sets in a low-dimensional model for shear flows, *SIAM J. Appl. Dyn. Syst.* **4**, 352 (2005).
- [59] L. V. Kantorovich, Mathematical methods of organizing and planning production, *Management Science* **6**, 366 (1960).
- [60] P. H. Westfall, Kurtosis as peakedness, 1905–2014. *R.I.P.*, *Amer. Stat.* **68**, 191 (2014).
- [61] [github/MagriLab](https://github.com/MagriLab).