Machine learning prediction of tipping in complex dynamical systems

Shirin Panahi,[1] Ling-Wei Kong,[1] Mohammadamin Moradi,[1] Zheng-Meng Zhai,[1] Bryan Glaz,[2] Mulugeta Haile,[3] and Ying-Cheng Lai[1,4,*]

[1]*School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287, USA*
[2]*Army Research Directorate, DEVCOM Army Research Laboratory, 2800 Powder Mill Road, Adelphi, MD 20783-1138, USA*
[3]*Army Research Directorate, DEVCOM Army Research Laboratory, 6340 Rodman Road, Aberdeen Proving Ground, MD 21005-5069, USA*
[4]*Department of Physics, Arizona State University, Tempe, Arizona 85287, USA*

(Dated: November 1, 2024)

Anticipating a tipping point, a transition from one stable steady state to another, is a problem of broad relevance due to the ubiquity of the phenomenon in diverse fields. The steady-state nature of the dynamics about a tipping point makes its prediction significantly more challenging than predicting other types of critical transitions from oscillatory or chaotic dynamics. Exploiting the benefits of noise, we develop a general data-driven and machine-learning approach to predicting potential future tipping in nonautonomous dynamical systems and validate the framework using examples from different fields. As an application, we address the problem of predicting the potential collapse of the Atlantic Meridional Overturning Circulation (AMOC), possibly driven by climate-induced changes in the freshwater input to the North Atlantic. Our predictions based on synthetic and currently available empirical data place a potential collapse window spanning from 2040 to 2065, in consistency with the results in the current literature.

I. INTRODUCTION

A tipping point in nonlinear and complex dynamical systems is referred to as a transition from one stable steady state supporting the normal functioning of the system to another that can often be catastrophic and corresponds to system collapse [1]. This can happen as a system parameter passes through a critical point. For example, in ecosystems, before tipping the system is in a survival state with healthy species populations, while the state after the tipping is associated with extinction [1–8]. In the past decade, tipping point in ecosystems has been extensively studied [1–4, 6–26]. The phenomenon of tipping can also arise in other contexts such as epidemic outbreak [27], a sudden transition from normal to depressed mood in bipolar patients [28], alterations in the stability of the Amazon rain forest [29], an increase in the carbon emission from Boreal permafrost [30], and the melting of Arctic sea ice [31]. A likely scenario by which a tipping point can occur is when a parameter of the system varies with time - nonautonomous dynamical systems. Suppose the system is in a normal functioning state at the present. Due to the parameter change, at a certain time in the future a critical point will be crossed, leading to a catastrophic tipping. The global climate change is causing ecosystems and climate systems of different scales to become such nonautonomous dynamical systems with the increasing risk of tipping. Articulating effective methods to reliably anticipate tipping is an urgent problem with broad implications and applications.

In this paper, we develop a reservoir-computing framework tailored to anticipating tipping in nonautonomous dynamical systems and demonstrate its predictive power using examples from different fields. A particular application that provided the main motivation for our work is predicting the possible collapse of the Atlantic Meridional Overturning Circulation (AMOC) [32–34] that supports moderate and livable temperature conditions in Western Europe [35]. The AMOC transports warmer, upper waters in the Atlantic northward and returns colder, deeper waters southward [34]. Studies suggested that, since about 30 years ago, there has been a tendency for the AMOC to weaken [36, 37]. At the present, the AMOC is still in a "healthy" steady state that maintains a stable circulation of the pertinent ocean flows. A potential halt of the circulation would signify a collapse of the AMOC with dire consequences, which corresponds to another stable steady state of the underlying dynamical system. Such a collapse meets the criterion of a tipping point, i.e., a transition from one stable steady state to another.

It is worth emphasizing that the phenomenon of tipping in its original context [1–8] is characteristically distinct from the more commonly studied critical transitions from an oscillatory state to some final state. Examples of such transitions include a crisis through which a chaotic attractor is destroyed and replaced by a chaotic transient [38], the onset of synchronization from a desynchronization state [39], amplitude death [40], and the encountering with a periodical window [41]. While machine learning, in particular reservoir computing, has been applied to predicting such critical transitions [42–45], a shared characteristic among the existing works is the system's oscillatory behavior before the transition.

* Ying-Cheng.Lai@asu.edu

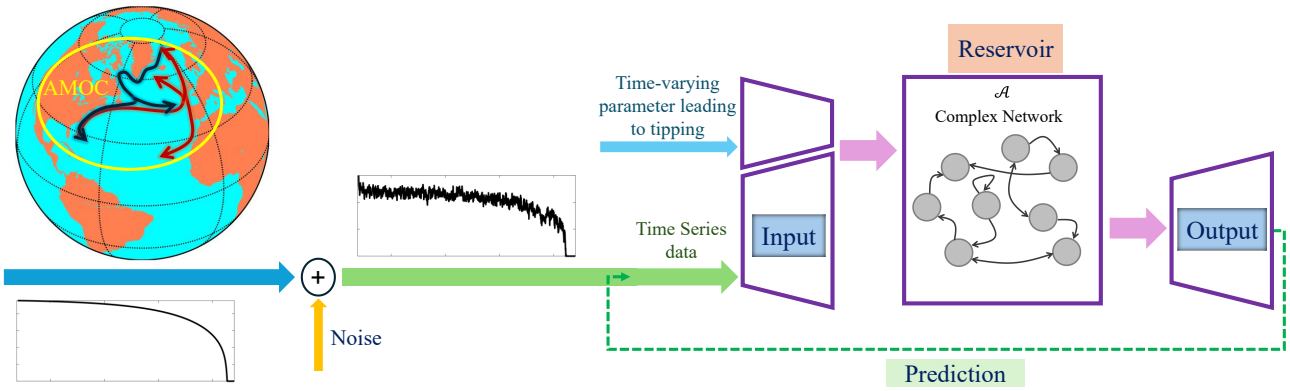

FIG. 1. Schematic illustration of the machine-learning framework for anticipating tipping in nonautonomous dynamical systems. The system begins in a stable steady state with no deterministic oscillations in the dynamical variables. Dynamic noise is leveraged to perturb the system, enabling the machine-learning model to detect changes and predict the tipping point even when the system is in a parameter regime prior to tipping.

This is advantageous because the time series for training the neural networks contain the temporal variations necessary for the machine to learn the dynamics of the system. Predicting a tipping point is significantly more challenging because, prior to the tipping, the system is in a stable steady state with no oscillations in the dynamical variables. (See Appendix A for a more detailed account of the notion of tipping in the literature.)

Our solution for machine-learning based prediction of tipping is exploiting dynamic noise. In particular, time series measured from real-world systems are noisy, and the inherent random oscillations are naturally suited for machine-learning training. In developing a machine-learning prediction framework, synthetic data are needed for validation. In this case, we generate time series with random perturbations about the deterministic steady state through stochastic dynamical modeling. While the presence of noise may potentially compromise the prediction accuracy, it serves a dual purpose by facilitating an adequate exploration of the phase space by the neural network dynamics, unveiling latent features that would otherwise remain obscured under noise-free conditions. A recent work has established that dynamical noise and/or measurement noise in the training dataset can be beneficial to the training process through a stochastic-resonance mechanism [46]. In addition, optimal calibration of noise levels can mitigate the risk of overfitting and promote generalization, allowing the reservoir computer to adapt to varying environmental conditions and data distributions. Incorporating a parameter channel into reservoir computing [42] to accommodate the time-varying parameter, we demonstrate that the reservoir computer can be trained to predict the occurrence of tipping in the future. To show the efficacy of our prediction framework, we present examples from climatic systems and ecological networks. For the problem of anticipating a potential collapse of the AMOC with synthetic and currently available empirical data, our machine-learning scheme places a collapse window spanning from 2040 to 2065, in consistence with the results in the current literature.

II. METHODS

A. Nonlinear dynamical mechanism of tipping

In nonlinear dynamics, a typical bifurcation leading to tipping is the forward or backward saddle-node bifurcation. Consider the situation of two coexisting stable steady states (or attractors): a normal "healthy" state and a catastrophic or "low" state, where each attractor has its own basin of attraction. As the bifurcation parameter increases with time, the healthy attractor can disappear through a backward saddle-node bifurcation, after which the low state is the only attractor in the phase space, signifying a tipping point. In the past, considerable efforts were devoted to anticipating tipping by identifying early warning indicators or signals [47–51]. A known phenomenon is enhanced fluctuations where, as the tipping point is approached, the variances of the measured values of the dynamical variables tend to increase. The reason is that, as the system moves toward a fold bifurcation, the dominant eigenvalue of the Jacobian matrix evaluated at the steady state approaches zero, making the landscape flatter and closer to a random walk about the steady-state attractor. Small noise will then generate large deviations of the trajectory from the attractor. In a recent work, a deep-learning based time-series classification scheme was introduced to determine if a tipping event is about to occur and the bifurcation [51].

B. Challenges with anticipating tipping

Oscillatory behaviors in the data in the pre-critical regime have the benefit of system trajectory's visiting a substantial portion of the phase space, thereby facilitating training by enabling the neural network to effectively learn the phase-space behavior or the dynamical climate of the target system. Differing from existing works on predicting critical transitions from an oscillatory dynamical state to a collapsed state, we aim to predict tipping from one stable steady state to another. In a noise-free situation, in the pre-tipping regime the system is in a stable steady state without oscillations in its dynamical variables. Introducing stochasticity or noise into the system leading to randomly oscillating dynamical variables provides a solution for neural-network training. We exploited dynamic noise in the data for training, where validation and hyperparameter optimization are performed based on data in the pre-critical regime. During the test or prediction phase, the reservoir computer operates as a closed-loop, deterministic dynamical system capable of predicting how the dynamical climate of the system changes with the time-varying bifurcation parameter. Since no data from the target system in the post-tipping regime were used for training (in a realistic situation, such data are not available), it is not possible for the reservoir computer to correctly predict the detailed system behavior after the tipping. However, the neural machine is capable of generating characteristic changes in the output variables at the tipping transition, making its anticipation possible. (See Appendix B for more details.)

C. Parameter-adaptable reservoir computing

We adopt parameter-adaptable reservoir computing [42] for anticipating tipping. A basic reservoir computer comprises three layers: an input layer, a hidden recurrent layer, and an output layer. Figure 2 illustrates the basic structure of parameter-adaptable reservoir computing that extends conventional reservoir computing by incorporating an additional parameter channel for the bifurcation parameter b. During the training, the input time series vector $\mathbf{u}(t)$ and the parameter b are concurrently projected onto the hidden layer through the time-series input matrix W_{in} and the parameter input matrix W_b, respectively. The hidden layer consists of N 1D dynamical neurons. Concatenating the dynamical states of all the neurons leads to an N-dimensional vector - the hidden state $\mathbf{r}(t)$ at each time step. The neural network in the hidden layer is recurrent with the connection matrix $W_{\mathbf{r}}$ and short-term memory. The output matrix W_{out} projects the hidden state $\mathbf{r}(t)$ to the output layer, generating the output vector $\mathbf{v}(t)$. The iteration equations of the parameter-adaptable reservoir computer are

$$\mathbf{r}(t) = (1-\alpha_r)\mathbf{r}(t-\Delta t) + \alpha_r \tanh\left[W_{\mathbf{r}}\mathbf{r}(t-\Delta t) + W_{\text{in}}\mathbf{u}(t) + W_b(k_b b + b_b)\right], \quad (1)$$

$$\mathbf{v}(t) = W_{\text{out}}\mathbf{f}(\mathbf{r}(t)), \quad (2)$$

where $\alpha_r \in (0,1]$ is the leakage parameter defining a temporal scale of the reservoir network, Δt is the time step, $\tanh(\cdot)$ is the hyperbolic tangent function serving as the nonlinear activation function in the hidden layer, k_b and b_b are the gain and bias of the parameter b, respectively, and $\mathbf{f}(\cdot)$ is a nonlinear output function of the reservoir computer.

A feature of the reservoir network is the random generation and the subsequent fixation of the input matrices W_{in} and W_b, along with the recurrent network matrix $W_{\mathbf{r}}$. These matrices remain fixed during training with only the output matrix W_{out} undergoing optimization. This design choice eliminates the need for back propagation in time during the training, alleviating computational cost and mitigating potential difficulties such as vanishing and exploding gradients. Following the random generation of the three matrices, the training process begins by inputting the time series $\mathbf{u}(t)$ and the corresponding control parameter b through the input layer. The dynamical evolution of the neural network follows

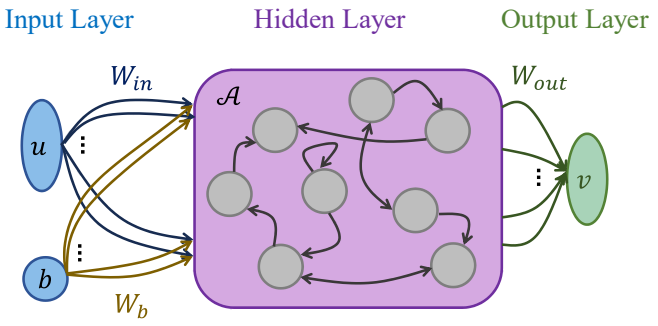

FIG. 2. Illustration of parameter-adaptable reservoir computing.

Eq. (1). This process is also referred to as the "listening phase" or "echoing phase", as if the driving training signals are echoing in the hidden state. During the training, Eq. (2) is not invoked as the output matrix has not been trained yet. Multiple trials of the time series

data from the target system, each associated with a distinct parameter value, are presented as the training data. Upon completion of the echoing phase for a specific trial for a particular b value, the parameter-adaptable reservoir computer is re-initialized for a new echoing phase for another training parameter value. The hidden state behaviors observed during the training are recorded, whose variations are implicitly linked to the corresponding parameter value b since it affects the dynamical evolution of the state in the hidden layer.

Let the length of each trial of the training time series be T_{train} (in the unit of the number of steps) and the number of trials of training with different parameter values be n_b ($n_b = 4$ in our work). The "echoing" results $\mathbf{r}(t)$ are concatenated into a matrix R of dimensions $N \times n_b T_{\text{train}}$. Applying the nonlinear function $\mathbf{f}(\cdot)$, we obtain the transformed matrix $R' = \mathbf{f}(R)$ that captures the echoing hidden state for subsequent linear regression. A training target is essential. We focus on reservoir networks whose output represents one-step prediction, where $\mathbf{v}(t)$ is equal to $\mathbf{u}(t + \Delta)$, making the training target the stacking of all training time series with a one-step difference from the input data, denoted as V. Finally, a ridge regression is conducted between R' and V to determine the output matrix:

$$W_{\text{out}} = V \cdot R'^T (R' \cdot R'^T + \beta_r I)^{-1}, \tag{3}$$

where β_r is the coefficient of $L - 2$ regularization.

An alternative training approach involves supplying a time series with a non-constant parameter value, a nonstationary time series with the corresponding time-varying parameter $b(t)$ as the training data [42]. This configuration is more practical in various scenarios and is employed in our work.

Having successfully trained a parameter-adaptable reservoir computer, we can now use it to make predictions for a specific parameter value b of interest, which is in a parameter regime different from that for training. The reservoir computer autonomously extrapolates the learned dynamics during training, generating predictions of the system dynamics at some unobserved parameter value. The iteration equation during prediction is given by

$$\mathbf{r}(t) = (1 - \alpha_r)\mathbf{r}(t - \Delta t) + \alpha_r \tanh(W_{\text{r}} \cdot \mathbf{r}(t - \Delta t) + W_{\text{in}}\mathbf{v}(t - \Delta t) + W_b(k_b b + b_b)), \tag{4}$$

$$\mathbf{v}(t) = W_{\text{out}}\mathbf{f}(\mathbf{r}(t)), \tag{5}$$

where $\mathbf{u}(t)$ in Eq. (1) is replaced by $\mathbf{v}(t - \Delta t)$. Given the recurrent structure of reservoir computing, it is necessary to properly initialize the hidden states in order to make any predictions. As one may observe from Eq. (1), a previous state $\mathbf{r}(t-\Delta t)$ is needed to calculate $\mathbf{r}(t)$. For short-term validation, we initialize the prediction by replicating the final hidden state obtained during the training as a one-step previous state. This allows for a direct comparison of the validation result with the actual time series, facilitating the calculation of an error metric, such as the root-mean-square error (RMSE), as the validation error. To conduct long-term testing on the "climate," where a more diverse ensemble of predictions is required to capture the behaviors of the target system, we introduce additional randomness. Specifically, we utilize a randomly selected short segment from the actual time series data to "warm up" the reservoir hidden state. To prevent the reservoir computer from becoming stuck in a single attractor, especially in the presence of multistability in the dynamics, we introduce additive observational noise to the "warm-up" data.

III. RESULTS

A. Anticipating a potential collapse of the AMOC

Our machine-learning framework is designed to tackle the challenge of anticipating tipping in nonautonomous dynamical systems in general, as shown in Fig. 1 where, prior to tipping, the system is in a stable steady state with no deterministic oscillations in the dynamical variables. We use four types of data: (1) synthetic data from one-dimensional (1D) AMOC fingerprint model, (2) synthetic data from a 2D conceptual model of AMOC, (3) synthetic data from the Community Earth System Model, and (4) empirical AMOC fingerprint data. As will be demonstrated, the predicted time window of a potential future AMOC collapse from the four types of data are consistent with each other.

1. Anticipating AMOC collapse from 1D synthetic fingerprint data

Due to the difficulty of continuously monitoring the AMOC and the limited availability of long-term observational data, analyzing certain fingerprints of the AMOC provides a viable method to gain insights [53]. For example, sea surface temperature (SST) has been employed as

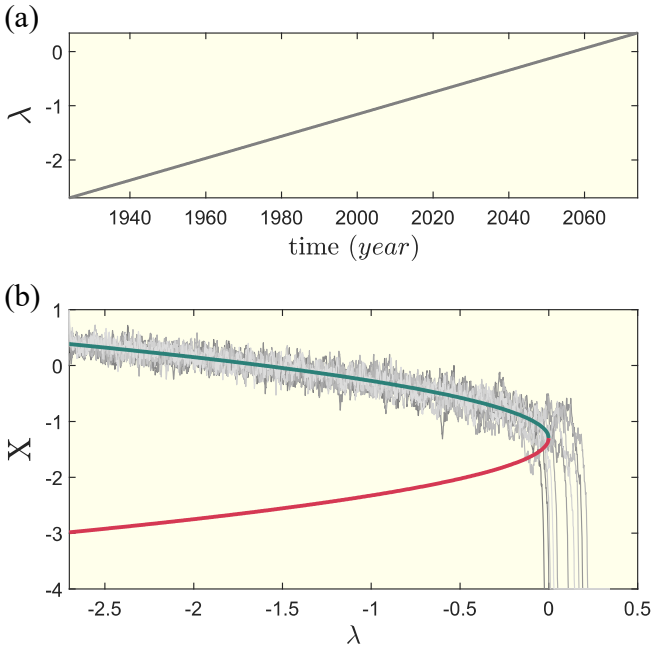

FIG. 3. Random realizations of a tipping point transition in the 1D stochastic AMOC fingerprint model (6). (a) Time-varying bifurcation parameter λ that increases exponentially with time, while other parameters are the best-estimated values extracted from the empirical fingerprint data [52]: $A = 0.95$, $m = -1.3$, $\lambda_0 = -2.7$, $\sigma = 0.3$, $t_0 = 1924$, and $\lambda_c = 0$. (b) Ten random model realizations, with the dashed green and red curves indicating the stable and unstable equilibria. In the underlying deterministic system, a backward saddle-node bifurcation and hence a tipping point occurs at $\lambda_c = 0$. In the presence of stochastic driving, the value of λ at which the system collapses, characterized by the dynamical variable X's approaching a large negative value, varies among the realizations, but they are near $\lambda = 0$ on the positive side.

a promising proxy for assessing the AMOC strength [54–57]. Quite recently, a 1D stochastic SST model [52] with parameter values estimated from the real data was constructed to understand the tipping dynamics of the AMOC. It was suggested that the AMOC may be approaching a potential collapse through a tipping point, which can occur as early as 2025. The model is described by the following stochastic nonlinear differential equation with a generic bifurcation parameter λ:

$$\dot{X}_t = -[A(X_t - m)^2 + \lambda)] + \sigma dB_t, \tag{6}$$

where X_t is a stochastic dynamical variable exhibiting a tipping transition, A is a time scale parameter, m is defined as $\mu - \sqrt{|\lambda|/A}$ with μ representing the stable fixed point of the process, B_t is a Brownian motion, and σ is the noise amplitude. Initially, the system is in a statistically stable state with constant $\lambda = \lambda_0$. At time t_0, λ begins to increase toward the critical point λ_c. As λ increases, the dynamical variable X_t exhibits fluctuations but its mean value decreases continuously. Despite the fluctuations, X_t eventually collapses to a large negative value, signifying the collapse of the AMOC. For fixed model parameters at the most likely estimated values from the AMOC fingerprint data ($A = 0.95$, $m = -1.3$, $\lambda_0 = -2.7$, and $\sigma = 0.3$, and $t_0 =$ Year 1924), the underlying deterministic system exhibits a backward saddle-node bifurcation, corresponding to the coalescing point of the stable and unstable equilibrium points, as shown in Fig. 3, where a tipping point occurs at $\lambda_c = 0$ (see Appendix C for details).

Climate change is a driving force to slow down and eventually halt the AMOC, making the underlying dynamical system nonautonomous. The nonautonomous version of Eq. (6) can be obtained by making the bifurcation parameter λ time-dependent. In particular, the impact of climate change was modeled [52] by an exponential increase in $\lambda(t)$ with time from some initial value $\lambda_0 < 0$, as illustrated in Fig. 4(a). As λ increases towards the bifurcation point $\lambda_c = 0$, the system approaches a tipping point at the time T_c. In the time interval, $[0, T_c]$, the AMOC variable $X(t)$ fluctuates about the stable equilibrium. After the tipping at T_c, $X(t)$ rapidly decreases to a large negative value, signifying the AMOC collapse. The value of the tipping time T_c varies across different realizations.

We now demonstrate that a trained reservoir computer is able to predict the tipping time T_c. For each realization, we divide the data into two distinct segments: training and testing, as highlighted in Figs. 4(a) and 4(b) in purple and blue, respectively, where the end of the purple data segment marks the present time (in year). Note that, up to the present time, the AMOC has been stable, where the dynamical variable $X(t)$ fluctuates about the healthy stable steady state. If there was no noise, $X(t)$ would be a smooth and a slowly decreasing function of time, as exemplified in Fig. 3, and it is not possible to train the reservoir computer with the non-oscillatory time series. What makes training possible is noise rendering oscillatory and random the time series $X(t)$. Figure 4(c) presents one prediction run, where the blue trace is the testing data in the time interval between now and year 2065 (the ground truth), and the dashed red trace is the reservoir-computing prediction. For this particular realization, the predicted AMOC collapse time is between the years 2062 and 2063. At about the same time, the predicted $X(t)$ exhibits an abnormal behavior that is drastically and characteristically different from that prior to the tipping, indicating a successful prediction of the tipping point. Note that, since the reservoir computer has never "seen" the blue testing data segment that includes the collapse of $X(t)$ to some negative value, it is not possible for the machine to predict the value of $X(t)$ after tipping. Nevertheless, the predicted abnormal behavior is indicative of some critical behavior in the system. Figure 4(d) shows a histogram of the predicted values of T_c from 1000 reservoir-computing realizations. For the 1D AMOC fingerprint model, the parameter-adaptable reservoir computer predicts that a collapse of the AMOC is likely to occur between the years 2055 and

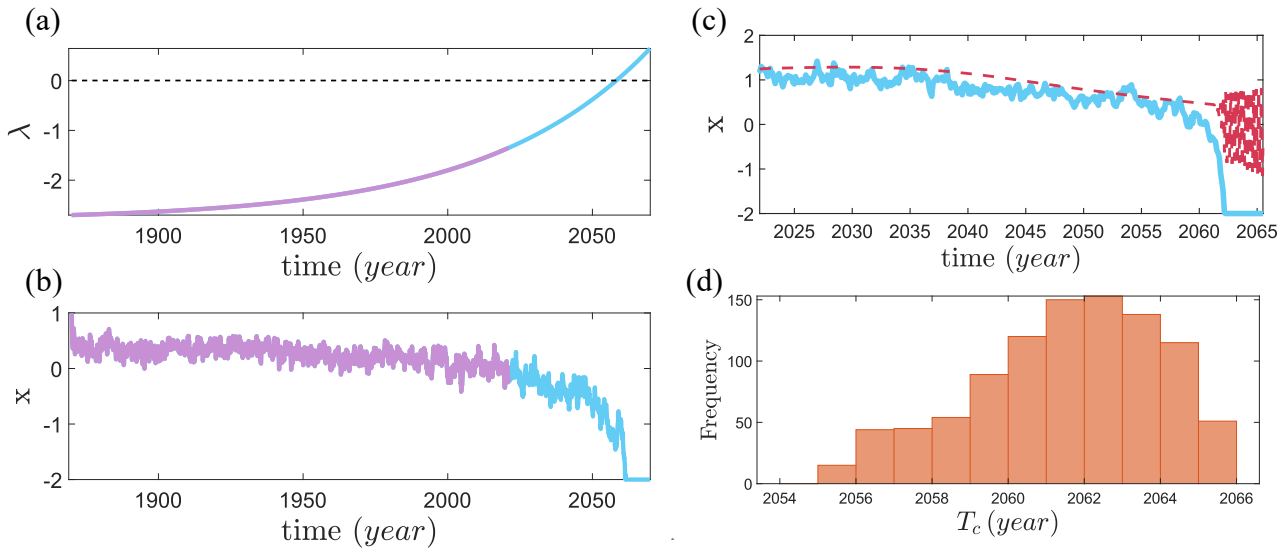

FIG. 4. Reservoir-computing prediction of the time window of AMOC collapse from the 1D time-dependent fingerprint model. (a) The exponential growth with time of the bifurcation parameter $\lambda(t)$, which starts from the value $\lambda_0 = -2.7$ in the year 1870. The horizontal dashed line indicates the tipping point $\lambda_c = 0$. (b) A realization of the time series $X(t)$, where the purple (blue) segment represents the training and testing (validation and prediction) data, respectively. (c) The testing data (blue) and reservoir-computing prediction (red) in the time window from year 2022 to year 2065. For this particular realization, the AMOC variable $X(t)$ collapses between the years 2062 and 2063 (blue, real data). The reservoir computer predicts an abnormal behavior in $X(t)$ at about the same critical time T_c, signifying a tipping point. (d) Histogram of the predicted AMOC collapse time T_c obtained from 1000 machine realizations. Tipping is likely to occur between year 2055 and year 2066.

2066, which is consistent with the result in Ref. [52].

The histogram of the collapse time T_c in Fig. 4(d) was obtained from 1000 machine realizations, but the training and testing data are from one specific realization of the 1D AMOC fingerprint model. For different model realizations, the tipping time T_c is different, so are the predictions. Table I lists the prediction results from 10 model realizations. It can be seen that in all cases, the predicted mean value of the collapse year is close to that of the original data, providing further validation of our reservoir-computing prediction scheme.

TABLE I. Predicted tipping time from 20 synthetic datasets

Dataset	Model T_c	T_c from 1000 machine realizations	
		Mean	Std (years)
1	2061	2062	4
2	2054	2057	6
3	2056	2057	3
4	2070	2069	6
5	2064	2062	5
6	2059	2060	4
7	2058	2060	5
8	2062	2061	4
9	2060	2062	6
10	2065	2066	4

2. *Predicting AMOC collapse based on synthetic data from a 2D conceptual model*

A recent study [33] addressed the phenomenon of tipping within climate systems, providing insights into how time-varying parameters can lead to abrupt and potentially catastrophic transitions in the climate. The mechanism of tipping was illustrated using a 2D conceptual model [33] with two state variables, denoted as x and y. The dynamics produced by this model closely resemble those observed in the ocean model. The 2D model is described by the following equations:

$$\frac{dx}{dt} = (-r^4 + 2r^2 - \beta)x - \omega\hat{y}, \tag{7}$$

$$\frac{d\hat{y}}{dt} = (-r^4 + 2r^2 - \beta)\hat{y} + \omega x, \tag{8}$$

where $r^2 = x^2 + \hat{y}^2$, and $\hat{y}$ is defined as $y - \gamma\beta$. The parameter β is a bifurcation parameter describing the freshwater forcing parameter, while ω represents the frequency. The dependence of y on β is parameterized by γ. The system possesses one stable and one unstable limit cycle for $\beta < 1$. For $\beta = 1$, the limit cycles merge and disappear in a saddle-node bifurcation [33].

To make the dynamical system (7) nonautonomous, we assume that the bifurcation parameter β is time-dependent: starting from an initial value β_0, it linearly increases towards a critical value β_c. Dynamical noise is introduced into the system by additive stochastic terms

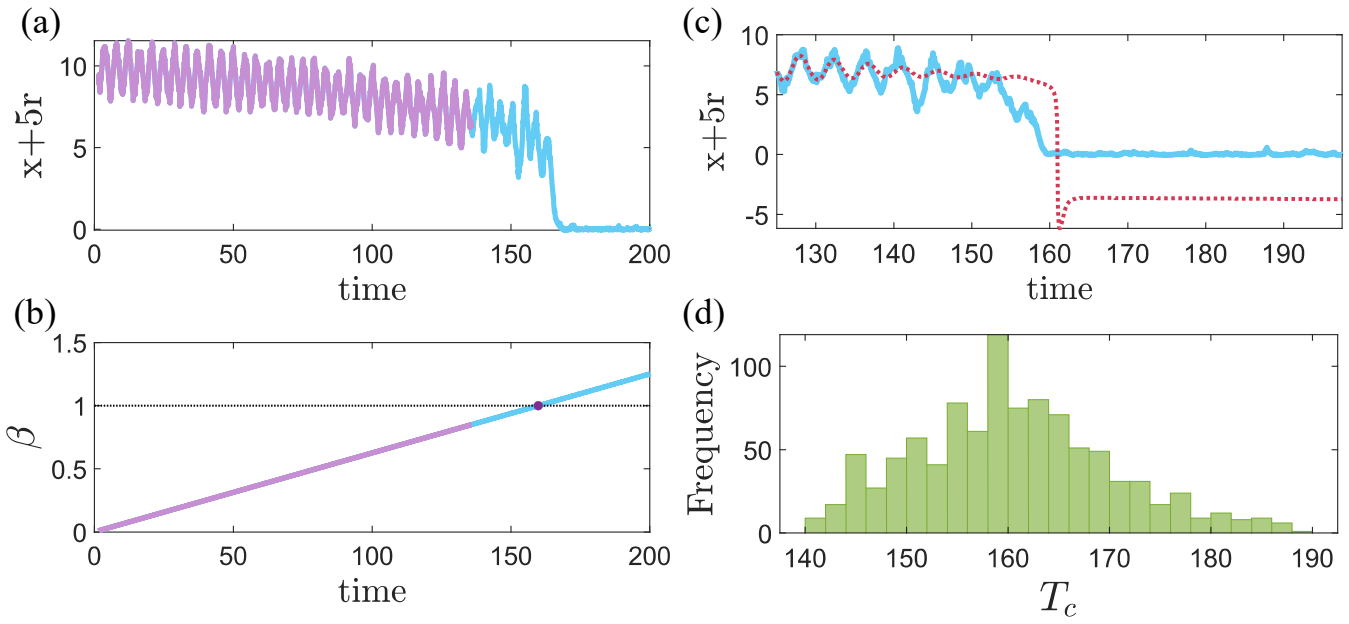

FIG. 5. Reservoir-computing prediction of the time window of AMOC collapse from the 2D time-dependent conceptual AMOC model. (a) One realization of 2D conceptual AMOC model for $\gamma = 3$ and $\sigma = 0.1$. (b) Time-varying freshwater forcing parameter β. (c) An example of testing data (solid blue trace) and reservoir-computing prediction (dash-dotted red trace). (d) Histogram of the predicted critical point from 1000 random reservoir realizations.

in (7) of independent Gaussian random processes of zero mean and amplitude σ. Figures 5(a) and 5(b) show a single realization of the evolution of the conceptual AMOC fingerprint $(x + 5r)$ and the β parameter over time, respectively. As β increases, a tipping point occurs at $\lambda_c = 0$, where the system undergoes a sudden shift in the attractor, transitioning from random oscillations about the stable limit cycle to a stable equilibrium. The specific tipping time varies among the independent stochastic realizations.

To anticipate a possible tipping point, we partition the data into two sets: training data (highlighted in purple) and testing data (in blue). The training data consists of a portion of the time series of $x + 5r$ as input and β as a control parameter associated with oscillations about the stable limit cycle. During the testing phase, the trained reservoir computer is employed alongside the remaining control parameter data to predict the tipping. Figure 5(c) shows an illustrative example, where the real testing data are in blue and the corresponding predicted data are represented by red. Training data are collected for parameter values $\beta \in [0.01, 0.79]$. To ensure the prediction efficacy, we repeat the whole process for 1000 random realizations of the reservoir computer. Figure 5(d) presents a histogram of the anticipated tipping point values. In all the realizations, a tipping point is anticipated to occur in the future within the time interval $T_c \in [140, 190]$, which contains the ground truth value $T_c = 159$ from direct simulation of the 2D stochastic system.

3. *Predicting AMOC collapse using the synthetic data from the Community Earth System Model*

In a quite recent study of the Community Earth System Model (CESM) [58], an AMOC tipping event with significant climate consequences was revealed. CESM is a coupled climate model for simulating various components of the Earth's climate system simultaneously, making it possible to explore the dynamics under the past, present, and future climate conditions. An analysis of the output data of CESM revealed a tipping point as characterized by the minimum of the AMOC-induced freshwater transport at the Southern boundary of the Atlantic [58].

We use the simulated AMOC strength data from Ref. [58] to test our reservoir-computing based framework for predicting tipping. In the CESM model, the freshwater flux forcing (F_H) linearly increases at the rate 3×10^{-4} Sv year-1 until the model year 2200, where a maximum of $F_H = 0.66$ Sv is reached, as shown in Fig. 6(a). The AMOC strength, defined as the total Meridional volume transport, is shown in Fig. 6(b), where the vertical dashed line indicates a tipping point $T_c = 1758$. The purple segment is used to train the reservoir computer and the blue segment is the testing data. The reservoir-computing output is shown in red in Fig. 6(c), where an abnormal behavior in $x(t)$ occurs at about the same critical time T_c as the model tipping time. To characterize the prediction performance, we repeat the process using 1000 machine-learning realizations. The resulting histogram of the predicted AMOC collapse time T_c is shown in Fig. 6(d), which indicates that tipping is likely

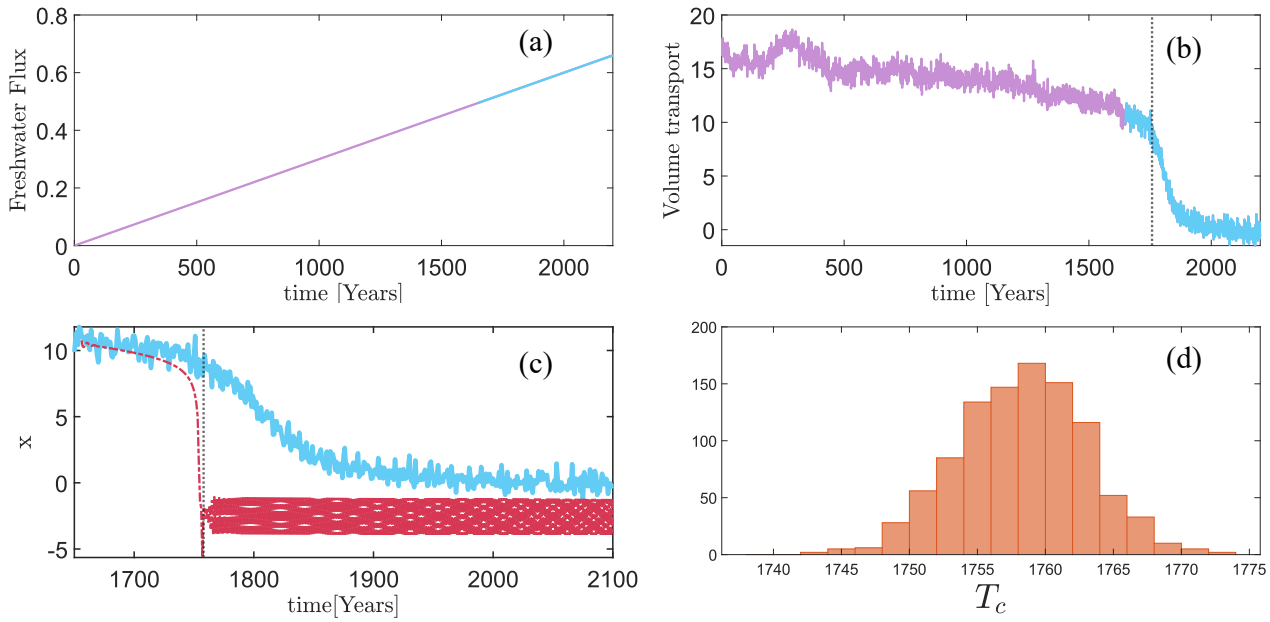

FIG. 6. Reservoir-computing prediction of the time window of potential AMOC collapse from the CESM synthetic data. (a) Linear time-varying freshwater flux. (b) AMOC strength, where the purple and blue segments represent the training and testing (validation and prediction) data, respectively. The horizontal dashed line indicates the tipping point $T_c = 1758$. (c) Testing data (blue) and reservoir-computing prediction (red). AMOC strength ($x(t)$) collapses at model year 1758 (blue, real data). The reservoir computer predicts an abnormal behavior in $x(t)$ at about the same critical time T_c, signifying a tipping point. (d) Histogram of the predicted AMOC collapse time T_c obtained from 1000 reservoir network realizations. Tipping is likely to occur between model years 1740 and 1775.

to occur between the model years 1740 and 1775.

4. *Predicting AMOC collapse using empirical fingerprint data*

It is necessary to conduct tests using empirical AMOC data. We use AMOC fingerprint sea-surface temperature (SST) datasets with the same exponential growth of the bifurcation parameter [52], as shown in Fig. 7(a). Figure 7(b) shows a segment of the SST data up to the present time (in purple color), which is used for training, and a typical realization of the reservoir-computing predicted time series (red). Prior to reaching the critical point $\lambda_c = 0$, the predicted AMOC fingerprint exhibits a smooth behavior that is essentially a continuation of the training data, indicating no collapse. About $\lambda_c = 0$, the machine-learning prediction becomes highly irregular, signifying a collapse. Figure 7(c) shows a histogram of the predicted critical time of AMOC collapse from 1000 reservoir-computing realizations. The range of possible collapse time is from year 2040 to year 2066, with the median around year 2053. This result is consistent with those in Fig. 4(d) and in Ref. [52].

To further demonstrate the generality and power of our parameter-adaptable reservoir computing framework for predicting tipping in complex and nonautonomous dynamical systems, we tested the following datasets from alternative AMOC models and ecological networks that exhibit a tipping point in the conventional sense of coexisting stable fixed-point attractors: (1) mutualistic pollinator-plant networks, (2) a plant-herbivore model, and (3) a climate model. For models (1) and (2), a bifurcation parameter is assumed to vary continuously with time. For model (3), the observed time series are collected from a sequence of different constants or nearly constant parameter values. We have also tested a three-box AMOC model, with detailed results presented in Appendix D.

B. Anticipating tipping in pollinator-plant mutualistic networks

We demonstrate the capability of our parameter-adaptable reservoir-computing approach to anticipate tipping points in real mutualistic networks. Specifically, we study two real-world pollinator-plant networks (Web of Life database): Network A from Flores, Acores [59] with 10 plant species and 12 pollinator species; Network B from an empirical study in Hestehaven, Denmark [60], which has 8 plant species and 42 pollinator species. The dynamical variables of these systems are the abundances of the plant and pollinator species. We take κ, the average pollinator species decay rate, as the bifurcation parameter. Due to environmental changes, κ varies slowly

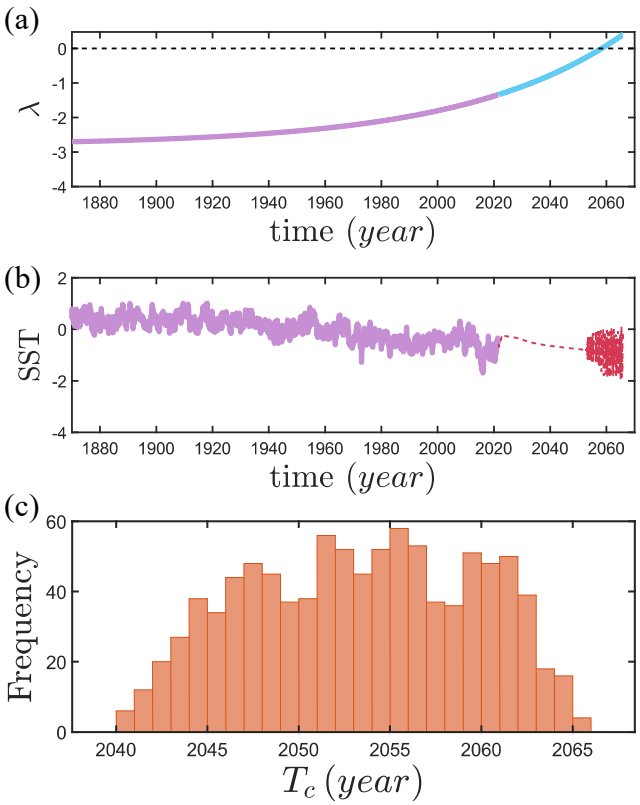

FIG. 7. Reservoir-computing prediction of AMOC collapse time using empirical fingerprint data. (a) An exponential growth of the bifurcation parameter λ (from Ref. [52]). (b) Available AMOC fingerprint SST data from year 1875 to the present year (purple). This data segment is noisy and employed in training the reservoir computer. The red trace is one example of the predicted SST behavior, which is smooth until the critical value $\lambda_c = 0$ for collapse is reached. (c) A histogram of the predicted AMOC collapse time from 1000 reservoir-computing realizations. The time range of potential AMOC collapse is between the year 2040 and the year 2066.

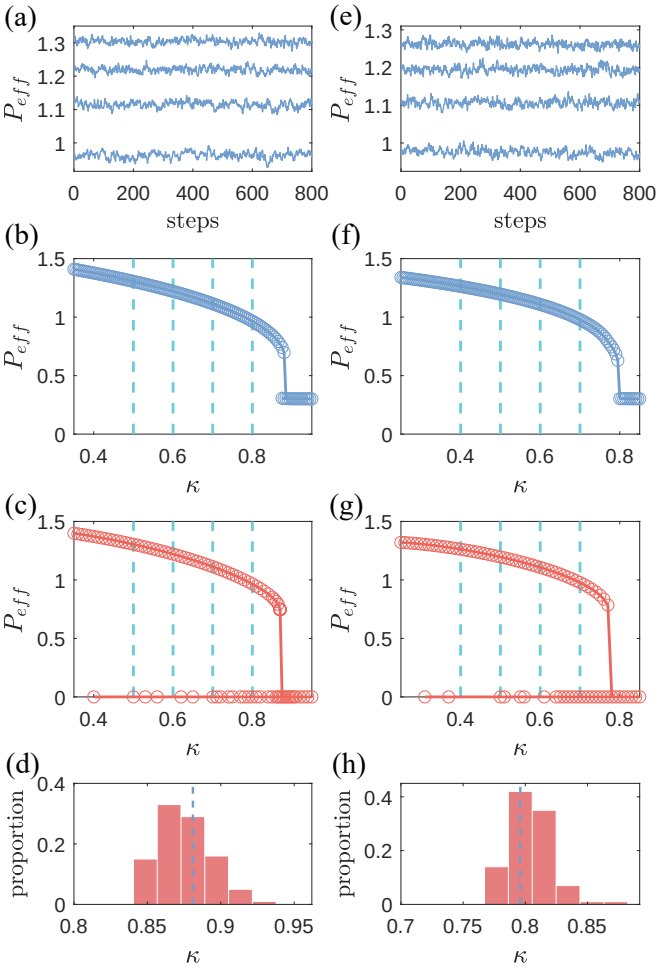

FIG. 8. Anticipating tipping with noisy data in two real plant-pollinator networks. (a) Noisy training data from mutualistic network A. The four trials of data correspond to four different training values of the control parameter κ. (b) Tipping in network A: at the tipping point $\kappa_c \approx 0.881$, the effective abundance of the plant species P_{eff} goes through an abrupt regime shift to a lower level through a saddle-node bifurcation. (c) Reservoir-computing predicted tipping. The vertical dashed cyan lines correspond to the values of κ in the training data set, all in a safe regime. Occasional predicted negative abundance values are counted as zero. (d) Histogram of the anticipated tipping point from an ensemble of 100 prediction runs, where the vertical blue dashed line denotes the true value κ_c. (e) Noisy training data from network B. (f) Tipping in network B at $\kappa_c = 0.796$. (g) Reservoir-computing predicted tipping, where the vertical dashed cyan lines correspond to the values of κ in the training data set. (h) Histogram of the anticipated tipping point from an ensemble of 100 prediction runs, with the ground truth κ_c indicated by the vertical blue dashed line.

with time. Tipping occurs in both systems, as shown in Figs. 8(b) and 8(f), where the equilibria of the effective plant abundance P_{eff} averaged over all the plant species for different κ values are plotted. The tipping point for network A (B) is $\kappa_c = 0.881$ ($\kappa_c = 0.796$), after which P_{eff} decreases abruptly.

The training data are time series of the abundance of the species in the mutualistic network. The reservoir computer is trained based on noisy time series from a number of distinct values of the bifurcation parameter in the pre-tipping ("safe") regime, where the networked system is under correlated, demographic noise. Specifically, for network A (B), the parameter values are $\kappa = 0.5, 0.6, 0.7$, and 0.8 ($\kappa = 0.4, 0.5, 0.6$, and 0.7). As the mutualistic networks are high-dimensional dynamical systems, training using the time series from all species is computationally costly. We employ a previously developed dimension-reduction method for mutualistic networks [18] and take the effective plant and pollinator abundances, denoted as P_{eff} and A_{eff}, respectively, as the training data. Figures 8(a) and 8(e) show the P_{eff} components of the training data, which are noisy around the equilibria.

In the prediction phase, we extend the bifurcation parameter into an untrained nearby regime. Figures 8(c)

and 8(g) show the prediction results for networks A and B, respectively, where the training parameter values are marked by the vertical cyan dashed lines. The predicted tipping points for both networks are near the true values. Note that, during prediction the reservoir computer is a deterministic dynamical system designed to capture the deterministic components of the dynamics of the target system. As a result, the machine-generated trajectories correspond to smooth, damped oscillations converging towards the fixed points, which deviate significantly from the noisy training trajectories. The final equilibrium point is extracted as the predicted values of the fixed point, as shown by the red circles in Figs. 8(c) and 8(g). Given the stochastic nature of the training datasets, we conduct 100 random testing trials for each mutualistic network, where both the training data and the relevant matrices underlying reservoir computing differ across the testing trials. Figures8(d) and 8(h) show the histograms of the anticipated tipping point values. Tipping is consistently anticipated in all the trials. For network A, the predicted tipping points occur within the interval $\kappa \in [0.8405, 0.9215]$ about the true value $\kappa_c = 0.881$. Out of 100 trials, 66 anticipate the tipping point within the narrower interval $\kappa \in [0.8608, 0.9012]$. For network B, 98 out of 100 trials predict the tipping point to fall within the interval $\kappa \in [0.748, 0.844]$ about the true value $\kappa_c = 0.796$. In addition, 66 out of 100 trials foresee the tipping point within the narrower interval $\kappa \in [0.772, 0.820]$.

We address two pertinent issues. First, to remove any doubt that the observed collapses in the predicting phase might be artifacts stemming from the instability of the reservoir computer when operating in an untrained parameter region, we obtain extrapolation results on the opposite side of the untrained parameter region, i.e., with κ smaller than the training values. For networks A and B, the parameter regions tested are $\kappa \in (0.35, 0.50)$ and $\kappa \in (0.25, 0.50)$, respectively. No collapse is observed on this side of the parameter region for any trial, as the reservoir computer predicts stable species abundances for all trials. This justifies that the tipping behaviors predicted on the other side of the parameter region are not artifacts. Second, while we have shown that the tipping of the target attractor can be anticipated from a forecasting approach, it is difficult to predict the actual dynamics after the tipping. For the ground truth results shown in Fig. 8, there is an abrupt decrease in the abundance of the plant species after tipping but without an immediate total extinction. The reservoir computer is not able to predict such behaviors correctly. In general, the target system can be in a new regime after tipping, which can be far away from the training region in the phase space. Expanding the dynamics from the training region to such a distant new regime can be difficult and unreliable given the nonlinearity in the target system. For a similar reason, the parameter region of multistability is also not accurately predicted, since the dynamical features of the lower state have never been "seen" by the reservoir computer.

C. Anticipating tipping in a plant-herbivore system

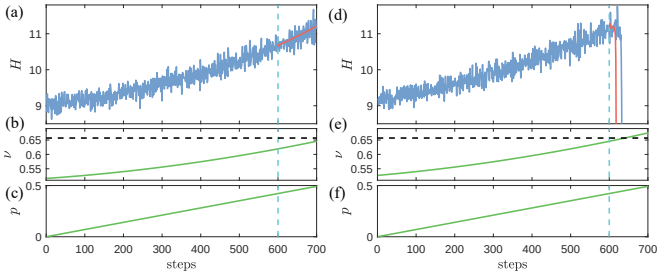

FIG. 9. Anticipating tipping point with an unknown continuously changing environmental factor in a plant-herbivore system. (a) Simulated herbivore abundance in a parameter region away from tipping. Training and testing data from the target system are shown as the noisy blue curve, separated by the vertical dashed cyan line. The red segment to the right of the vertical dashed cyan line is the reservoir-computing predicted abundance, which on average agrees with the blue data. (b) Time variation of the environmental factor $\nu(t)$ that has not yet reached the tipping point value (marked by the horizontal dashed black line) in the testing interval. (c) The surrogate control parameter $p(t)$, which is assumed to increase linearly with time. (d) Simulated time series with tipping in the testing interval, which is successfully anticipated by reservoir computing (the red curve). (e) Parameter $\nu(t)$ reaching the tipping point $\nu_c = 0.657$ at the crossing of dashed black lines, leading to the extinction behavior in panel (d). (f) The surrogate control parameter $p(t)$.

We study a plant-herbivore system [26, 61] in a tipping parameter regime with a saddle-node bifurcation. The dynamical variables are the biomass densities of the plant P and the herbivore H. Two system parameters, the plant growth rate and herbivore mortality rate, are time-varying [26] according to a common environmental factor ν. Increasing ν across a tipping point at $\nu_c = 0.657$ results in a sudden decrease of the herbivore biomass H and its subsequent extinction, dynamically induced by a saddle-node bifurcation at ν_c. We assume that the environmental factor $\nu(t)$ increases quadratically as a function of the time. In a realistic situation, the functional form of $\nu(t)$ is unknown. For simplicity, we assume that $\nu(t)$ increases linearly with time [44].

We generate two different datasets (including training and testing) for comparison. For the first set, $\nu(t)$ is below the tipping point in both the training and testing data. For the second dataset, $\nu(t)$ is below but close to the tipping point ν_c during training, but it crosses the tipping point in the testing data, leading to a sudden extinction of the herbivore. Demographic noise is applied to the system to generate noisy training data. In both cases, the surrogate parameter $p(t)$ has the same rate of linear increase with time. Figures 9(a-c) show the testing results for the first dataset, while Figs. 9(d-f)

are for the second dataset. For the pre-tipping dataset, the reservoir computer predicts the correct dynamical behavior of no tipping, as shown in Fig. 9(a). Likewise, for the second dataset, the tipping behavior is correctly predicted, as shown in Fig. 9(d). We repeat the process of training and predicting on the two settings of $\nu(t)$ 1,000 times with different random seeds for dynamical noise and reservoir-computing matrices. For the pre-tipping case, within the same testing time interval as in Fig. 9(a), 126 trials (12.6%) provide false positive predictions, most of which are near the end of the testing interval where the bifurcation parameter is closer to the tipping point. The remaining 874 trials (87.4%) are all true negatives. For the tipping case, only 72 trials (7.2%) fail to anticipate the collapse within twice the average actual time of collapse.

D. Anticipating tipping in a climate model with discrete control parameter scheme

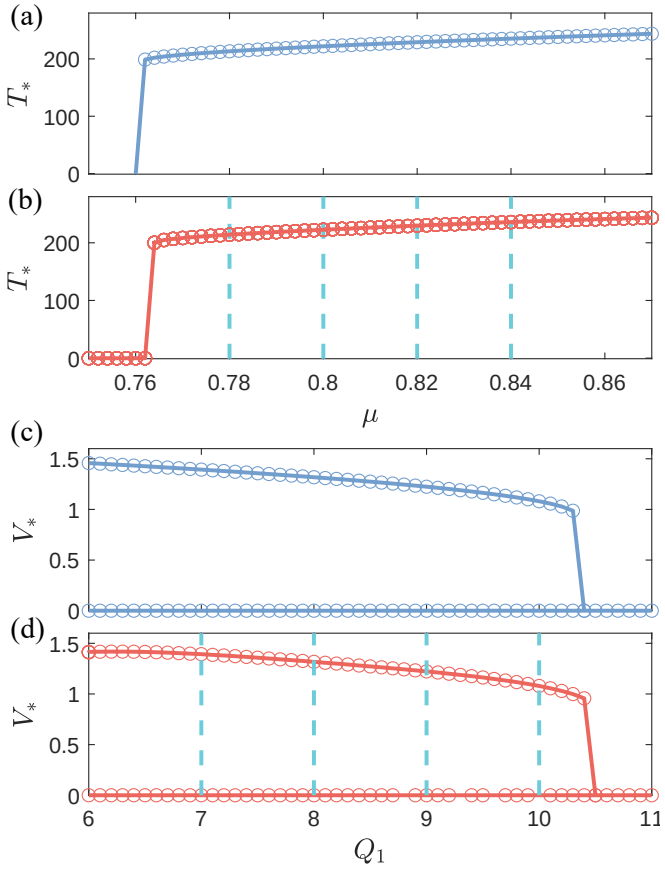

FIG. 10. Anticipating tipping in a climate model. (a) Tipping behavior of the equilibrium temperature T_* with respect to the bifurcation parameter μ. (b) Successful prediction of the tipping behavior by parameter-aware reservoir computing, where the bifurcation-parameter values for training are marked by the vertical cyan dashed lines.

We consider the climate model for the ice-albedo feedback and represents the global temperature as a zero-dimensional average field [62, 63]. The deterministic version of the model is given by [62, 63]:

$$\begin{aligned}\frac{dT}{dt} = f_c(T) &= \frac{1}{4}\mu I_0(1-(a_2-b_2T^2)) - e_{SA}\sigma_S T^4 \\ &= -e_{SA}\sigma_S T^4 + \frac{\mu I_0 b_2}{4}T^2 + \frac{\mu I_0(1-a_2)}{4},\end{aligned} \tag{9}$$

where $\frac{1}{4}\mu I_0(1-(a_2-b_2T^2))$ is the incoming solar radiation subtracted by the reflected part proportional to the albedo $(a_2-b_2T^2)$, and $e_{SA}\sigma_S T^4$ is the outgoing radiation described by the Stefan-Boltzmann law. Following Ref. [63], we treat parameter μ as the time-varying bifurcation parameter. Other parameter values are $I_0 = 1366$, $e_{SA} = 0.62$, $\sigma_S = 5.6704 \times 10^{-8}$, $b_2 = 1.69 \times 10^{-5}$, and $a_2 = 1.6927$. The stochastic version of the system is given by

$$dT = f_c(T)dt + \sigma_{\text{noise}}dW, \tag{10}$$

where dW is a normalized Wiener (white noise) process. To generate the training data, we numerically integrate this stochastic equation by the Heun method [64] with an additive noise of strength $\sigma_{\text{noise}} = 0.1$ and time step of $\Delta t = 0.004$. A tipping point occurs at $\mu_c = 0.761$, as shown in Fig. 10(a). Figure 10(b) shows the prediction results by our parameter-adaptable reservoir computer. It can be seen that the tipping behavior has been correctly predicted, based on training data from the pre-tipping regime.

IV. BENEFICIAL ROLE OF NOISE IN ANTICIPATING TIPPING

An attractor made of a simple fixed point is zero-dimensional and is essentially featureless with limited information about the dynamics of the target system. A machine-learning model is unable to learn the dynamics based on time series with only constant values. In related previous works [65, 66], transient behaviors of the target systems before reaching the asymptotic states were exploited in training to overcome this difficulty, where the target system is initialized at locations in the phase space away from the asymptotic states. Another possibility is exploiting dynamical noise that drives the target system around the fixed point. A certain level of dynamical noise can in fact be advantageous for parameter-adaptable reservoir computing, as it provides the opportunity for the machine to explore a larger phase-space region. As the target system approaches a tipping point, the dominant eigenvalue of the Jacobian matrix around the fixed point is approaching zero, so the system landscape becomes flatter and noise can effectively enlarge the phase-space region of exploration for the reservoir computer to learn the dynamics.

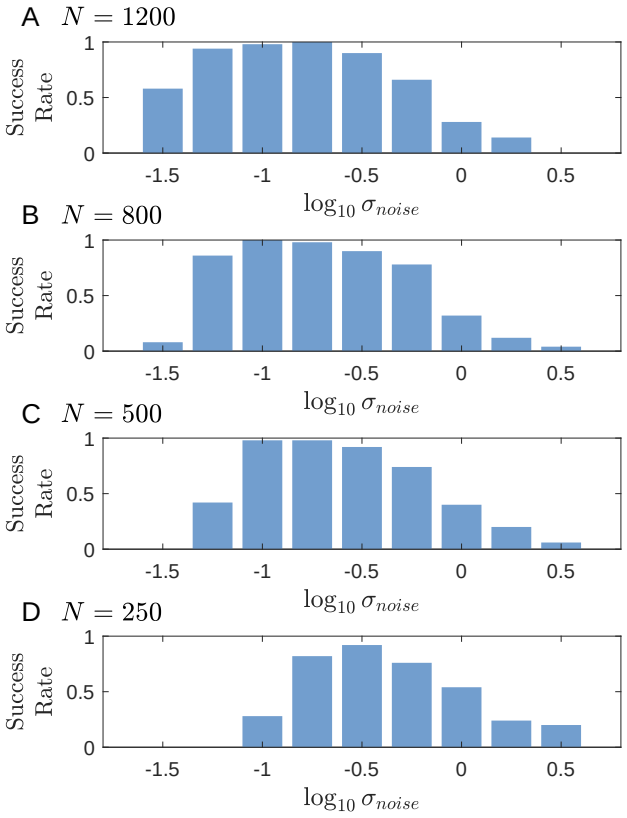

FIG. 11. Anticipation performance with different noise levels and varying complexity of reservoir computing. The rates of successful anticipation for different levels of $\sigma_{\rm noise}$ are obtained for reservoir network of size (a) $N = 1,200$, (b) $N = 800$, (c) $N = 500$, and (d) $N = 250$. An optimal region of the noise strength arises for stable and relatively accurate anticipation, which increases with the size of the reservoir network. A larger reservoir network can also exploit training data at a lower noise level, while a smaller network performs better at exploiting training data with a higher noise level.

Our computations reveal an optimal region of the noise strength for stable and relatively accurate anticipation results. Some too large or too small noise level $\delta_{\rm noise}$ tends to reduce the performance. As the reservoir network becomes larger so that it possesses a higher level of dynamical complexity, the optimal $\delta_{\rm noise}$ region becomes wider. Figure 11 shows histograms of the anticipation accuracy for different noise levels on the climate model with different sizes of the reservoir network, where a success rate is defined as the fraction of the testing results that successfully anticipate the tipping and predict the position of the tipping point with a relative error of less than 50%. For the climate model, a predicted μ_c within the interval of $[0.751, 0.772]$ is considered as successful. Each success rate is calculated from an ensemble of 50 statistical realizations of the reservoir network and training data. For weak dynamical noise, the training data are close to some constant values, making it difficult for the reservoir computer to grasp the dynamical features, leading to the anticipation that the target system is near a stable steady state. While the trend by which the equilibrium point changes slowly with the bifurcation parameter can be anticipated as it requires only the zero-order information about the equilibrium in the training set, the reservoir computer can hardly predict a tipping point, which requires higher-order dynamical features of the target system. The small dynamical fluctuations can also be overwhelmed by observational noises.

Large dynamical noise is harmful, too. Intuitively, in this case, the finite training data may not be statistically representative for the machine-learning framework to learn the deterministic component of the target system near the equilibrium without overfilling. Extracting the deterministic component from its noisy trajectory is similar to the problem of information transmission through a noisy communication channel [67]. To see this, consider a one-dimensional dynamical system $dx = f(x)dt + \sigma_{noise}dW$, where dW is a normalized Wiener (white noise) process and $f(t)$ is the deterministic component of the dynamics. Let there be a stable fixed point at $x_0 = 0$, about which we have $f(x) = -x + o(x)$. The system's capacity corresponds to the information channel capacity and is given by [67]

$$C = \frac{1}{2\sigma_{noise}^2} \int f(x)^2 P(x) dx \tag{11}$$

$$= \frac{1}{2\sigma_{noise}^2} \int (x^2 - 2xo(x)) P(x) dx, \tag{12}$$

where $P(x)$ is the distribution of x under dynamical noise. The capacity C measures the maximal rate of information transmission. If we ignore high-order terms $o(x)$, $P(x)$ is a normal distribution $P(x) \sim \mathcal{N}(0, \sigma_{\rm noise}^2/2)$. However, this would result in a constant C independent of σ_{noise}, because C measures the maximal information transmission rate about the features of $f(x)$ in the entire phase space. Anticipating the tipping point requires focusing on the local behavior of $f(x)$ about x_0, and the information about $f(x)$ away from x_0 may not be useful. As a crude approximation, we integrate Eq. 12 within a neighborhood $x \in (-d, d)$ of the fixed point and find that C will quickly approach zero as the noise level increases through d. (A more accurate approximation is to use a mask distribution in the integration to put larger weights near the fixed point.) This provides a heuristic understanding of the disadvantage of large noise.

In realistic applications, it is not feasible to decide or control the noise level of the observed data. However, as shown in Fig. 11, the optimal noise region also depends on the complexity of the reservoir network, where a larger network is better at exploiting the relatively small fluctuations caused by a lower noise level than a smaller network, and can also make the optimal noise region wider. Overfitting can be a problem with large reservoir networks. As a result, the optimal noise level shifts to the lower side as the reservoir network size increases.

V. DISCUSSION

Recent years have witnessed significant efforts in developing machine-learning models for predicting critical transitions in nonlinear dynamical systems. A tacit assumption in these works is that oscillatory time series are available for training the neural network. For critical transitions such as crises, synchronization onset and amplitude death, this requirement can indeed be met. Tipping, by its historical origin from nonlinear ecosystems, is a different type of critical transition in that the system is in some sort of stable steady state before and after the transition. In a deterministic system, the available time series are non-oscillatory. The lack of pre-tipping oscillations means that the usual temporal variations used for training are absent, making it significantly more challenging to predict the impending shift using machine learning. This is a reason that most previous works used *detection-based* approaches to extracting early warning signals or features from observed time series before the tipping.

We developed a *forecasting-based*, machine-learning framework to anticipate tipping in nonautonomous dynamical systems by taking advantage of noise. Our prediction is based on the presently available time-series data in a stable steady state but under the influence of noise. For synthetic data from a dynamical model, we incorporate stochasticity into the governing equations to generate noisy time series. For empirical data from the real word, most likely they are already noisy. In any case, the random oscillations associated with the noisy time series make it possible to train a machine-learning model, such as reservoir computing. We tested our parameter adaptable reservoir-computing scheme on a variety of systems from diverse fields, all sharing the common tipping scenario: sudden transition from one steady state to another as a bifurcation parameter passes through a critical point. The main application is predicting the potential collapse of the AMOC. Using simulated and empirical fingerprint data, our results suggest that the AMOC could halt in a time window centered about the year 2055, with the earliest possible occurrence in year 2040. These are consistent with the recent results based on a statistical optimization approach [52].

Our machine-learning method is predicated on the availability of a known time-varying parameter that drives the system towards tipping [42]. A simple assumption, in the absence of detailed information, is that the control parameter changes gradually and approaches its unknown critical value linearly over time [52, 58]. However, this linear assumption does not fully capture the complexities of real-world scenarios. The exact nature of the time-varying parameter and its real-time changes remain ambiguous. It has been found that the AMOC is sensitive to variations in the ocean's freshwater forcing [52, 58, 68] that can manifest through surface freshwater fluxes such as precipitation or through the input of freshwater from river runoff and ice melt, including significant contributions from the Greenland Ice Sheet. More sophisticated models suggested that the freshwater flux exhibits a quasi-exponential behavior [69, 70]. We have studied the case where the time-varying parameter $\lambda(t)$ changes exponentially over time. This assumption aligns more closely with observed behaviors and enhances the accuracy of our machine-learning prediction, enabling better anticipation of a potential tipping of the AMOC.

Another technical issue is whether the predicted collapse is merely an artifact caused by the reservoir computer operating in an untrained parameter region. To address this concern, we conducted simulations to test the extrapolation results on the other side of the untrained parameter region using synthetic AMOC data. For example, we tested λ values smaller than those in the trained parameter interval $\lambda \in (-2.5, -1.5)$ in the 1D AMOC model. In the simulations, we conducted 500 testing trials. For all these trials, no collapse was observed in this parameter region. The reservoir computer consistently and persistently predicted that the system remained in a healthy, stable steady state with no indication of an impending collapse. This consistent behavior across numerous trials strongly suggests that the predicted collapse is not an artifact of the reservoir computer functioning outside its trained parameter region, but rather robust prediction of the system's dynamics.

DATA AVAILABILITY

The data generated in this study, including both the training time series and the weights of the reservoir computers, can be found in the repository: https://github.com/SPanahi/RC_Tipping; The CESM data are also available from: https://github.com/RenevanWesten/SA-AMOC-Collapse/tree/main/Data

CODE AVAILABILITY

The codes used in this paper can be found in the repository: https://github.com/SPanahi/RC_Tipping

ACKNOWLEDGMENTS

This work was supported by the Air Force Office of Scientific Research under Grant No. FA9550-21-1-0438 and by the US Army Research Office under Grant No. W911NF-24-2-0228.

Appendix A: Notion of "tipping point" in the literature

In recent years, tipping points have garnered significant attention across diverse scientific disciplines [3, 4, 6, 11, 16, 17, 71–73]. There is a public interest in the

phenomenon of tipping as well, e.g., Malcolm Gladwell's popular book entitled "The Tipping Point" [74]. In climate science, a tipping point manifests itself as an abrupt transition from one regime to another within some time frame [75]. In ecology, the term denotes a critical threshold, often regarded as a point of no return, where minor changes in the environmental conditions can drive the system into a fundamentally altered state [76]. In dynamical systems, a common situation for a tipping point is multistability. Tipping can also occur in monostable fast-slow dynamical systems [77].

Despite the broad spectrum of domains in which tipping can arise, Kuehn's work [78] identified several common attributes shared by most tipping phenomena. These include (1) a sudden qualitative shift in the system's behavior, (2) faster changes compared to regular dynamics, (3) the crossing of a specific threshold near a transition, (4) the emergence of a new state significantly distant from the preceding state, (5) the presence of noise in a deterministic system, (6) a slow recovery from perturbations, (7) escalating variance as the transition approaches, (8) more asymmetric noisy fluctuations, and (9) an increase in autocorrelation prior to a transition. Another attribute of tipping is some hysteresis behavior associated with multistability [73]. In such an instance, the trajectory leading to the tipping point differs significantly during forward and backward shifts, influencing the system's response and recovery. The hysteresis near a tipping point, particularly in ecosystems, poses a significant challenge, as it complicates efforts to restore the system to its previous state following the tipping transition.

While diverse types of tipping points have been reported in the literature [47], a comprehensive classification of tipping points in dynamical systems was proposed in 2011 [63], based on the tipping mechanism. The most extensively studied type is bifurcation-induced tipping (B-tipping) [79], where a small alteration in the system's parameters leads to an abrupt and qualitative change in the system's state. Another type is noise-induced tipping (N-tipping) associated with transitions between states due to noisy fluctuations [80]. A more recently discovered type is rate-induced tipping (R-tipping), where a time-varying input or parameter of the system causes the system to "tip" away from its normal states [63]. Overall, B-tipping is triggered by a critical level of external inputs or a critical value of a system parameter, N-tipping occurs due to the presence of noisy fluctuations, and R-tipping arises when the moving stable state cannot adapt to an external input or a time-varying parameter of the system [81].

Appendix B: Anticipating tipping versus predicting critical transitions

Our work differs from the recent work on machine-learning prediction of critical transitions in nonlinear dynamical systems [42], where the former is predicting tipping from one stable steady state to another and the latter is predicting a transition from an oscillatory dynamical state to a collapsed state. More specifically, oscillatory behaviors in the data in the pre-critical regime have the benefit of system trajectory's visiting a substantial portion of the phase space, thereby facilitating training by enabling the neural network to effectively learn the phase-space behavior or the dynamical climate of the target system. However, for tipping in a deterministic system, in the pre-tipping regime the system is in a stable steady state without oscillations in its dynamical variables. In this case, stochasticity or noise leading to randomly oscillating dynamical variables is essential to neural-network training. In our study, we exploit dynamic noise in the data for training, where validation and hyperparameter optimization are performed based on data in the pre-critical regime. During the test or prediction phase, the reservoir computer operates as a closed-loop, deterministic dynamical system capable of predicting how the dynamical climate of the system changes with the bifurcation parameter. Since no data from the target system in the post-tipping regime were used for training (in a realistic situation, such data are not available), it is not possible for the reservoir computer to correctly predict the system's behavior after the tipping. However, the neural machine is capable of generating characteristic changes in the output variables at the tipping transition, making its anticipation possible.

A marked difference between our approach and the method in Ref. [42] stems from variations in the target systems under consideration. The assumption in Ref. [42] is that stochastic effects in the target system are relatively small compared to the deterministic component. It was demonstrated that even with relatively small observational noise, the results remained largely unaffected. In contrast, our study assigns crucial roles to stochastic effects. Omitting stochastic noises renders the training data with no variations in time and devoid of distinctive features. During the prediction phase, our closed-loop reservoir network operates as a deterministic system. Consequently, evaluating metrics such as root mean square errors during validation or hyperparameter optimization provides limited insights into the reservoir computer's performance. Validation can only discern stability in the prediction with no abrupt collapse in a specific hyperparameter region. While this stable region is generally broader than the optimized region in Ref. [42], it does not hinder the ability to successfully anticipate tipping.

Appendix C: A detailed account of the 1D AMOC fingerprint model

The AMOC is a crucial component of the global ocean circulation system, playing a vital role in regulating global climate patterns. However, due to the challenges

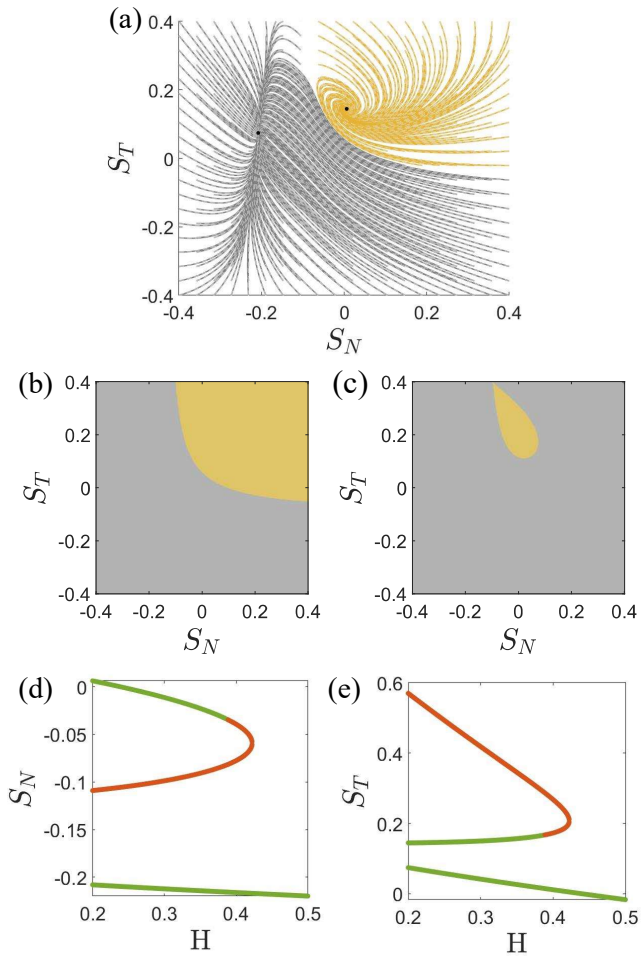

FIG. 12. Phase-space structure and bifurcation diagram of the three-box AMOC model (D1). (a) An example of the vector field of the model with the parameter values specified in Tabs. II and III). (b,c) Basin of attraction of the stable equilibria for two distinct values of the freshwater forcing parameter: $H = 0.2$ and $H = 0.35$, respectively. The basins associated with the high (low) state is highlighted in yellow (gray). Changing the freshwater forcing parameter results in a decrease (increase) in the basin associated with the "on (off)-state." (d) A bifurcation diagram under the atmospheric conditions of doubled pre-industrial CO_2.

associated with continuously monitoring the AMOC and the limited availability of long-term observational data, researchers have turned to fingerprint analysis techniques to gain insights into the dynamics of the AMOC [53]. One such fingerprint analysis technique involves using the SST data as a proxy for assessing the AMOC strength [54–57], where SST is an important fingerprint of the AMOC due to its sensitivity to changes in the ocean circulation patterns and its ability to capture the variations in heat transport within the North Atlantic [82–84]. In general, the fingerprints offer the possibility of detecting changes in the AMOC earlier than direct observations and extend time series data into the past, potentially enabling anticipation and understanding of the shifts in this critical ocean circulation system.

The 1D AMOC stochastic fingerprint model was introduced recently to describe the AMOC dynamics with a time-varying control parameter [52]. The basic assumption is that the AMOC is in equilibrium before undergoing a transition, where the bifurcation parameter, denoted as λ, undergoes slow evolution towards an unknown critical value. In spite of its simplicity, the 1D model produces time series in close alignment with observed AMOC fingerprint data. The primary drivers such as freshwater flux or the logarithm of atmospheric CO_2 concentration are not reflected in the model, but the fingerprint data generated are robust. In the vicinity of tipping, the model is described by a dynamical variable X_t governed by the following stochastic differential equation:

$$\dot{X}_t = -(A(X_t - m)^2 + \lambda)) + \sigma dB_t, \tag{C1}$$

$$\lambda = \begin{cases} \lambda_0 & t < t_0 \\ \lambda_0(1 - \frac{t-t_0}{t_c-t_0}), \end{cases} \tag{C2}$$

where A is a time scale parameter, $m = \mu - \sqrt{|\lambda|/A}$ with μ being the stable fixed point, B_t denotes a Brownian motion, and σ is the noise amplitude. The underlying deterministic system exhibits a tipping point triggered by a saddle-node bifurcation at $\lambda = \lambda_c = 0$. Initially, the system is in a statistically stable state with a constant $\lambda = \lambda_0$. Starting from time t_0, λ undergoes a linear change towards λ_c. The actual tipping time t_c exhibits fluctuations due to the stochastic forcing and varies across different realizations. Figure 3(b) shows ten different realizations of the tipping event for $A = 0.95$, $m = -1.3$, $\lambda_0 = -2.7$, $\sigma = 0.3$, $t_0 = 1924$ and $\lambda_c = 0$.

Appendix D: Three-box AMOC model

TABLE II. Salinity and flux initial conditions for the three-box AMOC model with doubled atmospheric CO_2 ($2\times$ CO_2)

Salinity	Flux ($m^3\,s^{-1}$)	
$S_{N0} = 0.034912$	$F_{N0} = 0.486 \times 10^6$	$F_{N1} = 0.1311 \times 10^6$
$S_{T0} = 0.035435$	$F_{T0} = -0.997 \times 10^6$	$F_{T1} = 0.6961 \times 10^6$
$S_{S0} = 0.034427$	$F_{S0} = 1.265 \times 10^6$	
$S_{IP0} = 0.034668$	$F_{IP0} = -0.754 \times 10^6$	
$S_{B0} = 0.034538$		

There are a number of models of the AMOC in the literature, ranging in complexity from simple box models to the intricate atmosphere-ocean general circulation models. These models differ in their structure and parameters, such as the five-box model [85] or the reduced three-box model [86]. In the original five-box model, the boxes correspond to different ocean regions: North Atlantic (N), Tropical Atlantic (T), Indo-Pacific (IP), Southern Ocean (S), and bottom waters (B). The AMOC

flow (q) is directly proportional to the density gradient of the temperatures and salinity in the N and S boxes. In the reduced three-box model, variations in the salinity of the Southern Ocean (S_S) and bottom waters (S_B) are treated as constants over time due to their relatively small and slow rate of change compared to the salinity variations in the North Atlantic (S_N), Tropical Atlantic (S_T), and the Indo-Pacific (S_{IP}). Taking the constant salt (C) and S_{IP} as a dependent variable, the governing equations of the AMOC dynamics are

$$q = \frac{\lambda\left[\alpha\left(T_s - T_0\right) + \beta(S_N - S_s)\right]}{1 + \lambda\alpha\mu}, \tag{D1}$$

where

$$V_N \frac{dS_N}{dt} = q\left(S_T - S_N\right) + K_N\left(S_T - S_N\right) - (F_{N0} + F_{N1}H))S_0, \tag{D2}$$

$$V_T \frac{dS_T}{dt} = q\left[\gamma S_S + (1-\gamma)\, S_{IP} - S_T\right] + K_S\left(S_S - S_T\right) + K_N\left(S_N - S_T\right) - (F_{T0} + F_{T1}H))S_0, \tag{D3}$$

for $q \geq 0$ and

$$V_N \frac{dS_N}{dt} = |q|\left(S_B - S_N\right) + K_N\left(S_T - S_N\right) - (F_{N0} + F_{N1}H))S_0, \tag{D4}$$

$$V_T \frac{dS_T}{dt} = |q|\left(S_N - S_T\right) + K_S\left(S_S - S_T\right) + K_N\left(S_N - S_T\right) - (F_{T0} + F_{T1}H))S_0, \tag{D5}$$

for $q < 0$. The parameter H corresponds to the freshwater fluxes and S_{IP} is determined by the following equation:

$$C = V_N S_N + V_T S_T + V_S S_S + V_{IP} S_{IP} + V_B S_B. \tag{D6}$$

The parameter values of the three-box model are listed in Tables II and III.

TABLE III. Parameter values of the three-box AMOC model with doubled atmospheric CO_2 ($2\times$ CO_2)

Volume (m^3)	Parameters	Flux Parameters ($m^3\, s^{-1}$)	Parameters
$V_N = 0.3683 \times 10^{17}$	$S_0 = 0.035$	$K_N = 1.762 \times 10^6$	$\lambda = 1.62 \times 10^7 m^6 kg^{-1} s^{-1}$
$V_T = 0.5418 \times 10^{17}$	$T_0 = 3.870°C$	$\eta = 33.264 \times 10^6$	$\gamma = 0.36$
$V_S = 0.6097 \times 10^{17}$	$T_S = 7.919°C$	$K_S = 1.872 \times 10^6$	$\mu = 22 \times 10^{-8\circ} C^{-1} m^{-3} s^{-1}$
$V_{IP} = 1.4860 \times 10^{17}$	$\alpha = 0.12\, kg°C^{-1} m^{-3}$	$K_{IP} = 99.977 \times 10^6$	
$V_B = 9.9250 \times 10^{17}$	$\beta = 790.0 kg m^{-3}$	$H \in [0.2, 0.5]$	

Figure 12(a) presents an example of an AMOC fingerprint: tropical Atlantic salinity in relation to north Atlantic salinity. The three-box model has two stable equilibria as indicated by the black dots. The locations of these equilibria and their basins of attraction depend on the model's parameter values. For example, Figs. 12(b) and 12(c) show the basin of attraction for two different values of the parameter H. The basins associated with the high and low states are shown in yellow and gray, respectively. As the freshwater forcing parameter H is increased from 0.2 to 0.35, the basin of the high state diminishes till it undergoes extinction at $H_c = 0.388$. Figure 12(d) shows a bifurcation diagram of the three-box model with respect to the parameter $H \in [0.2, 0.5]$, where the two stable equilibria are represented by the two solid green curves. The upper (lower) equilibrium corresponds to the "on state" ("off state)." At a critical value ($H_{cr} = 0.388$), the system undergoes a saddle-node bifurcation where the high stable equilibrium and the unstable equilibrium collide and annihilate together, leaving the low stable equilibrium as the only attractor of the system.

We now use reservoir computing to predict the tipping time. In general, augmenting the freshwater input from melting glaciers can impact the AMOC system and bring it into closer to a tipping point, as illustrated in

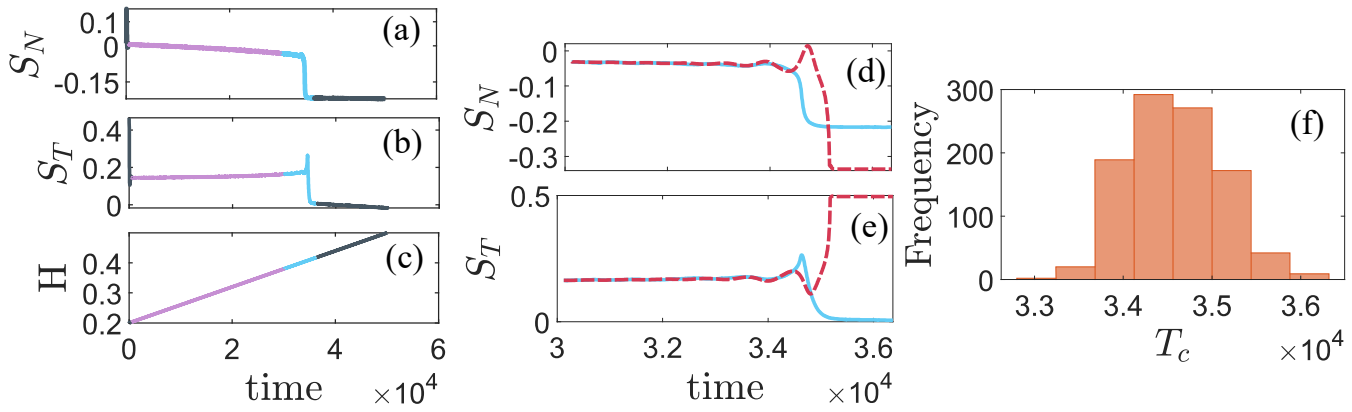

FIG. 13. Reservoir-computing prediction of the time window of AMOC collapse in the time-dependent three-box AMOC model. The fresh-water flux parameter $H(t)$ is time dependent, making the system nonautonomous. (a-c) Different realizations of the AMOC fingerprint. The data are divided into three segments: transients, training, and testing, depicted in black, purple, and blue, respectively. The vertical red dotted line indicates the critical time T_c of tipping under dynamical noise of amplitude $\sigma = 5 \times 10^{-5}$. The critical point is $H_c = 0.359$. (d,e) An illustrative example of the testing data as represented by the solid blue line, while the reservoir-computing prediction is indicated by the dash-dotted line. (f) Histogram of the predicted critical point T_c from 1000 random realizations of the reservoir computer.

Figs. 13(a-c) by three different time-series realizations for different values of H. The training data consist of time series for S_N and S_T for $H \in [0.20, 0.38]$, representing the system's behavior associated with the "on-state" equilibrium. To introduce random fluctuations about the fixed points in the training data, we incorporate certain levels of dynamical noise into the deterministic system. During the testing phase, the control parameter continues to increase into the untrained regime, enabling a potential tipping to be anticipated. Some representative prediction results are shown in Figs. 13(d) and 13(e). From an ensemble of 1000 random reservoir realizations, we obtain a histogram the anticipated tipping point, as shown in Fig. 13(f), where tipping occurs for all the trials and the anticipated tipping time falls within the interval $T_c \in [3.3 \times 10^4, 3.6 \times 10^4]$ about the ground truth from the model ($T_c = 3.48 \times 10^4$).

The results in Fig. 13 imply that the system might undergo a tipping in roughly the next 350 centuries, which are inconsistent with those from the recent study [52] that predicts the potential collapse of the AMOC around the mid-century under the current scenario of emissions. Discrepancies between different models and studies are not uncommon in climate science, where variations in the complexity, assumptions, and input parameters among the models can lead to divergent predictions. Additionally, uncertainties in climate projections may stem from different sources, including variations in the modeling techniques, emission scenarios, and natural climate variability.

[1] M. Scheffer, *Ecology of Shallow Lakes* (Springer Science & Business Media, 2004).

[2] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara, Early-warning signals for critical transitions, Nature **461**, 53 (2009).

[3] M. Scheffer, Complex systems: foreseeing tipping points, Nature **467**, 411 (2010).

[4] D. B. Wysham and A. Hastings, Regime shifts in ecological systems can occur with no warning, Ecol. Lett. **13**, 464 (2010).

[5] J. M. Drake and B. D. Griffen, Early warning signals of extinction in deteriorating environments, Nature **467**, 456 (2010).

[6] L. Chen, R. Liu, Z.-P. Liu, M. Li, and K. Aihara, Detecting early-warning signals for sudden deterioration of complex diseases by dynamical network biomarkers, Sci. Rep. **2**, 342 (2012).

[7] C. Boettiger and A. Hastings, Quantifying limits to detection of early warning for critical transitions, J. R. Soc. Interface **9**, 2527 (2012).

[8] L. Dai, D. Vorselen, K. S. Korolev, and J. Gore, Generic indicators for loss of resilience before a tipping point leading to population collapse, Science **336**, 1175 (2012).

[9] J. M. Drake and B. D. Griffen, Early warning signals of extinction in deteriorating environments, Nature **467**, 456 (2010).

[10] P. Ashwin, S. Wieczorek, R. Vitolo, and P. Cox, Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system, Philo Trans. R. Soc. A Math. Phys. Eng. Sci. **370**, 1166 (2012).

[11] T. M. Lenton, V. N. Livina, V. Dakos, E. H. van Nes, and M. Scheffer, Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness, Phil. Trans. Roy. Soc. A **370**, 1185 (2012).

[12] A. D. Barnosky, E. A. Hadly, J. Bascompte, E. L. B. J. H.

Brown, M. Fortelius, W. M. Getz, J. Harte, A. Hastings, P. A. Marquet, N. D. Martinez, A. Mooers, P. Roopnarine, G. Vermeij, J. W. Williams, R. Gillespie, J. Kitzes, C. Marshall, N. Matzke, D. P. Mindell, E. Revilla, and A. B. Smith, Approaching a state shift in earth's biosphere, Nature **486**, 52 (2012).

[13] C. Boettiger and A. Hastings, Tipping points: From patterns to predictions, Nature **493**, 157 (2013).

[14] J. M. Tylianakis and C. Coux, Tipping points in ecological networks, Trends. Plant. Sci. **19**, 281 (2014).

[15] J. J. Lever, E. H. Nes, M. Scheffer, and J. Bascompte, The sudden collapse of pollinator communities, Ecol. Lett. **17**, 350 (2014).

[16] T. S. Lontzek, Y.-Y. Cai, K. L. Judd, and T. M. Lenton, Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy, Nat. Clim. Change **5**, 441 (2015).

[17] S. Gualdia, M. Tarziaa, F. Zamponic, and J.-P. Bouchaudd, Tipping points in macroeconomic agent-based models, J. Econ. Dyn. Contr. **50**, 29 (2015).

[18] J. Jiang, Z.-G. Huang, T. P. Seager, W. Lin, C. Grebogi, A. Hastings, and Y.-C. Lai, Predicting tipping points in mutualistic networks through dimension reduction, Proc. Nat. Acad. Sci. (USA) **115**, E639 (2018).

[19] B. Yang, M. Li, W. Tang, S. Liu, Weixinand Zhang, L. Chen, and J. Xia, Dynamic network biomarker indicates pulmonary metastasis at the tipping point of hepatocellular carcinoma, Nat. Commun. **9**, 678 (2018).

[20] J. Jiang, A. Hastings, and Y.-C. Lai, Harnessing tipping points in complex ecological networks, J. R. Soc. Interface **16**, 20190345 (2019).

[21] M. Scheffer, *Critical Transitions in Nature and Society*, Vol. 16 (Princeton University Press, 2020).

[22] Y. Meng, J. Jiang, C. Grebogi, and Y.-C. Lai, Noise-enabled species recovery in the aftermath of a tipping point, Phys. Rev. E **101**, 012206 (2020).

[23] Y. Meng, Y.-C. Lai, and C. Grebogi, Tipping point and noise-induced transients in ecological networks, J. R. Soc. Interface **17**, 20200645 (2020).

[24] Y. Meng and C. Grebogi, Control of tipping points in stochastic mutualistic complex networks, Chaos **31**, 023118 (2021).

[25] Y. Meng, Y.-C. Lai, and C. Grebogi, The fundamental benefits of multiplexity in ecological networks, J. R. Soc. Interface **19**, 20220438 (2022).

[26] P. E. O'Keeffe and S. Wieczorek, Tipping phenomena and points of no return in ecosystems: beyond classical bifurcations, SIAM J. Appl. Dyn. Syst. **19**, 2371 (2020).

[27] C. Trefois, P. M. Antony, J. Goncalves, A. Skupin, and R. Balling, Critical transitions in chronic disease: transferring concepts from ecology to systems medicine, Cur. Opin. Biotechnol. **34**, 48 (2015).

[28] A. Bayani, F. Hadaeghi, S. Jafari, and G. Murray, Critical slowing down as an early warning of transitions in episodes of bipolar disorder: A simulation study based on a computational model of circadian activity rhythms, Chronobiol. Int. **34**, 235 (2017).

[29] C. A. Nobre and L. D. S. Borma, 'tipping points' for the amazon forest, Curr. Opin. Env. Sust. **1**, 28 (2009).

[30] K. R. Miner, M. R. Turetsky, E. Malina, A. Bartsch, J. Tamminen, A. D. McGuire, A. Fix, C. Sweeney, C. D. Elder, and C. E. Miller, Permafrost carbon emissions in a changing arctic, Nat. Rev. Earth Environ. **3**, 55 (2022).

[31] P. Wadhams, Arctic ice cover, ice thickness and tipping points, Ambio **41**, 23 (2012).

[32] M. W. Buckley and J. Marshall, Observations, inferences, and mechanisms of the Atlantic Meridional Overturning Circulation: A review, Rev. Geophys. **54**, 5 (2016).

[33] J. Lohmann and P. D. Ditlevsen, Risk of tipping the overturning circulation due to increasing rates of ice melt, Proc. Natl. Acad. Sci. (USA) **118**, e2017989118 (2021).

[34] L. C. Jackson, A. Biastoch, M. W. Buckley, D. G. Desbruyères, E. Frajka-Williams, B. Moat, and J. Robson, The evolution of the North Atlantic Meridional overturning Circulation since 1980, Nat. Rev. Earth Environ. **3**, 241 (2022).

[35] K. E. Trenberth, Y. Zhang, J. T. Fasullo, and L. Cheng, Observation-based estimates of global and basin ocean meridional heat transport time series, J. Clim. **32**, 4567 (2019).

[36] A. Biastoch, C. W. Böning, J. Getzlaff, J.-M. Molines, and G. Madec, Causes of interannual–decadal variability in the meridional overturning circulation of the midlatitude North Atlantic Ocean, J. Clim. **21**, 6599 (2008).

[37] S. Yeager and G. Danabasoglu, The origins of late-twentieth-century variations in the large-scale North Atlantic circulation, J. Clim. **27**, 3222 (2014).

[38] C. Grebogi, E. Ott, and J. A. Yorke, Fractal basin boundaries, long-lived chaotic transients, and unstable-unstable pair bifurcation, Phys. Rev. Lett. **50**, 935 (1983).

[39] H. Fan, L.-W. Kong, Y.-C. Lai, and X. Wang, Anticipating synchronization with machine learning, Phys. Rev. Res. **3**, 023237 (2021).

[40] R. Xiao, L.-W. Kong, Z.-K. Sun, and Y.-C. Lai, Predicting amplitude death with machine learning, Phys. Rev. E **104**, 014205 (2021).

[41] D. Patel, D. Canaday, M. Girvan, A. Pomerance, and E. Ott, Using machine learning to predict statistical properties of non-stationary dynamical processes: System climate, regime transitions, and the effect of stochasticity, Chaos **31**, 033149 (2021).

[42] L.-W. Kong, H.-W. Fan, C. Grebogi, and Y.-C. Lai, Machine learning prediction of critical transition and system collapse, Phys. Rev. Res. **3**, 013090 (2021).

[43] L.-W. Kong, H. Fan, C. Grebogi, and Y.-C. Lai, Emergence of transient chaos and intermittency in machine learning, J. Phys. Complex. **2**, 035014 (2021).

[44] D. Patel and E. Ott, Using machine learning to anticipate tipping points and extrapolate to post-tipping dynamics of non-stationary dynamical systems, Chaos **33** (2023).

[45] L.-W. Kong, Y. Weng, B. Glaz, M. Haile, and Y.-C. Lai, Reservoir computing as digital twins for nonlinear dynamical systems, Chaos **33**, 033111 (2023).

[46] Z.-M. Zhai, L.-W. Kong, and Y.-C. Lai, Emergence of a resonance in machine learning, Phys. Rev. Res. **5**, 033127 (2023).

[47] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara, Early-warning signals for critical transitions, Nature **461**, 53 (2009).

[48] C. Boettiger, N. Ross, and A. Hastings, Early warning signals: the charted and uncharted territories, Theor. Ecol. **6**, 255 (2013).

[49] I. A. van de Leemput, M. Wichers, A. O. Cramer, D. Borsboom, F. Tuerlinckx, P. Kuppens, E. H. van Nes, W. Viechtbauer, E. J. Giltay, S. H. Aggen, *et al.*, Critical slowing down as early warning for the onset and termination of depression, Proc. Natl. Acad. Sci. (USA) **111**,

87 (2014).
[50] N. Boers, Early-warning signals for dansgaard-oeschger events in a high-resolution ice core record, Nat. Commun. **9**, 1 (2018).
[51] T. M. Bury, R. Sujith, I. Pavithran, M. Scheffer, T. M. Lenton, M. Anand, and C. T. Bauch, Deep learning for early warning signals of tipping points, Proc. Natl. Acad. Sci. (USA) **118**, e2106140118 (2021).
[52] P. Ditlevsen and S. Ditlevsen, Warning of a forthcoming collapse of the Atlantic Meridional Overturning Circulation, Nat. Commun **14**, 4254 (2023).
[53] L. C. Jackson and R. A. Wood, Fingerprints for early detection of changes in the AMOC, J. Clim. **33**, 7027–7044 (2020).
[54] J. Robson, D. Hodson, E. Hawkins, and R. Sutton, Atlantic overturning in decline?, Nat. Geosci. **7**, 2–3 (2013).
[55] R. Zhang, On the persistence and coherence of subpolar sea surface temperature and salinity anomalies associated with the Atlantic multidecadal variability, Geophys. Res. Lett. **44**, 7865–7875 (2017).
[56] R. K. Haskins, K. I. C. Oliver, L. C. Jackson, R. A. Wood, and S. S. Drijfhout, Temperature domination of AMOC weakening due to freshwater hosing in two GCMs, Clim. Dyn. **54**, 273–286 (2019).
[57] M. Ben-Yami, V. Skiba, S. Bathiany, and N. Boers, Uncertainties in critical slowing down indicators of observation-based fingerprints of the Atlantic overturning circulation, Nat. Commun. **14**, 8344 (2023).
[58] R. van Westen, M. Kliphuis, and H. A. Dijkstra, Physics-based early warning signal shows that AMOC is on tipping course, Sci. Adv. **10**, 10.1126/sciadv.adk1189 (2024).
[59] J. M. Olesen, L. I. Eskildsen, and S. Venkatasamy, Invasion of pollination networks on oceanic islands: importance of invader complexes and endemic super generalists, Divers. Distrib. **8**, 181 (2002).
[60] A. Montero, *The Ecology of Three Pollination Networks*, Ph.D. thesis, Master's thesis (Aarhus Univ, Aarhus, Denmark) (2005).
[61] M. Scheffer, E. H. Van Nes, M. Holmgren, and T. Hughes, Pulse-driven loss of top-down control: the critical-rate hypothesis, Ecosys. **11**, 226 (2008).
[62] K. Fraedrich, Catastrophes and resilience of a zero-dimensional climate system with ice-albedo and greenhouse feedback, Q. J. R. Meteorol. Soc. **105**, 147 (1979).
[63] P. Ashwin, S. Wieczorek, R. Vitolo, and P. Cox, Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system, Philos. Transact. A Math. Phys. Eng. Sci. **370**, 1166 (2012).
[64] C. Van den Broeck, J. Parrondo, R. Toral, and R. Kawai, Nonequilibrium phase transitions induced by multiplicative noise, Phys. Rev. E **55**, 4084 (1997).
[65] J. Z. Kim, Z. Lu, E. Nozari, G. J. Pappas, and D. S. Bassett, Teaching recurrent neural networks to infer global temporal structure from local examples, Nat. Mach. Intell. **3**, 316 (2021).
[66] H. Fan, L. Wang, Y. Du, Y. Wang, J. Xiao, and X. Wang, Learning the dynamics of coupled oscillators from transients, Phys. Rev. Res. **4**, 013137 (2022).
[67] A. Frishman and P. Ronceray, Learning force fields from stochastic trajectories, Phys. Rev. X **10**, 021009 (2020).
[68] M. Taherkhani, S. Vitousek, P. L. Barnard, N. Frazer, T. R. Anderson, and C. H. Fletcher, Sea-level rise exponentially increases coastal flood frequency, Sci. Rep. **10**, 1 (2020).
[69] N. C. Swart and J. C. Fyfe, The influence of recent antarctic ice sheet retreat on simulated sea ice area trends, Geophys. Res. Lett. **40**, 4328–4332 (2013).
[70] T. Gorte, N. S. Lovenduski, C. Nissen, and J. T. M. Lenaerts, Antarctic ice sheet freshwater discharge drives substantial southern ocean changes over the 21st century, Geophys. Res. Lett. **50**, 10.1029/2023gl104949 (2023).
[71] P. Lamberson, S. E. Page, *et al.*, Tipping points, Quart. J. Polit. Sci. **7**, 175 (2012).
[72] M. Milkoreit, J. Hodbod, J. Baggio, K. Benessaiah, R. Calderón-Contreras, J. F. Donges, J.-D. Mathias, J. C. Rocha, M. Schoon, and S. E. Werners, Defining tipping points for social-ecological systems scholarship—an interdisciplinary literature review, Environ. Res. Lett. **13**, 033005 (2018).
[73] V. Dakos, B. Matthews, A. P. Hendry, J. Levine, N. Loeuille, J. Norberg, P. Nosil, M. Scheffer, and L. De Meester, Ecosystem tipping points in an evolving world, Nat. Ecol. Evol. **3**, 355 (2019).
[74] M. Gladwell, *The Tipping Point: How Little Things Can Make a Big Difference* (Little, Brown, 2006).
[75] T. M. Lenton, H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber, Tipping elements in the earth's climate system, Proc. Natl. Acad. Sci. (USA) **105**, 1786 (2008).
[76] J. M. Tylianakis and C. Coux, Tipping points in ecological networks, Tre. Plant Sci. **19**, 281 (2014).
[77] A. Vanselow, S. Wieczorek, and U. Feudel, When very slow is too fast-collapse of a predator-prey system, J. Theo. Biol. **479**, 64 (2019).
[78] C. Kuehn, A mathematical framework for critical transitions: Bifurcations, fast–slow systems and stochastic dynamics, Physica D **240**, 1020 (2011).
[79] J. M. T. Thompson and J. Sieber, Predicting climate tipping as a noisy bifurcation: a review, Int. J. Bif. Chaos **21**, 399 (2011).
[80] Y. Meng, Y.-C. Lai, and C. Grebogi, Tipping point and noise-induced transients in ecological networks, J. R. Soc. Interface. **17**, 20200645 (2020).
[81] S. Wieczorek, P. Ashwin, C. M. Luke, and P. M. Cox, Excitability in ramped systems: the compost-bomb instability, Proc. R. Soc. A Math. Phys. Eng. Sci. **467**, 1243 (2011).
[82] G. McCarthy, D. Smeed, W. Johns, E. Frajka-Williams, B. Moat, D. Rayner, M. Baringer, C. Meinen, J. Collins, and H. Bryden, Measuring the Atlantic Meridional Overturning Circulation at 26°n, Prog. Oceanogr. **130**, 91–111 (2015).
[83] G. D. McCarthy, I. D. Haigh, J. J.-M. Hirschi, J. P. Grist, and D. A. Smeed, Ocean impact on decadal Atlantic climate variability revealed by sea-level observations, Nature **521**, 508–510 (2015).
[84] M. S. Lozier, F. Li, S. Bacon, F. Bahr, A. S. Bower, S. A. Cunningham, M. F. de Jong, L. de Steur, B. deYoung, J. Fischer, S. F. Gary, B. J. W. Greenan, N. P. Holliday, A. Houk, L. Houpert, M. E. Inall, W. E. Johns, H. L. Johnson, C. Johnson, J. Karstensen, G. Koman, I. A. Le Bras, X. Lin, N. Mackay, D. P. Marshall, H. Mercier, M. Oltmanns, R. S. Pickart, A. L. Ramsey, D. Rayner, F. Straneo, V. Thierry, D. J. Torres, R. G. Williams, C. Wilson, J. Yang, I. Yashayaev, and J. Zhao, A sea change in our view of overturning in the subpolar North Atlantic, Science **363**, 516–521 (2019).

[85] R. A. Wood, J. M. Rodríguez, R. S. Smith, L. C. Jackson, and E. Hawkins, Observable, low-order dynamical controls on thresholds of the Atlantic Meridional Overturning Circulation, Clim. Dyn. **53**, 6815 (2019).

[86] H. Alkhayuon, P. Ashwin, L. C. Jackson, C. Quinn, and R. A. Wood, Basin bifurcations, oscillatory instability and rate-induced thresholds for Atlantic Meridional Overturning Circulation in a global oceanic box model, Proc. Math. Phys. Eng. Sci. **475**, 20190051 (2019).