

Machine-learning nowcasting of the Atlantic Meridional Overturning Circulation

Zheng-Meng Zhai,¹ Mohammadamin Moradi,¹ Shirin Panahi,¹ Zhi-Hua Wang,² and Ying-Cheng Lai^{1,3, a)}

¹*School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85287, USA*

²*School of Sustainable Engineering and the Built Environment, Arizona State University, Tempe, AZ 85287, USA*

³*Department of Physics, Arizona State University, Tempe, Arizona 85287, USA*

(Dated: 17 June 2024)

The Atlantic Meridional Overturning Circulation (AMOC) is a significant component of the global ocean system, which has so far ensured a relatively warm climate for the North Atlantic and mild conditions in regions such as Western Europe. The AMOC is also critical for the global climate. The complexity of the dynamical system underlying the AMOC is so vast that a long-term assessment of the potential risk of AMOC collapse is extremely challenging. However, short-term prediction can lead to accurate estimates of the dynamical state of the AMOC and possibly to early warning signals for guiding policy making and control strategies towards preventing AMOC collapse in the long term. We develop a model-free, machine-learning framework to predict the AMOC dynamical state in the short term by employing four datasets: MOVE and RAPID (observational), AMOC fingerprint (proxy records), AMOC simulated fingerprint and CESM AMOC (synthetic). We demonstrate the power of our framework in predicting the variability of the AMOC within the maximum prediction horizon of 12 or 24 months. A number of issues affecting the prediction performance are investigated.

I. INTRODUCTION

The Atlantic Meridional Overturning Circulation (AMOC) is a significant component of the global ocean system, transporting warmer, upper waters in the Atlantic northwards and colder, deeper waters southwards¹. More specifically, the AMOC is defined as the zonally and vertically integrated northward volume transport in terms of Sverdrups (Sv, $10^6 \text{m}^3 \text{s}^{-1}$), i.e., as a function of latitude and depth². The AMOC is the main reason that the climate of the North Atlantic is able to remain relatively warm, facilitating mild and livable conditions in regions such as Western Europe³. From the perspective of global climate system, the AMOC governs the frequency of Atlantic hurricanes and storms, tropical monsoons, and even the global carbon equilibrium^{4,5}. Alarmingly, studies of the subpolar AMOC suggested strong evidence of the weakening of the AMOC from the mid-1990s^{6,7}. In recent years, measurements and model studies revealed considerable variabilities in the AMOC on time scales ranging from daily to multi-decadal¹. There are multiple reasons for the variabilities: carbon emissions from human activities^{8,9} as well as the internal interactions within the climate systems and external forcing such as volcanic eruptions and solar radiation. A fairly recent modeling study suggested, shockingly, that the AMOC may be currently on the verge of a potential collapse¹⁰, which can cause a significant tipping phenomenon in the Earth's climate system¹¹.

Due to the global climate change and its tendency to continue to accelerate, the likelihood of AMOC collapse is increasing. However, due to the vast complexity of the non-linear dynamical system underlying the AMOC, long-term

prediction of the AMOC, i.e., to forecast when such a collapse might occur with certain level of confidence, is a daunting challenge. Even if a method is developed to predict that the collapse will occur in certain time period in the future, it would not provide specific criteria for devising countermeasures that should be implemented now to prevent the collapse in the future, particularly because the physical system underlying the AMOC is extremely complex and highly nonlinear, and subject to various stochastic forcing. Our point of view is that nowcasting - a term we use to coin short-term prediction of the dynamical behaviors of the AMOC, is also important and pertinent. In particular, we focus on the prediction horizon of one to two years and ask whether the detailed evolution of some key physical variables characterizing the AMOC can be accurately predicted based on the available observational data, modeling, and fingerprints (proxy records) of the AMOC. An advantage of nowcasting is that the presently available observations can be used to update the forecasting in a continuous manner, thereby guaranteeing the prediction accuracy. This is in fact an urgent problem, as successful and reliable nowcasting of the AMOC dynamics can provide a detailed and comprehensive picture of the AMOC evolution, e.g., whether it has deviated from the normal course. The ability to accurately assess the AMOC dynamics in the near future through reliable nowcasting can potentially lead to the discovery of critical early warning signals, based on which effective policy change and control strategies can be devised to reverse any harmful deviations of the dynamical evolution of the AMOC. The purpose of this paper is to present a machine-learning framework to demonstrate that accurate nowcasting of AMOC in a future time window between 12 and 24 months are feasible.

Past research demonstrated considerable variabilities in the AMOC, spanning spatially from the subtropics to the subpolar regions, and temporally across time from daily, intra-

^{a)}Electronic mail: Ying-Cheng.Lai@asu.edu

annual to interannual and decadal scales. The AMOC shows pronounced fluctuations across different timescales, and the variabilities can be as large as 100 % of its mean value on intra-annual and seasonal scales, whereas ranging between 10% – 30% on interannual to decadal scales¹². In addition, the AMOC behavior also varies in different regions, e.g., it is largely dictated by high-frequency wind forces spanning from sub-annual to interannual periods in the subtropical region. However, in the subpolar region, the variabilities primarily occur at lower frequencies on scales from interannual to decadal, where both wind and buoyancy forces come into play as significant influences^{13,14}.

The AMOC variabilities can be assessed through observations, models and reanalyses, and proxy records. In particular, observations or measurements provide a real-time window into the AMOC dynamics. For example, the RAPID (Rapid Climate Change) monitoring program¹⁵, established across the Atlantic at 26°N in 2004, uses an array of moorings to capture data related to the flow flux and temperature of the currents continuously. Based on the data, the AMOC is assessed by the combination of the strength of the ocean current through the strait, the near-surface Ekman transport from wind stress, and Gulf Stream transport from submarine cables¹⁶. Another program, The OSNAP (Overturning in the subpolar North Atlantic Program) focuses on the subpolar North Atlantic, aiming at capturing the intricate interactions among the currents, the atmosphere, and the cryosphere by stretching a line of observational instruments from Labrador to Scotland¹⁷. However, this program began in 2014, and due to its monthly timescale, there are currently not sufficient data points to evaluate the interannual or decadal variability. In addition to the RAPID and OSNAP programs, the MOVE (Meridional Overturning Variability Experiment) array, anchored in the subtropical North Atlantic near 16°N, provides crucial information about the AMOC variabilities in the upper to the deeper layers of the ocean. The positioning strategies in MOVE allow the dynamics of both warm, northward surface waters and the cold, southward deeper waters to be captured. While the observations from moorings are valuable datasets for analyzing the AMOC, there are issues such as missing data, noise, resolution and data point limitations. Complementing the real data analysis, models and reanalyses can be used to gain insights into the AMOC dynamics in terms of the structure of the ocean^{18–20}. Such models can simulate the dynamics of the ocean, atmosphere, and even the global climate system under different scenarios, which are particularly vital for analyzing and forecasting the AMOC changes in response to global warming or other large-scale disruptions. The third method is proxy records that serve as indirect indicators of the historical variabilities of the AMOC. These proxies, rooted in the mechanistic and statistical connections with AMOC, can leverage relationships discerned from models due to the scarcity of long-term observational data. One of the widely used proxies is derived from sea surface temperatures (SSTs) and subsurface temperatures^{10,21–23}.

Different available datasets measure physical quantities pertinent to the AMOC system on different timescales, e.g., intra-annual or interannual even decadal scales, in different

regions of the North Atlantic. As a result, nowcasting or short-term prediction does not mean that the prediction horizon would or should be the same for different physical quantities. While methods were proposed in the past to predict AMOC changes^{10,24,25}, the complex interplay of the physical variables and the vast datasets gathered demand a comprehensive computational framework to address the nowcasting problem. Our idea is to exploit machine learning that has proven to be powerful for analyzing experimental and observational datasets, uncovering hidden patterns, and making reliable predictions²⁶, e.g., in climate science²⁷. Our conviction is that a neural network can be trained with the historical data to learn the intricate dynamics of the AMOC and can then be used to forecast its future behavior, at least in short term. We exploit reservoir computing^{28–31}, a class of recurrent neural networks^{32,33}, which has been demonstrated recently being capable of accurate short-term prediction of the detailed dynamical evolution of nonlinear or even chaotic systems whose defining hallmark is sensitive dependence on initial conditions^{34–52}. Here, we use the aforementioned four datasets: MOVE and RAPID (observational), AMOC fingerprint (proxy records), AMOC simulated fingerprint and CESM AMOC (synthetic) to demonstrate accurate and reliable nowcasting of the AMOC dynamics in terms of the variabilities of the key underlying physical quantities within 12 to 24 months. We also address a number of issues affecting the prediction performance.

Our work has two unique features that go beyond the existing works. First, no prior studies have utilized a model-free approach to short-term prediction of the AMOC. As described, traditional methods are typically based on highly simplified, phenomenological models that may not adequately describe the measurements or observations. Our methodology leverages a model-free approach to predicting nonlinear dynamics using reservoir computing. The adaptability and versatility of this recurrent neural-network machine-learning architecture make it well-suited for predicting the AMOC dynamics with real or simulated data sets. Differing from the traditional methods, our method requires no prior knowledge about the dynamics, rendering it broadly applicable to complex dynamical systems in situations where a model is unknown or too complicated to be constructed. Second, we conducted an extensive study with a focus on short-term predictions of the AMOC using both synthetic and real-world datasets, demonstrating the robustness and effectiveness of the proposed method. The datasets studied include real measurements of the AMOC (MOVE and RAPID), data generated by phenomenological models (Simulated AMOC fingerprint data and CESM data), and proxy records (AMOC fingerprint), showcasing predictions of different aspects of the AMOC.

II. RESULTS

The AMOC characterizes the movement of the water in the North Atlantic, where warm surface water is transported northward into high northern latitudes and becomes cold and

returns southward through the deep ocean after releasing heat into the atmosphere and sinking. The AMOC can have profound impacts on regional and global climate patterns, including the temperature and precipitation in the western Europe, sea level along the east United States Coast, tropical monsoons, and Atlantic hurricane activity. To understand the intricacies of the AMOC and its impact on the climate, a variety of projects were set up to record and/or deduce the temperatures, salinity, transports, fluxes, and so on. The generated datasets represent the available information for studying the variability, trends, and the associated phenomena of the AMOC. At the present, a large number of AMOC datasets are available. We choose the following datasets: the AMOC fingerprint, MOVE data, and RAPID array measurements, for the reason that they are relatively complete and provide sufficiently long time series for machine learning. Figure 1 shows the various geographic positions from which the AMOC measurements were performed. In addition to the three datasets, we use a stochastic differential equation to simulate the AMOC fingerprint for validating our machine-learning model.




FIG. 1: Illustration of AMOC measurement positions. The color map displays the sea surface temperature (SST) in January 2023. Three measurement arrays in the Atlantic Ocean are shown in the map: AMOC fingerprint, RAPID, and MOVE, from top to bottom as the orange contour, purple and green line segments, respectively.

A. Data description and preprocessing

1. AMOC fingerprint

The AMOC has been continuously monitored since 2004 using a combination of tools⁵³, which seems to indicate a tendency for its strength to decline. However, to evaluate the trend and fluctuations, longer data sequences are necessary. Climate model simulations of the sea surface temperature (SST) suggested that, in the North Atlantic's Subpo-

lar gyre (SG) region, as illustrated by the orange outline in Fig. 1, SST can characterize the strength of the AMOC^{21–23} and therefore can act as an effective fingerprint of the AMOC. Figure 2(a) shows this preprocessed AMOC fingerprint for the period 1870–2020¹⁰, where the original data was from the Hadley Centre Sea Ice and Sea Surface Temperature data set (HADISST)⁵⁴. Here, the term “AMOC fingerprint” is defined¹⁰ as the SG anomaly minus twice the global mean anomaly so as to compensate for polar amplified global warming. More specifically, two factors need to be taken into account for compensating the SG anomaly: the seasonal cycle in the SST, which is governed by the surface radiation and is independent from the circulation, and the increasing trend in SST related to global warming. This requires that the global mean SST be taken away twice. The time interval between two adjacent data points in the AMOC fingerprint is one month. Altogether, the fingerprint dataset has 1812 points.




FIG. 2: AMOC time series data. (a) AMOC fingerprint. (b) A representative segment of the simulated AMOC fingerprint time series from Eq. (1). (c) Transport time series of the MOVE dataset in units of Sv ($10^6 \text{m}^3 \text{s}^{-1}$).

2. Simulated AMOC fingerprint data

While real datasets are critical for detecting and predicting the AMOC dynamical trend, often they have limited time periods of observations, rendering difficult obtaining statistically reliable results in some circumstances. In such a case, using simulated models to generate sufficient data can be useful for evaluating the capability of the proposed machine-learning prediction framework.

Assume that the dynamics underlying the AMOC fingerprint are in an equilibrium state (i.e., before a tipping point), the following one-dimensional stochastic differential equation (SDE) has been used in the literature¹⁰ as an empirical

model:

$$\frac{dX}{dt} = -A(X - m)^2 + \lambda + \xi, \quad (1)$$

where $\lambda < 0$ is a control or bifurcation parameter, which may cause the AMOC fingerprint to undergo a critical transition through a saddle-node bifurcation. The parameter m is defined as: $m \equiv \mu - |\lambda|/A$, where μ is the stable fixed point and A is a parameter defining the time scale of the system. The noise term ξ is a Gaussian stochastic process injected into the dynamics at each time step, which has zero mean and variance σ^2 . In Ref.¹⁰, the parameter values of the empirical SDE model were obtained through fitting with real data: $m = -1.5$, $A = 0.93$, $\sigma = 0.01$, and $\lambda = -2.8$. Figure 2(b) shows a representative segment of the simulated time series.

3. MOVE

At the present, the meridional overturning variability experiment (so-called MOVE) comprises two “geostrophic end-point moorings,” alongside a traditional current meter mooring on the slope. The primary objective of these installations was to gauge the transport variations across the section between the Lesser Antilles (Guadeloupe) and the Mid-Atlantic Ridge⁵⁵. The geostrophic transport can then be estimated according to the dynamic height and the bottom pressure variations between the moorings. Figure 2(c) shows a representative North Atlantic Deep Water (NADW) volume transport time series⁵⁶, derived from the pressure gradients observed at the MOVE section endpoints and the continental slope transport obtained from the current meters on the western MOVE moorings. The NADW time series exhibits significant interannual fluctuations that the southward (negative) transport undergoes. Furthermore, the decadal-scale variability can be observed unequivocally in the evolution of the time series. The processed time series started in January 2000 and had been accessed until May 2018, with the time interval of one day.

4. RAPID

The RAPID program aims to continuously obtain the strength and variability of the AMOC, enabling the relationship with the climate patterns and ocean carbon sink to be determined on the interannual and decadal time scales¹⁵. The RAPID array is located at 26°N, where the instruments are set across the Atlantic from Morocco to Florida, to measure the temperature, salinity, and current velocities from the surface to the floor of the ocean. By combining this array data with the observations from the Florida current and satellite-derived surface wind measurements, the overturning circulation can be calculated. Specifically, the AMOC data are obtained according to the combination of the velocity fields from three components: Florida Strait, Ekman, and density-driven transport^{16,57,58}, as shown in Fig. 3. Denoted as the Gulf Stream, the Florida Straits transport is calculated and calibrated as the




FIG. 3: Representative RAPID AMOC transport time series. Different colors indicate the AMOC (blue) and its components: the Florida Straits (green), the Ekman transport (orange), and the density-driven transport (purple). Transports are given in units of Sv ($10^6 \text{m}^3 \text{s}^{-1}$).

induced voltage, which can be used as a continuous indicator of the strength of the ocean current through the strait. Ekman transport is calculated from the wind stress acting on the ocean surface, which contributes to the largest variability in the AMOC. The density-driven transport denoted as Upper Mid-Ocean, calculates the current velocities by measuring the vertical profiles of the seawater density at a number of different longitudes. The data starts from April 2004 to February 2022, with a 12-hour time interval between two adjacent points.

5. CESM




FIG. 4: Representative CESM transport time series. (a) Linearly increasing freshwater flux with time. (b) Simulated AMOC with CESM.

The Community Earth System Model (CESM) was developed to simulate global climate system behaviors⁵⁹. It integrates various components of the Earth’s climate, including

the atmosphere, ocean, land, and sea ice to simulate complex interactions within the system. During the simulation, greenhouse gas concentrations, solar radiation, and aerosol levels are kept constant, while the primary forcing mechanism is the freshwater flux anomaly in the North Atlantic, between latitudes 20°N and 50°N, for generating the AMOC data. Specifically, the anomaly is a slowly linearly increasing function of time at the rate of 3×10^{-4} Sv per model year, as shown in Fig. 4(a), which describes the effects of the freshwater input from melting ice and increased precipitation. The generated data reveal variations in the AMOC strength under increasing freshwater forcing. Natural variability dominates the AMOC strength in the first 400 points, while a negative trend appears after $t = 800$ due to the increasing freshwater forcing. Figure 4(b) presents a representative segment of the simulated AMOC data from the CESM.

6. Data preprocessing

The sampling rates of the different available datasets vary, e.g., from 12 hours to one year. Moreover, the AMOC fingerprint, its simulated time series, the MOVE data, and the CESM data are one-dimensional, but the RAPID data are four-dimensional. For all the simulated and real datasets, we use min-max normalization to preprocess the time series so that they are in the range $[0,1]$. For each dataset, we divide it into three segments for training (50%), validation (25%), and testing or prediction (25%), respectively. In particular, for training, both the input and target output are provided to determine the pertinent neural-network weights - all elements of the output matrix (open-loop reservoir computing). During the validation, input is still present but the true output is absent (still open-loop operation but with a fixed output matrix). For testing (prediction), the output of the reservoir network is connected to the input to execute closed-loop operation. We perform multi-step predictions: at each time step, we predict T_p steps ahead. The following measure of root-mean-square errors is used to characterize the validation and testing performances:

$$\text{RMSE}(y, \hat{y}) = \sqrt{\frac{1}{T_p} \sum_{t=i}^{T_p+i} [y(t) - \hat{y}(t)]^2} \quad (2)$$

where $y(t)$ and $\hat{y}(t)$ are the real and predicted signals, respectively. RMSEs can be obtained by taking the average over the whole validation or testing length (time interval).

B. Results of recursive prediction

The training of the reservoir neural network was done on two computers with RTX 4000 NVIDIA GPU using Python. There are two commonly used prediction methods: direct and recursive, both generating multiple-step predictions. In particular, let T_p be the prediction horizon, i.e., the number of time steps of forward prediction. For direct prediction, at

each time step, the trained reservoir computer generates T_p consecutive data points at once. The input is then updated using the real data, to prepare for the prediction at the next time step, and so on. For recursive prediction, at each time step, the output of the reservoir computer is fed back to the input, generating a closed-loop dynamical system capable of self evolution. The system evolves forward for T_p time steps, generating a T_p -step prediction. We present the results from recursive prediction here, while reporting these from direct prediction in Appendix. With respect to the four available datasets, the nowcasting results from the AMOC fingerprint, MOVE, and the CESM data are presented here, while those from simulated AMOC fingerprint and RAPID can be found in Appendix.

1. Recursive prediction of AMOC fingerprint data




FIG. 5: Recursive prediction of AMOC fingerprint data. (a) The ground truth (blue), trained (orange) and predicted (green) AMOC fingerprint with the prediction horizon $T_p = 24$ months. (b) Training (orange) and prediction (green) errors. (c) Distribution of the prediction error.

For nowcasting of the AMOC fingerprint, we use reservoir neural networks of size $n = 500$. The hyperparameter val-

TABLE I: Optimal hyperparameter values for prediction with the four AMOC datasets

Dataset	a	γ	ρ	p	w_b	β
Fingerprint	0.17	1.32	0.45	0.19	3.04	-5.30
Simulated fingerprint	0.05	2.35	0.05	0.65	0.32	-2.30
MOVE	0.95	0.67	1.74	0.50	1.65	-1.66
RAPID	1.0	0.01	2.27	1.0	1.53	-3.84
CESM	0.14	0.28	3.73	0.62	3.62	-1.70




FIG. 6: Recursive prediction of AMOC fingerprint with different prediction horizons. (a, c, e) Reservoir-computing predicted fingerprint and the ground truth, for prediction horizons $T_p = 12, 24,$ and 36 months, respectively. (b, d, f) The corresponding prediction errors. (g) Ensemble-averaged testing RMSE versus the prediction horizon T_p , where the error bars represent the standard deviation calculated from 50 independent realizations.

ues obtained through Bayesian optimization are displayed in Tab. I. For the available AMOC fingerprint time series that has 1812 monthly data points, we use 900 points for training and the 300 points afterward for validation to obtain the hyperparameter values through Bayesian optimization. The testing or prediction phase corresponds to the time period associated with validation segment and the remaining 612 data points (912 points in total), which are used to calculate the RMSEs.

Figure 5 shows the nowcasting results with the prediction horizon $T_p = 24$ months, where the ground truth as well as the trained and predicted data are presented in Fig. 5(a), the corresponding RMSEs are shown in Fig. 5(b), and the distribution of the RMSEs in the testing phase is demonstrated in Fig. 5(c). Note that, for better visualization, the displayed prediction result in Fig. 5(a) is presented as the moving-window average of window size $T_p = 24$ time steps. It can be seen that the prediction result is faithful in that its moving-window average is fully embedded in the ground truth. In fact, the RMSE is relatively small: about 0.14. For com-




FIG. 7: Recursive prediction of AMOC fingerprint data with 95% confidence interval. The ground truth (blue), trained (orange) and predicted (green) AMOC fingerprint with the prediction horizon $T_p = 24$ months. The green shaded area denotes the 95% confidence interval calculated from 50 independent sets of optimized hyperparameters.

parison, results for three different values of the prediction horizon are shown in Figs. 6(a-c) for $T_p = 12, 24,$ and 36 months, respectively, with the corresponding RMSEs shown in Figs. 6(d-f). The average RMSEs for the three cases are 0.13, 0.14, and 0.16, respectively. In general, as the prediction horizon increases, the prediction deteriorates, as exemplified in Fig. 6(g) - the result of a systematic computation of the RMSE versus the prediction horizon where, for each value of the prediction horizon, the average RMSE value over 50 realizations is obtained by randomizing the training length between 880 and 900.

To assess the reliability of the prediction, we incorporate uncertainty quantification. This entails training multiple models, each with independently optimized sets of hyperparameters, and aggregating the predictions to estimate the confidence intervals. Specifically, we optimize the hyperparameters with Bayesian optimization 50 times, generate an ensemble of predictions, and calculate the 95% prediction interval by approximately 1.96 standard deviations from the mean value of these predictions at each time step. Figure 7 illustrates the recursive prediction of the AMOC fingerprint data along with the 95% confidence interval, denoted by the green shaded area. It can be seen that the confidence interval effectively captures the trend of the fingerprint data, demonstrating the reliability of the predictions.

2. Recursive prediction of AMOC MOVE data

For nowcasting of the MOVE data, the hyperparameter values obtained through Bayesian optimization during validation are listed in Tab. I. The MOVE time series has 6710 points (i.e., 6710 days). The segments of the date for training and validation consist of 3000 and 1000 points, respectively. The validation segment together with the remaining (altogether 3710 points) is for testing. Figure 8(a) shows the trained




FIG. 8: Recursive prediction of MOVE dataset. (a) The ground truth (blue), training (orange) and predicted (green) MOVE dataset with the prediction horizon $T_p = 40$ days. (b) Training (orange) and prediction (green) errors. (c) Distribution of the prediction error.

and predicted time series, together with the ground truth, for the prediction horizon $T_p = 40$ days. The corresponding RMSEs and their distribution are displayed in Figs. 8(b) and 8(c), respectively, with the ensemble-averaged RMSE of about 0.12. Figures 9(a-c) show the predicted MOVE time series for $T_p = 20, 40$, and 100 days, respectively, with the corresponding RMSEs shown in Figs. 9(d-f). Figure 9(g) shows the ensemble-averaged RMSE (together with the error bars) versus T_p . As T_p increases, the RMSE increases, rendering infeasible long-term prediction. Figure 10 shows the recursive prediction of the MOVE data along with the 95% confidence interval, as shaded by the green area. The confidence interval effectively captures the trend of the fingerprint data, demonstrating the reliability of the predictions.

3. Recursive prediction of CESM data

For nowcasting of the CESM data, the hyperparameter values obtained through Bayesian optimization during validation are listed in Tab. I. The CESM time series contains 1700 points. The segments of the data for training and validation consist of 1000 and 300 points, respectively. The validation segment together with the remaining (altogether 700 points) is for testing. Figure 11(a) shows the trained and predicted time series, together with the ground truth, for the prediction horizon $T_p = 24$ days. The corresponding RMSEs and their distribution are displayed in Figs. 11(b) and 11(c), respectively, with the ensemble-averaged RMSE of about 0.06. Figures 12(a-c) show the predicted CESM time series for $T_p = 12, 24$, and 36 days, respectively, with the corresponding RMSEs shown in Figs. 12(d-f). Figure 12(g) shows the ensemble-averaged RMSE (together with the error bars) versus T_p . As T_p increases, the RMSE increases, rendering in-




FIG. 9: Recursive prediction MOVE dataset with varying prediction horizons. (a, c, e) Reservoir-computing predicted transport together with the ground truth for the prediction horizons $T_p = 20, 40$, and 100 days, respectively. (b, d, f) The corresponding RMSEs. (g) Ensemble-averaged testing RMSE versus T_p , with error bars being the standard deviation calculated from 50 independent simulations.




FIG. 10: Recursive prediction of MOVE data with 95% confidence interval. The ground truth (blue), trained (orange) and predicted (green) AMOC fingerprint with the prediction horizon $T_p = 40$ months. The green shaded area denotes the 95% confidence interval calculated from 50 independent sets of optimized hyperparameters.




FIG. 11: Recursive prediction of CESM data. (a) The ground truth (blue), training (orange) and predicted (green) CESM dataset with the prediction horizon $T_p = 24$ points. (b) Training (orange) and prediction (green) errors. (c) Distribution of the prediction error.

feasible long-term prediction.

To characterize the uncertainties in the prediction results, we optimize the hyperparameters with Bayesian optimization 50 times, generate an ensemble of predictions, and calculate the 95% prediction interval by approximately 1.96 standard deviations from the mean value at each time step. Figure 13 illustrates the recursive prediction of the CESM data along with the 95% confidence interval, denoted by the green shaded area. It can be seen that the confidence interval effectively captures the trend of the fingerprint data, demonstrating the reliability of the short-term prediction results.

III. DISCUSSION

In summary, we articulated a reservoir-computing based, model-free framework to predict the real-time dynamical evolution of some key physical variables characterizing the AMOC and demonstrated the feasibility of accurate short-term prediction, using five different types of datasets: MOVE and RAPID (observational), AMOC fingerprint (proxy records), AMOC simulated fingerprint and CESM data (synthetic model). The two observational datasets and the proxy records are collected from different regions and describe variability on different timescales, even with distinct physical properties. Quantitatively, we found a monotonic increase of the prediction error with the horizon where the error tends to increase relatively more rapidly as the horizon increases, e.g., from 12 to 36 months. Empirically, we conclude that short-term prediction within 24 months can be achieved, while any longer horizon would deem the prediction inaccurate and unreliable. Our machine-learning framework thus will not be able to address the question of any possible AMOC collapse in the future. Rather,




FIG. 12: Recursive prediction CESM data with varying prediction horizons. (a,c,e) Reservoir-computing predicted transport together with the ground truth for the prediction horizons $T_p = 12, 24$, and 36 points, respectively. (b,d,f) The corresponding RMSEs. (g) Ensemble-averaged testing RMSE versus T_p , with error bars being the standard deviation calculated from 50 independent simulations.




FIG. 13: Recursive prediction of CESM data with 95% confidence interval. The ground truth (blue), trained (orange) and predicted (green) AMOC fingerprint with the prediction horizon $T_p = 24$ months. The green shaded area denotes the 95% confidence interval calculated from 50 independent sets of optimized hyperparameters.

the significance of the demonstrated achievable nowcasting of some key physical variables associated with the AMOC lies in the capability to detect any unusual trend or early warning for about 24 months ahead of time. This can provide a window for devising control/mitigation strategies through policy making to reverse any adverse trend of the AMOC.

Our general point of view is that long-term prediction of a possible collapse of the AMOC with the currently available data will remain to be a formidable challenge. While mathematical models of the AMOC with parameter values extracted from data suggested different collapse scenarios, including a fairly recent one predicting a potential collapse around the middle of the century based on one-dimensional stochastic differential equations¹⁰, the predictions are only suggestive. In view of the grave consequence of AMOC weakening and a potential collapse, even the speculations cannot be afforded to be ignored. The dynamical system underlying the AMOC is vastly complex and highly nonlinear with sensitivity on initial conditions, so predicting the state evolution of this dynamical system in the long run is fundamentally ruled out. It may be argued, however, that an AMOC collapse is a phenomenon of critical or tipping transition, thereby requiring no detailed knowledge about the state evolution.

At the present, there are two main approaches to predicting future critical transitions in nonlinear dynamical systems. The first is based on finding the system equations from data^{60,61}. A more recent approach^{62–71} was based on sparse optimization. If the accurate equations governing the underlying system can be found, when a bifurcation parameter changes with time, a possible collapse of the system can be assessed. However, the applicability of the sparse optimization approach is limited to systems with a simple or “sparse” equation structure^{62,63}. The second approach is based on machine learning, and it has been demonstrated that reservoir computing can be exploited to predict critical transitions^{42,45,48}. An advantage of this approach is that, in principle, it can be applied to any system, regardless of the underlying mathematical structure of the governing equations. The disadvantage is that the amount of data required for training can be quite demanding.

The focus of this work is on the short-term prediction of the AMOC using machine learning. Given the importance of the AMOC to the global climate, predicting the AMOC’s behavior even in short term is important and of considerable interest. However, the complexity of the dynamical system underlying the AMOC makes accurate predictions challenging. In response to the challenges, we propose a model-free approach to short-term AMOC prediction. The proposed method aims to estimate the AMOC dynamics effectively into the near future, potentially providing early warning signals that can inform policy-making and control strategies. With respect to the inherent uncertainties of any long-term predictions, short-term predictions can be more accurate and reliable to provide actionable insights.

A significant challenge in AMOC research is data insufficiency and inconsistency, particularly with observational records such as those from RAPID or OSNAP. Machine-learning methods generally require large amounts of high-

quality data for effective training, creating a substantial hurdle given the limited and often inconsistent AMOC data. Proxy records such as those derived from ocean sediments or sea ice may provide historical data to benefit prediction. Short-term AMOC predictions may also be improved by overcoming data insufficiency through the method of transfer learning. The main idea involves utilizing synthetic data generated from mathematical models to initially train the neural networks and then, fine-tune the well-trained model using a limited number of real-world measurements. By so doing, the model may be capable of highly accurate prediction of the short-term dynamics of the AMOC, even with limited measurement data.

Long-term prediction of the AMOC is currently a significant challenge, particularly concerning the probability of its collapse — a critical or tipping transition event. In nonlinear and complex dynamics, at the present there are two possible approaches to predicting future critical transitions: finding the governing equations and machine learning. The first approach is effective if the underlying system equations meet the sparsity requirement but it is impractical for the AMOC. The machine-learning approach, while promising, faces its own challenges. For example, for long-term prediction, machine-learning models such as reservoir computing require knowledge about some key system parameters that vary over time. To determine such time-varying parameters for the AMOC from data is currently an open problem. A plausible speculation is that global climate change may be the main driving force behind a potential AMOC collapse, but the available climate data pertinent to the AMOC are quite limited (within the past 100 years or so) and are far from sufficient for extrapolating how some key parameters may change into the future. Although machine-learning models can be trained with limited measurement data, their effectiveness is hampered by the lack of detailed knowledge about how these parameters vary over time. Without this information, these models cannot reliably predict a potential future collapse of the AMOC.

IV. METHOD

A. Reservoir computing

The core of reservoir computing is a recurrent neural network (RNN) with a non-Markovian type of dynamics in a single hidden layer. Differing from the conventional RNNs where the network link weights are trained by a gradient-based method, in reservoir computing the input matrix elements and the internal weights of the RNN are randomly initialized and remain fixed during training: all required to be trained is the weights of the output matrix that can be determined by a standard linear regression. This unique character makes the training process computationally efficient, while capturing the information or dynamical climate of the target dynamical system and embedding it into the dynamics of the hidden-layer neural network^{28–31,72}.

Figure 14 shows the architecture of reservoir computing.




FIG. 14: Architecture of reservoir computing. A reservoir computer has three layers: the input, hidden, and the output layers, characterized by the matrices W_{in} , A , and W_{out} , respectively. The input, hidden state, and the output vectors are $\mathbf{u}(t)$, $\mathbf{r}(t)$ and $\mathbf{o}(t)$, respectively.

The low-dimensional input signal $\mathbf{u}(t)$ is mapped by the input matrix W_{in} into the high-dimensional internal state vector $\mathbf{r}(t)$ in the hidden layer, which is updated step by step according to the following rule:

$$\mathbf{r}(t+1) = (1 - \alpha) \cdot \mathbf{r}(t) + \alpha \cdot \tanh [A \cdot \mathbf{r}(t) + W_{in} \cdot \mathbf{u}(t) + W_{bias}], \quad (3)$$

where α is the leakage parameter that determines the memory loss after each time step, the activation is achieved through a hyperbolic tangent function ($\tanh$), the adjacency matrix A gives the links and connection structure of the reservoir network, and the components of the bias vector W_{bias} are an equal constant w_b , whose role is to shift the values inside the activation function from zero. The input signal of whole training length $[I(1), I(2), \dots, I(t)]$ is concatenated into a matrix U and the recorded reservoir dynamical states $[\mathbf{r}(2), \mathbf{r}(3), \dots, \mathbf{r}(t+1)]$ can be concatenated into a matrix R . The output matrix is determined by using Tikhonov regularization⁷³ as

$$W_{out} = U \cdot R^\tau \cdot (R^\tau \cdot R^\tau + \beta I)^{-1}, \quad (4)$$

where $R' = \text{fs}(R)$ is modified from R such that the elements in all its even rows are squared to minimize overfitting³⁴, R^τ is the transpose of R , β is the regularization coefficient, and I is the identity matrix. The output is

$$\mathbf{o}(t) = W_{out} \cdot \mathbf{r}(t). \quad (5)$$

The reservoir-computing architecture possesses a small number of hyperparameters: the leakage parameter α , the scaling factor γ of the input matrix, the spectral radius ρ and link probability p of the reservoir network, the bias constant w_b of the bias matrix, and the regularization coefficient β . The values of these hyperparameters can have a significant effect on the performance, so it is necessary to find their optimal values, which is done commonly through Bayesian optimization by using⁷⁴, e.g., the algorithm from Python (*bayesian-optimization*)⁷⁵.

DATA AND CODE AVAILABILITY

Data and codes are available from GitHub: <https://github.com/Zheng-Meng/AMOC>

ACKNOWLEDGMENTS

This work was supported by the Air Force Office of Scientific Research (AFOSR) under Grant No. FA9550-21-1-0438.

AUTHOR CONTRIBUTIONS

All designed the research project, the models, and methods. Z.-M. Z. performed the computations. All analyzed the data. Z.-M.Z. and Y.-C.L. wrote the paper. M.M. and Y.-C.L. edited the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

CORRESPONDENCE

To whom correspondence should be addressed: Ying-Cheng.Lai@asu.edu

Appendix A: Simulated AMOC fingerprint

The simulated AMOC fingerprint dataset has 10^5 data points. For nowcasting with this dataset, we set the size of the reservoir network to be $n = 500$. The hyperparameter values obtained through Bayesian optimization are listed in Tab. I in the main text. The training length is set to be 5,000, the length of the validation time series (after the training data) is 3,000, and the testing data length is 5,600 (including the validation data).

Figure 15 shows the nowcasting results with the prediction horizon $T_p = 40$, where Fig. 15(a) displays the ground truth as well as the trained and predicted time series. The corresponding training and testing RMSEs are shown in Fig. 15(b), and the distribution of the testing RMSEs is presented in Fig. 15(c). The ensemble-averaged RMSE is about 0.10. Three additional examples are shown in Fig. 16 for $T_p = 20, 40$, and 100, respectively, where the left column shows the predicted AMOC fingerprint time series and the ground truth with the corresponding RMSEs shown in the right row. The averaged RMSE for the three cases are 0.08, 0.1, and 0.11, respectively. Figure 16(g) shows the testing RMSE versus the prediction horizon, where the mean and the standard deviation are obtained through an ensemble of simulated AMOC fingerprint data. As the prediction horizon increases, the RMSE increases.

Appendix B: RAPID

The total length of the RAPID time series is 12183 (about 6000 days). The training, validation, and testing (including validation) data segments have 6000, 3000, and 6183 points, respectively. For nowcasting with the RAPIC dataset, we set the size of the reservoir network to be $n = 500$. The hyperparameter values obtained through Bayesian optimization are listed in Tab. I in the main text. The RAPID dataset is four-dimensional, so the input and output of the reservoir computer are four-dimensional.

Figure 17 shows the nowcasting results for the RAPID dataset, where the prediction horizon is $T_p = 40$ (20 days). In particular, Figs. 17(a,c,e,g) display the trained and predicted time series as well as the ground truth for MOC, Upper Mid-Ocean, Gulf Stream, and Ekman, respectively, with the corresponding RMSEs shown in Figs. 17(b,d,f,h). The ensemble-averaged RMSEs associated with the four dimensions are about 0.09, 0.10, 0.12, and 0.09, respectively, suggesting accurate short-term prediction. Several examples of predicted RAPID time series with different prediction horizons are shown in Fig. 18, for $T_p = 20, 40$, and 70 (10, 20, and 70 days), respectively. Each column of Fig. 18 represents the prediction results and the corresponding errors. The averaged RMSEs for MOC for $T_p = 20, 40, 70$ are 0.06, 0.09, and 0.10, respectively. The averaged RMSEs for the Upper Mid-Ocean for $T_p = 20, 40, 70$ are 0.08, 0.10, and 0.13, respectively. The averaged RMSEs for the Gulf Stream for $T_p = 20, 40, 70$ are 0.09, 0.12, and 0.14, respectively. The averaged RMSEs for the Ekman for $T_p = 20, 40, 70$ are 0.07, 0.09, and 0.09, respectively. Figure 19 shows the RMSE versus the prediction horizon, where a shorter testing length of 5000 points is used due to the large values of T_p in the range. In addition, to validate the robustness of the prediction results, the training length was randomly set between 5900 and 6000. The RMSE is small for short prediction horizon, but it increases with the prediction horizon.

Appendix C: Comparison of two machine-learning prediction methods

There are two common methods in time-series multi-step forecasting with machine learning: recursive and direct prediction, both with their own unique advantages and constraints. The training process of the two methods is the same, and the difference lies in the prediction phase. In particular, as shown in Fig. 20(a), the recursive method, also known as the closed-loop iteration method, predicts one step ahead at each time step and then feeds the predicted value back into the machine-learning model as the input for the next time step. This procedure is iterated until a desired number of steps, i.e., the prediction horizon T_p , is reached. During the prediction, the hidden state of the reservoir network is recorded since the current and previous real data are needed at the beginning of each prediction step. For example, at time step t , the time series information and the hidden state $\mathbf{r}$ at $0, 1, \dots, t$ are needed to predict the next T_p time steps iteratively. The

state vector $\mathbf{r}$ is updated until $\mathbf{r}_{t+T_p}$ time steps are reached. To make the prediction at time step $t + 1$, the hidden state $\mathbf{r}_{t+1}$ is needed (not $\mathbf{r}_{t+T_p}$). The recursive prediction method has the advantage of train easiness and flexibility, but error accumulation during the iteration is an issue. In comparison, in direct prediction, the reservoir machine is trained to predict T_p time steps ahead, without the iterative feedback loop employed in recursive prediction. In the testing phase, the direct method predicts T_p at once, as illustrated in Fig. 20(b). Direct prediction has the advantage of reduced error propagation, but it is susceptible to overfitting, especially with limited data.

For the four AMOC datasets, while we have presented the prediction results from the recursive prediction method, the direct prediction method performs better for two datasets. Figures 21(a) and 21(b) show the ensemble-averaged testing RMSE versus the prediction horizon T_p of the AMOC fingerprint and MOVE dataset, respectively. For these two datasets, the recursive method outperforms the direct method. The corresponding results for the simulated AMOC fingerprint and RAPID datasets are shown in Figs. 21(c) and 21(d), respectively. For these two datasets, the direct method outperforms the recursive method.

Appendix D: Effects of low-pass filtering

An issue with the datasets is noise, which can be removed to certain extent by low-pass filtering. Will partial noise removal affect the machine-learning performance of short-term prediction of the AMOC? To address this question, we employ a moving average method by using a rolling window of size L_w to smooth the training data segment. The validation and testing data remain in their raw, unfiltered state. Representative results are shown in Fig. 22. It can be seen that, for the four datasets, low-pass filtering does not lead to any apparent performance improvement. In fact, the RMSEs tend to increase with the window width L_w . A plausible explanation is that the noise-filtering procedure distorted the similarity between the training data and the real noisy data, leading to information loss and making it harder for the machine to learn the trend of the AMOC.




FIG. 15: Nowcasting of simulated AMOC fingerprint data. (a) The ground truth (blue), trained (orange) and predicted (green) time series with the prediction horizon $T_p = 40$. (b) RMSEs with the trained (orange) and predicted (green) data. (c) Distribution of the predicted error.




FIG. 16: Examples of nowcasting of simulated AMOC fingerprint with different prediction horizons. (a, c, e) Reservoir-computing predicted fingerprint together with the ground truth, for prediction horizons $T_p = 20$, $T_p = 40$, $T_p = 100$, respectively. (b, d, f) The corresponding prediction errors. (g) Ensemble-averaged testing RMSE versus the prediction horizon T_p , with the standard deviation calculated from 50 independent simulations.




FIG. 17: Nowcasting with RAPID dataset. (a, c, e, g) The ground truth (blue), trained (orange) and predicted (green) MOC, Upper Mid-Ocean, Gulf Stream, and Ekman, respectively, with the prediction horizon $T_p = 40$ (i.e., 20 days). (b, d, f, h) The corresponding trained (orange) and predicted (green) errors.




FIG. 18: Nowcasting of RAPID dataset with different prediction horizons. The three columns correspond to prediction horizon $T_p = 20$, $T_p = 40$, and $T_p = 100$ (10, 20, and 70 days), respectively. The first, third, fifth, and seventh rows give the four predicted quantities associated with the RAPID data and the ground truth, with the corresponding RMSEs shown in the second, fourth, sixth, and eighth rows.

This is the author's peer-reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.




FIG. 19: Quantifying short-term prediction of the RAPID dataset through the RMSE. Shown is the ensemble-averaged testing RMSE versus the prediction horizon T_p , with the standard deviations calculated from 50 independent realizations.




FIG. 20: Illustration of two short-term prediction methods: recursive and direct prediction.




FIG. 21: Errors with recursive prediction and direct prediction. (a-d) Ensemble-averaged testing RMSE versus the prediction horizon T_p for AMOC fingerprint, MOVE, simulated AMOC fingerprint, and RAPID dataset, respectively. The error bars are obtained from 50 independent realizations.




FIG. 22: Effect of low-pass filtering on the performance of short-term prediction. (a-d) Ensemble-averaged testing RMSE versus the moving-window size L_w for AMOC fingerprint, MOVE, simulated AMOC fingerprint, and RAPID dataset, respectively. Data-smoothing or low-pass filtering does not lead to any apparent performance improvement.

- ¹L. 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**, 241–254 (2022).
- ²M. W. Buckley and J. Marshall, "Observations, inferences, and mechanisms of the Atlantic Meridional Overturning Circulation: A review," *Rev. Geophys.* **54**, 5–63 (2016).
- ³K. 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**, 4567–4583 (2019).
- ⁴K. Bellomo, M. Angeloni, S. Corti, and J. von Hardenberg, "Future climate change shaped by inter-model differences in Atlantic meridional overturning circulation response," *Nat. Commun.* **12**, 3659 (2021).
- ⁵J. L. Sarmiento and C. Le Quere, "Oceanic carbon dioxide uptake in a model of century-scale global warming," *Science* **274**, 1346–1350 (1996).
- ⁶A. Biastoch, C. W. Boëning, 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**, 6599–6615 (2008).
- ⁷S. Yeager and G. Danabasoglu, "The origins of late-twentieth-century variations in the large-scale North Atlantic circulation," *J. Clim.* **27**, 3222–3247 (2014).
- ⁸M. B. Menary, J. Robson, R. P. Allan, B. B. Booth, C. Cassou, G. Gastineau, J. Gregory, D. Hodson, C. Jones, J. Mignot, *et al.*, "Aerosol-forced AMOC changes in CMIP6 historical simulations," *Geophys. Res. Lett.* **47**, e2020GL088166 (2020).
- ⁹M. Klockmann, U. Mikolajewicz, and J. Marotzke, "The effect of greenhouse gas concentrations and ice sheets on the glacial AMOC in a coupled climate model," *Clim. Past* **12**, 1829–1846 (2016).
- ¹⁰P. Ditlevsen and S. Ditlevsen, "Warning of a forthcoming collapse of the Atlantic meridional overturning circulation," *Nat. Commun.* **14**, 4254 (2023).
- ¹¹T. 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. U.S.A.* **105**, 1786–1793 (2008).
- ¹²B. I. Moat, D. A. Smeed, E. Frajka-Williams, D. G. Desbruyères, C. Beaulieu, W. E. Johns, D. Rayner, A. Sanchez-Franks, M. O. Baringer, D. Volkov, *et al.*, "Pending recovery in the strength of the meridional overturning circulation at 26 N," *Ocean Sci.* **16**, 863–874 (2020).
- ¹³H. R. Pillar, P. Heimbach, H. L. Johnson, and D. P. Marshall, "Dynamical attribution of recent variability in Atlantic overturning," *J. Clim.* **29**, 3339–3352 (2016).
- ¹⁴S. M. Larson, M. W. Buckley, and A. C. Clement, "Extracting the buoyancy-driven Atlantic meridional overturning circulation," *J. Clim.* **33**, 4697–4714 (2020).
- ¹⁵E. Frajka-Williams, I. J. Ansorge, J. Baehr, H. L. Bryden, M. P. Chidichimo, S. A. Cunningham, G. Danabasoglu, S. Dong, K. A. Donohue, S. Elipot, *et al.*, "Atlantic meridional overturning circulation: Observed transport and variability," *Front. Mar. Sci.* **6**, 260 (2019).
- ¹⁶G. D. McCarthy, D. A. Smeed, W. E. Johns, E. Frajka-Williams, B. I. Moat, D. Rayner, M. O. Baringer, C. S. Meinen, J. Collins, and H. L. Bryden, "Measuring the Atlantic meridional overturning circulation at 26 N," *Prog. Oceanogr.* **130**, 91–111 (2015).
- ¹⁷M. S. Lozier, F. Li, S. Bacon, F. Bahr, A. S. Bower, S. Cunningham, M. F. de Jong, L. de Steur, B. deYoung, J. Fischer, *et al.*, "A sea change in our view of overturning in the subpolar North Atlantic," *Science* **363**, 516–521 (2019).
- ¹⁸W. Large and S. Yeager, "The global climatology of an interannually varying air–sea flux data set," *Clim. Dyn.* **33**, 341–364 (2009).
- ¹⁹A. Storto, A. Alvera-Azca, M. A. Balmaseda, A. Barth, M. Chevallier, F. Counillon, C. M. Domingues, M. Drevillon, Y. Drillet, G. Forget, *et al.*, "Ocean reanalyses: Recent advances and unsolved challenges," *Front. Mar. Sci.* **6**, 418 (2019).
- ²⁰B. Fox-Kemper, A. Adcroft, C. W. Boëning, E. P. Chassignet, E. Curchitser, G. Danabasoglu, C. Eden, M. H. England, R. Gerdes, R. J. Greatbatch, *et al.*, "Challenges and prospects in ocean circulation models," *Front. Mar. Sci.* **6**, 65 (2019).
- ²¹L. Caesar, S. Rahmstorf, A. Robinson, G. Feulner, and V. Saba, "Observed fingerprint of a weakening Atlantic Ocean overturning circulation," *Nature* **556**, 191–196 (2018).
- ²²L. Jackson and R. Wood, "Fingerprints for early detection of changes in the AMOC," *J. Clim.* **33**, 7027–7044 (2020).
- ²³M. Latif, E. Roeckner, M. Botzet, M. Esch, H. Haak, S. Hagemann, J. Jungclaus, S. Legutke, S. Marsland, U. Mikolajewicz, *et al.*, "Reconstructing, monitoring, and predicting multidecadal-scale changes in the North Atlantic thermohaline circulation with sea surface temperature," *J. Clim.* **17**, 1605–1614 (2004).
- ²⁴S. Mahajan, R. Zhang, T. L. Delworth, S. Zhang, A. J. Rosati, and Y.-S. Chang, "Predicting Atlantic meridional overturning circulation (AMOC) variations using subsurface and surface fingerprints," *Deep-Sea Res. II: Top. Stud. Oceanogr.* **58**, 1895–1903 (2011).
- ²⁵W. Weijer, W. Cheng, O. A. Garuba, A. Hu, and B. Nadiga, "CMIP6 models predict significant 21st century decline of the Atlantic meridional overturning circulation," *Geophys. Res. Lett.* **47**, e2019GL086075 (2020).
- ²⁶M. Moradi, Z.-M. Zhai, A. Nielsen, and Y.-C. Lai, "Random forests for detecting weak signals and extracting physical information: A case study of magnetic navigation," *APL mach. learn.* **2**, 016118 (2024).
- ²⁷K. Kashinath, M. Mustafa, A. Albert, J. Wu, C. Jiang, S. Esmailzadeh, K. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh, *et al.*, "Physics-informed machine learning: Case studies for weather and climate modelling," *Philos. Trans. R. Soc. A* **379**, 20200093 (2021).
- ²⁸H. Jaeger, "The "echo state" approach to analysing and training recurrent neural networks—with an erratum note," Bonn, Germany: German National Research Center for Information Technology GMD Technical Report **148**, 13 (2001).
- ²⁹W. Maass, T. Natschläger, and H. Markram, "Real-time computing without stable states: A new framework for neural computation based on perturbations," *Neu. Comp.* **14**, 2531–2560 (2002).
- ³⁰M. Lukoševičius and H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Compu. Sci. Rev.* **3**, 127–149 (2009).
- ³¹L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, "Information processing using a single dynamical node as complex system," *Nat. Commun.* **2**, 1–6 (2011).
- ³²T. Robinson, M. Hochberg, and S. Renals, "The use of recurrent neural networks in continuous speech recognition," in *Automatic Speech and Speaker Recognition: Advanced Topics* (Springer, 1996) pp. 233–258.
- ³³M. Hermans and B. Schrauwen, "Memory in linear recurrent neural networks in continuous time," *Neural Netw.* **23**, 341–355 (2010).
- ³⁴Z. Lu, J. Pathak, B. Hunt, M. Girvan, R. Brockett, and E. Ott, "Reservoir observers: Model-free inference of unmeasured variables in chaotic systems," *Chaos* **27**, 041102 (2017).
- ³⁵J. Pathak, Z. Lu, B. Hunt, M. Girvan, and E. Ott, "Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data," *Chaos* **27**, 121102 (2017).
- ³⁶J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, "Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach," *Phys. Rev. Lett.* **120**, 024102 (2018).
- ³⁷G. Tanaka, T. Yamane, J. B. He'roux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, "Recent advances in physical reservoir computing: A review," *Neu. Net.* **115**, 100–123 (2019).
- ³⁸J. Jiang and Y.-C. Lai, "Model-free prediction of spatiotemporal dynamical systems with recurrent neural networks: Role of network spectral radius," *Phys. Rev. Res.* **1**, 033056 (2019).
- ³⁹H. Fan, J. Jiang, C. Zhang, X. Wang, and Y.-C. Lai, "Long-term prediction of chaotic systems with machine learning," *Phys. Rev. Res.* **2**, 012080 (2020).
- ⁴⁰E. Bollt, "On explaining the surprising success of reservoir computing forecaster of chaos? The universal machine learning dynamical system with contrast to VAR and DMD," *Chaos* **31**, 013108 (2021).
- ⁴¹D. J. Gauthier, E. Bollt, A. Griffith, and W. A. Barbosa, "Next generation reservoir computing," *Nat. Commun.* **12**, 1–8 (2021).
- ⁴²L.-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).
- ⁴³H. Fan, L.-W. Kong, Y.-C. Lai, and X. Wang, "Anticipating synchronization with machine learning," *Phys. Rev. Res.* **3**, 023237 (2021).
- ⁴⁴J. 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. Machine Intell.* **3**, 316–323 (2021).

- ⁴⁵L.-W. Kong, H.-W. Fan, C. Grebogi, and Y.-C. Lai, “Emergence of transient chaos and intermittency in machine learning,” *J. Phys. Complex.* **2**, 035014 (2021).
- ⁴⁶R. Xiao, L.-W. Kong, Z.-K. Sun, and Y.-C. Lai, “Predicting amplitude death with machine learning,” *Phys. Rev. E* **104**, 014205 (2021).
- ⁴⁷D. 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).
- ⁴⁸L.-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).
- ⁴⁹Z.-M. Zhai, M. Moradi, L.-W. Kong, and Y.-C. Lai, “Detecting weak physical signal from noise: A machine-learning approach with applications to magnetic-anomaly-guided navigation,” *Phys. Rev. Appl.* **19**, 034030 (2023).
- ⁵⁰Z.-M. Zhai, M. Moradi, L.-W. Kong, B. Glaz, M. Haile, and Y.-C. Lai, “Model-free tracking control of complex dynamical trajectories with machine learning,” *Nat. Commun.* **14**, 5698 (2023).
- ⁵¹Y. Zhang and S. P. Cornelius, “Catch-22s of reservoir computing,” *Phys. Rev. Res.* **5**, 033213 (2023).
- ⁵²Z.-M. Zhai, M. Moradi, B. Glaz, M. Haile, and Y.-C. Lai, “Machine-learning parameter tracking with partial state observation,” *Phys. Rev. Res.* **6**, 013196 (2024).
- ⁵³D. A. Smeed, G. D. McCarthy, S. A. Cunningham, E. Frajka-Williams, D. Rayner, W. Johns, C. S. Meinen, M. O. Baringer, B. I. Moat, A. Duchez, *et al.*, “Observed decline of the Atlantic meridional overturning circulation 2004–2012,” *Ocean Sci.* **10**, 29–38 (2014).
- ⁵⁴N. Rayner, D. E. Parker, E. Horton, C. K. Folland, L. V. Alexander, D. Rowell, E. C. Kent, and A. Kaplan, “Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century,” *J. Geophys. Res. Atmos.* **108** (2003).
- ⁵⁵E. Frajka-Williams, M. Lankhorst, J. Koelling, and U. Send, “Coherent circulation changes in the Deep North Atlantic from 16 N and 26 N transport arrays,” *J. Geophys. Res.: Oceans* **123**, 3427–3443 (2018).
- ⁵⁶U. Send, M. Lankhorst, and T. Kanzow, “Observation of decadal change in the Atlantic meridional overturning circulation using 10 years of continuous transport data,” *Geophys. Res. Lett.* **38** (2011).
- ⁵⁷S. A. Cunningham, T. Kanzow, D. Rayner, M. O. Baringer, W. E. Johns, J. Marotzke, H. R. Longworth, E. M. Grant, J. J.-M. Hirschi, L. M. Beal, *et al.*, “Temporal variability of the Atlantic meridional overturning circulation at 26.5 N,” *Science* **317**, 935–938 (2007).
- ⁵⁸T. Kanzow, S. A. Cunningham, D. Rayner, J. J.-M. Hirschi, W. E. Johns, M. O. Baringer, H. L. Bryden, L. M. Beal, C. S. Meinen, and J. Marotzke, “Observed flow compensation associated with the MOC at 26.5 N in the Atlantic,” *Science* **317**, 938–941 (2007).
- ⁵⁹R. M. van Westen, M. Kliphuis, and H. A. Dijkstra, “Physics-based early warning signal shows that AMOC is on tipping course,” *Sci. Adv.* **10**, eadk1189 (2024).
- ⁶⁰J. P. Crutchfield and B. McNamara, “Equations of motion from a data series,” *Complex Sys.* **1**, 417–452 (1987).
- ⁶¹C. Yao and E. M. Bollt, “Modeling and nonlinear parameter estimation with Kronecker product representation for coupled oscillators and spatiotemporal systems,” *Physica D* **227**, 78–99 (2007).
- ⁶²W.-X. Wang, R. Yang, Y.-C. Lai, V. Kovanis, and C. Grebogi, “Predicting catastrophes in nonlinear dynamical systems by compressive sensing,” *Phys. Rev. Lett.* **106**, 154101 (2011).
- ⁶³W.-X. Wang, Y.-C. Lai, C. Grebogi, and J. Ye, “Network reconstruction based on evolutionary-game data via compressive sensing,” *Phys. Rev. X* **1**, 021021 (2011).
- ⁶⁴W.-X. Wang, R. Yang, Y.-C. Lai, V. Kovanis, and M. A. F. Harrison, “Time-series-based prediction of complex oscillator networks via compressive sensing,” *EPL (Europhys. Lett.)* **94**, 48006 (2011).
- ⁶⁵R.-Q. Su, X. Ni, W.-X. Wang, and Y.-C. Lai, “Forecasting synchronizability of complex networks from data,” *Phys. Rev. E* **85**, 056220 (2012).
- ⁶⁶R.-Q. Su, W.-X. Wang, and Y.-C. Lai, “Detecting hidden nodes in complex networks from time series,” *Phys. Rev. E* **85**, 065201 (2012).
- ⁶⁷R.-Q. Su, Y.-C. Lai, and X. Wang, “Identifying chaotic fitzhugh-nagumo neurons using compressive sensing,” *Entropy* **16**, 3889–3902 (2014).
- ⁶⁸R.-Q. Su, Y.-C. Lai, X. Wang, and Y.-H. Do, “Uncovering hidden nodes in complex networks in the presence of noise,” *Sci. Rep.* **4**, 3944 (2014).
- ⁶⁹Z. Shen, W.-X. Wang, Y. Fan, Z. Di, and Y.-C. Lai, “Reconstructing propagation networks with natural diversity and identifying hidden sources,” *Nat. Commun.* **5**, 4323 (2014).
- ⁷⁰R.-Q. Su, W.-W. Wang, X. Wang, and Y.-C. Lai, “Data based reconstruction of complex geospatial networks, nodal positioning, and detection of hidden node,” *R. Soc. Open Sci.* **3**, 150577 (2016).
- ⁷¹Y.-C. Lai, “Finding nonlinear system equations and complex network structures from data: A sparse optimization approach,” *Chaos* **31**, 082101 (2021).
- ⁷²P. R. Vlachas, J. Pathak, B. R. Hunt, T. P. Sapsis, M. Girvan, E. Ott, and P. Koumoutsakos, “Backpropagation algorithms and reservoir computing in recurrent neural networks for the forecasting of complex spatiotemporal dynamics,” *Neural Netw.* **126**, 191–217 (2020).
- ⁷³C. M. Bishop, “Training with noise is equivalent to Tikhonov regularization,” *Neural Comput.* **7**, 108–116 (1995).
- ⁷⁴Z.-M. Zhai, L.-W. Kong, and Y.-C. Lai, “Emergence of a resonance in machine learning,” *Phys. Rev. Res.* **5**, 033127 (2023).
- ⁷⁵F. Nogueira, “[Bayesian Optimization: Open source constrained global optimization tool for Python](#),” (2014–).