



RESEARCH ARTICLE

10.1029/2021MS002537

A Comparison of Data-Driven Approaches to Build Low-Dimensional Ocean Models

Key Points:

- The multi-level stochastic approach produces the most stable, accurate, and low-cost emulator of a double-gyre ocean model solution
- Artificial neural networks and long short term memory work better in a hybrid form with linear regression, providing the core dynamics, than in their standalone application
- Emulators incorporating memory effects and state-dependent noise show enhanced performance and deep learning can learn these effects

Correspondence to:




N. Agarwal,
n.agarwal17@imperial.ac.uk

Citation:

Agarwal, N., Kondrashov, D., Dueben, P., Ryzhov, E., & Berloff, P. (2021). A comparison of data-driven approaches to build low-dimensional ocean models. *Journal of Advances in Modeling Earth Systems*, 13, e2021MS002537. <https://doi.org/10.1029/2021MS002537>

Received 9 MAR 2021

Accepted 9 AUG 2021

Niraj Agarwal¹ , D. Kondrashov^{2,3} , P. Dueben⁴ , E. Ryzhov^{1,5}, and P. Berloff^{1,6}

¹Department of Mathematics, Imperial College London, London, UK, ²Department of Atmospheric and Oceanic Sciences, University of California Los Angeles, Los Angeles, CA, USA, ³Institute of Applied Physics of the Russian Academy of Sciences, Nizhny Novgorod, Russia, ⁴ECMWF, Reading, UK, ⁵Pacific Oceanological Institute, Vladivostok, Russia, ⁶Institute of Numerical Mathematics, Russian Academy of Sciences, Moscow, Russia

Abstract We present a comprehensive inter-comparison of linear regression (LR), stochastic, and deep-learning approaches for reduced-order statistical emulation of ocean circulation. The reference data set is provided by an idealized, eddy-resolving, double-gyre ocean circulation model. Our goal is to conduct a systematic and comprehensive assessment and comparison of skill, cost, and complexity of statistical models from the three methodological classes. The model based on LR is considered as a baseline. Additionally, we investigate its additive white noise augmentation and a multi-level stochastic approach, deep-learning methods, hybrid frameworks (LR plus deep-learning), and simple stochastic extensions of deep-learning and hybrid methods. The assessment metrics considered are: root mean squared error, anomaly cross-correlation, climatology, variance, frequency map, forecast horizon, and computational cost. We found that the multi-level linear stochastic approach performs the best for both short- and long-timescale forecasts. The deep-learning hybrid models augmented by additive state-dependent white noise came second, while their deterministic counterparts failed to reproduce the characteristic frequencies in climate-range forecasts. Pure deep learning implementations performed worse than LR and its simple white noise augmentation. Skills of LR and its white noise extension were similar on short timescales, but the latter performed better on long timescales, while LR-only outputs decay to zero for long simulations. Overall, our analysis promotes multi-level LR stochastic models with memory effects, and hybrid models with linear dynamical core augmented by additive stochastic terms learned via deep learning, as a more practical, accurate, and cost-effective option for ocean emulation than pure deep-learning solutions.

Plain Language Summary In weather and climate predictions, scientists use comprehensive ocean circulation models for representing the effects of the oceans on the atmosphere. These models simulate the three-dimensional ocean dynamics using millions of variables and, thus, require significant computational resources and running time. Therefore, there is a need for low-cost, data-driven ocean models with fewer variables that can reproduce essential oceanic circulations with reasonable accuracy. There are several popular data-driven approaches to build these models, but singling out the best one is difficult and significantly understudied. We have systematically assessed and compared the accuracy, stability, and computational cost of various data-driven models against the linear regression—a fundamental and easy-to-implement deterministic model, that is, it provides a fixed output for a fixed input. We considered several stochastic and deep-learning models for comparison; stochastic models combine a deterministic model with customized noise, whereas deep-learning models train a complex network of neurons similar to the human brain. We found that the stochastic models that properly include the core dynamics, time-delay effects, and model errors perform the best. The core dynamics provides the essential changes, time-delay effects are the changes due to correlation between successive ocean states, and model errors provide other possible causes of changes.

1. Introduction

Medium-range weather forecast models routinely use computationally expensive ocean general circulation models (OGCMs) that are coupled to the atmosphere model. However, the long timescales of ocean dynamics and the weak influence from the deeper layers of the ocean on the atmosphere for a weather forecast of, say, a couple of days justify the investigation of replacements of expensive OGCMs with low-dimensional

data-driven models that can run at negligible cost and emulate the upper ocean. Here, these models are referred to as ocean “emulators,” because they emulate statistical properties of the flow rather than simulating the dynamics derived from physical principles. Our definition of emulators is slightly different when compared to a number of studies that try to replace components of existing models to reduce computational cost. In our case, we do not aim to emulate a model component but rather the physical system that has generated the data—a dynamical ocean model. Ocean emulators are also useful (a) in long-time climate-type model simulations for process sensitivity studies, (b) in climate prediction, and (c) for improving ensemble forecast statistics. Furthermore, ocean emulators can potentially be down-scaled and used for data-driven parameterizations of mesoscale (and even sub-mesoscale) eddies for non-eddy-resolving and eddy-permitting comprehensive OGCMs. Finally, they can also be used as conceptual toy models for process-related studies (e.g., as kinematic flow emulators of material transport).

The physical ocean models have the fundamental advantage that they can operate even without training from data, which simply may not exist. The main problems with the physical models are that for some practical applications they can be prohibitively expensive or may not allow resolving all important features, some of the involved physics can be inaccurately accounted for, and numerical and discretization errors can be unacceptable. On the other hand, data-driven emulators, which are the focus of this study, can be much cheaper, accurate (for the data-trained regimes), and simple to deal with, which gives them a crucial advantage for many practical problems. However, they can be hard to interpret physically as many of them are ultimately used as a black box, for example, machine-learning-based methods. Low-cost emulators can be constructed in terms of the empirical orthogonal functions (EOFs) and their principal components (PCs; Lorenz, 1956), but the true governing equations in the EOF space are always unknown. Our approach to ocean emulation is supported by the existing and rapidly developing methodologies for statistical data-driven modeling (reviewed in Brenner et al., 2019; Rowley & Dawson, 2017). We consider three major statistical model types: linear regression (LR), stochastic, and deep learning.

LR belongs to the broader family of regression-only statistical models, where some polynomial surfaces are fit to the data while forbidding correlations with the residuals. The fitted polynomials are believed to capture “enough” of the dynamics, so that the residuals can be attributed to the uncertainty in the initial conditions and internal variability. The regression-only models benefit from significant speed-up as there is no need to identify the covariance matrix and the associated calculations, such as covariance inversion, which can be complex and computationally expensive for high-dimensional systems. The fast, simple, and easy-to-implement characteristics of regression-only models found their numerous applications in climate and environmental sciences (Holden et al., 2013; Sexton et al., 2012; Williamson, 2015). Here, we use them in the linear form, due to incomplete knowledge about the nonlinear basis functions of the ocean circulation tendencies, and develop the simplest ocean emulator that provides the baseline.

Next, stochastic models use parameterized noise signals to deal with missing/unknown physics, parameter uncertainties, inaccurate initial conditions, and noise-induced regime transitions (Berner et al., 2017; Sardeshmukh et al., 2001; Sura et al., 2005; Zanna et al., 2017) in chaotic dynamical systems. They produce random outputs, thus, allowing for ensemble statistics for uncertainty estimates. Stochastic terms can be added to any deterministic framework in two common ways, as either additive or multiplicative noises. The additive noise is directly added to the basic equations, whereas the multiplicative noise is added after multiplying it with the amplitude function depending on the predicted model variables. The latter models are significantly more complicated, and fitting their parameters is more difficult, yet both have distinct advantages and are applied in ocean and climate modeling for numerous purposes. The additive noise has been used to force linear dynamical models (DelSole & Hou, 1999; Farrell & Ioannou, 1993; Zhang & Held, 1999), to model effects of subgrid-scale turbulence (Berloff, 2005; Berloff et al., 2007; Berloff & McWilliams, 2003; D’Andrea & Vautard, 2001; DelSole, 2004; Farrell & Ioannou, 1995; Williams et al., 2016), to provide stochastic climate predictions (Majda et al., 1999; Seiffert & Von Storch, 2008; Seiffert & von Storch, 2010), and to derive stochastic primitive equations for the oceans and atmosphere (Ewald et al., 2007). On the other hand, the multiplicative noise strategy is considered most relevant for modeling non-Gaussian statistics, such as extreme events, tipping points in the dynamical systems (Franzke, 2012, 2013; Sura, 2011, 2013), uncertainty estimates in parameterization schemes (Buizza et al., 1999; Juricke et al., 2013, 2017; Ollinaho

et al., 2017), stochastic primitive equations (Debussche et al., 2012; Glatt-Holtz & Ziane, 2008), and nonlinear coupling between noise and model variables (Franzke et al., 2005; Sardeshmukh et al., 2003).

Recently, a multi-level framework of the additive nonlinear stochastic models—called Empirical Model Reduction (Kondrashov et al., 2005; Kravtsov et al., 2005, 2009; Kondrashov et al., 2013, EMR)—has also been developed. It allows to include inherently important time-delayed (memory) effects both in the additive state-dependent stochastic forcing and dynamical operator (Kondrashov et al., 2015), and allows for a more complicated temporal structure of the noise. Data-driven climate models based on EMR formulation have proven to be highly competitive in prediction and process studies (Chen et al., 2016; Ghil et al., 2018; Strouanine et al., 2010). The linear inverse modeling (LIM) approach (Penland, 1989, 1996; Penland & Sardeshmukh, 1995) is a specific case of the EMR with a linear propagator and additive white noise. Kondrashov and Berloff (2015) have shown that decadal oceanic variability can be successfully simulated by a linear EMR formulation and a change of basis, namely instead of principal component analysis (PCA) modes, using modes identified by multichannel singular spectrum analysis (M-SSA) which incorporates time-delayed embedding (Ghil et al., 2002). Similar to M-SSA, data-adaptive harmonic decomposition (DAHD) (Kondrashov et al., 2018, 2020; Ryzhov et al., 2020) relies on the eigendecomposition of the lag-covariance matrix. However, unlike M-SSA, DAHD modes form an orthonormal set of spatial patterns oscillating harmonically within the time-embedding window, and thus can be modeled by a system of coupled frequency-ranked nonlinear stochastic oscillators.

For this study in the class of stochastic methods, in order to make a fair comparison, we have made a deliberate choice to focus on the methods that are commonly used in PCA basis and implemented linear stochastic models for white noise, as well as the linear formulation of the multi-level EMR framework.

The third statistical model type investigated in this study is deep-learning models. These models approximate the intricate nonlinear functional relationships between the model inputs and outputs by training an extensive parametric network of interconnected nodes, using neither physical knowledge about the system nor the governing differential equations. With the recent advancements in computing power, simple feed-forward artificial neural networks (ANNs), convolutional neural networks (CNNs), and long short term memory (LSTM) models rose to prominence in many science disciplines and helped to find hidden patterns in multi-dimensional data sets. Feed-forward ANN is the most commonly used deep-learning model, and its design includes multiple layers of small blocks of equations communicating in a nonlinear fashion. These ANNs have been used broadly in oceanic and atmospheric studies, ranging from the idealized Lorenz63 (Lorenz, 1963) and Lorenz96 (Lorenz, 1996) models (Dueben & Bauer, 2018; Scher & Messori, 2019) to more realistic situations, such as developing subgrid-scale models (Karunasinghe & Liong, 2006; Maulik et al., 2019; Rasp et al., 2018), learning the inter-dependency between global climate and vegetation fields (Holden et al., 2015), super-parameterizations (Chattopadhyay, Subel, & Hassanzadeh, 2020), and spotting extreme events in complex weather data sets (Liu et al., 2016). CNNs (Krizhevsky et al., 2012) form another class of deep-learning methods that has been extensively used in geophysical fluid dynamics to identify (and regress) patterns in turbulent flow regimes by repeatedly convoluting the inputs with appropriate kernels (Bolton & Zanna, 2019; Chattopadhyay, Mustafa, et al., 2020; Liu et al., 2020; Weyn et al., 2020). LSTMs (Hochreiter & Schmidhuber, 1997) and reservoir computing (Jaeger & Haas, 2004; Nadiga, 2021; Pathak et al., 2018) are specific forms of Recurrent NNs that can progress learned information from one timestep to the next when applied in an iterative way. This improves the time evolution of the model which makes the LSTMs attractive for applications in oceanic and atmosphere modeling that show multi-scale and lagged interactions (Salman et al., 2018; Vlachas et al., 2018; Zhang et al., 2017).

Despite first attempts to interpret deep-learning models physically (McGovern et al., 2019; Portwood et al., 2021), they remain mostly black boxes. Therefore, the modern trend is to use them in combination with physical models (Karpadne, Atluri, et al., 2017; Reichstein et al., 2019)—referred to as hybrid methods—to potentially increase the forecast skills of an imperfect physical model. Also, the involved neural network may require less training and complexity (Jia et al., 2019). A few applications of hybrid methods using ANN and LSTMs are in Krasnopolsky and Fox-Rabinovitz (2006), Pawar et al. (2020), Rahman et al. (2018), and Watson (2019). Similarly, hybrid deep-learning stochastic approaches are also being developed (Mukhin et al., 2015; Seleznev et al., 2019), but this area remains understudied, and its full potential has not yet been explored. In this study, we test all of the above deep-learning models, except for the CNN, which is

best suited to image-based data sets and structured grids, whereas, here, we work in the EOF/PC space to achieve a significant model order reduction. We also propose novel hybrid stochastic formulations by utilizing residuals from the deep-learning procedure, thus effectively providing nonlinear state-dependent noise.

How do different emulators from the selected three classes compare against each other in terms of their skills? This is a difficult and significantly understudied question, which is central to our study. In particular, as most papers will only evaluate a single method for a specific data set, quantifiable intercomparison of different methods are often impossible across papers. In the context of geophysical applications, we can speculate the LR to be the least successful due to its purely linear form and deterministic nature, making it ineffective in accounting for the inherent uncertainties due to, for example, insufficient resolution, unresolved processes, parameterization errors, etc., and imprecise/incomplete knowledge of many geophysical processes (especially on the reduced-dimension space) and scale interactions caused by nonlinear terms of the differential equations. The stochastic models can deal with model and forecast uncertainties and, therefore, are expected to perform better than LR. However, it is hard to predict their performance relative to the deep-learning methods, which are generally deterministic (note that they can also be Bayesian) but capable of producing accurate and generalizable models. In this study, we aim to compare a number of ocean emulators for relatively simple ocean circulation data obtained from a long-term model simulation of an idealized ocean model. We do not use ocean observations as model data to avoid problems due to measurement errors, complex coast-lines, gaps in observations, biases between different observational products, etc. which would make a fair comparison between the different emulators more difficult. However, the study could easily be repeated with observational data sets. We aim to develop emulators, which are computationally cheap, able to reproduce the statistical characteristics of the reference flow, and are capable of providing simulations on climate-range time scales.

Section 2 provides details about the data sets; Section 3 discusses all the modeling frameworks; Section 4 presents the model assessment metrics and their outcomes, and Section 5 discusses the results and concludes.

2. Data Set

The reference data set used in this study was generated by a three-layer double-gyre quasigeostrophic ocean model (Berloff, 2015; Ryzhov et al., 2019), which provides an idealized representation of the North Atlantic wind-driven circulation dominated by the subtropical and subpolar gyres, and by the Gulf Stream current. A square ocean basin with the side 3,840 km is considered; a steady asymmetric wind forcing at the top layer is imposed; the partial-slip boundary condition is used on the lateral boundaries; stratification is imposed with the first and second Rossby deformation radii equal to 40 and 20.6 km, respectively; and the grid resolution is 7.5 km. Since the model is dynamically eddy-resolving, it qualitatively correctly reproduces the eastward jet extension of the western boundary currents and its adjacent recirculation zones, as well as the mesoscale eddy variability and interdecadal oscillations (of period 17 years). The reference solution of the statistically stationary flow regime is obtained in terms of the evolving potential vorticity and velocity streamfunction that are related to each other via elliptical inversion. The solution snapshots are saved after every 10 days, for a total of about 1,400 years (5×10^5 days). Both potential vorticity anomaly and streamfunction snapshots (Figure 1) show two asymmetric gyres of opposite circulations separated by the eastward jet region, which is characterized by the most vigorous flow variability, and, therefore, is in the focus of our study.

For the reference data set, we chose velocity streamfunction, because it is smoother than the potential vorticity anomaly (both of them represent the same dynamical regime). We deal with the upper isopycnal layer, as it contains the most intensive flow variability and is the most relevant for numerical weather predictions. Next, we re-organize the solution description and reduce the dimensionality of the raw data by using singular value decomposition (SVD), which allows us to find dominant spatial patterns, called empirical orthogonal functions (EOFs), and the corresponding temporal coefficients, called principal components (PCs). Before SVD, the reference 3D spatio-temporal data set ($\in \mathbb{R}^{m \times m \times n}$; $m = 513, n = 5 \times 10^4$ with the grid-dimension in each direction m and the number of records n) is reshaped into a matrix ($\in \mathbb{R}^{n \times m^2}$), where each row is a snapshot of the flow. It is then smoothed along the temporal dimension with a 100-days (i.e., 10 records) running-average window, in order to focus on the long-timescale tendencies (the window size is truncated at the endpoints), followed by an SVD decomposition to obtain n EOFs/PCs. These EOFs/

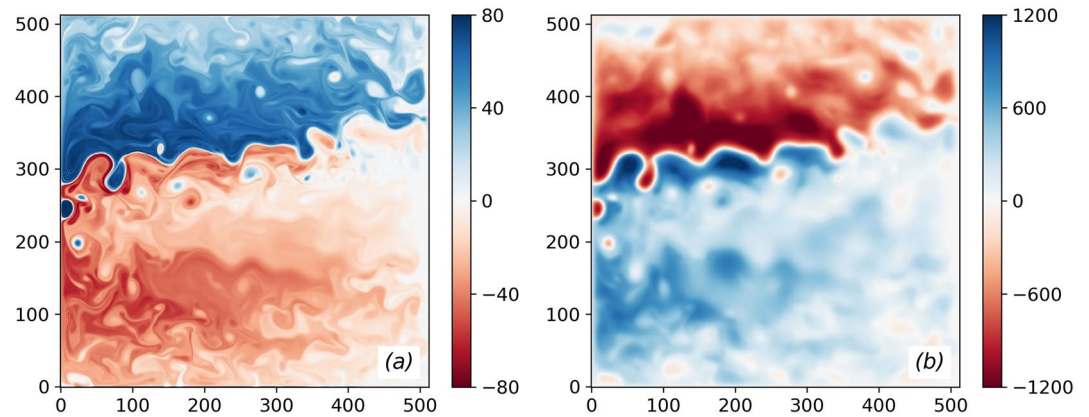


Figure 1. Illustration of the reference eddy-resolving ocean circulation solution. Snapshots of the top-layer (a) Potential vorticity anomaly and (b) Velocity streamfunction. In the case of (a), the anomaly is computed by subtracting the planetary vorticity from the full potential vorticity. Positive (red) and negative (blue) values of potential vorticity anomaly correspond to cyclonic and anticyclonic motions, respectively, and oppositely for the streamfunction. The color scales are in dimensionless units. Non-dimensionalization is done using the length and velocity scales equal to 7.5 km (grid interval) and 0.01 m/s, respectively, and CGS units.

PCs are used for the emulation and analyses in this work. Each EOF explains a fraction of the total temporal variance, and all EOFs are ranked so that this fraction decreases with k ; for complex geophysical data, this decreases exponentially (Ghil et al., 2002). For the purposes of our study, we considered the leading 150 EOFs and their PCs, which all together represent about 90% of the total variance. Our choice of relatively small (0.3% of the total number of EOFs), yet dynamically significant, number of EOFs is justified by our goal to gain a foothold in developing and applying a systematic methodology for comprehensive assessment of the model skills. However, we admit that we are trying to build a simplified statistical model of the QG dynamics, which itself is a simplification of the comprehensive general circulation model dynamics, therefore, our study does not present the real ocean situation, and the generalization of the results presented here are only possible up to a certain extent. Figure 2 shows the top-five and the bottom-most EOFs among the 150 EOFs/PCs considered here. The top-ranked EOFs represent the most dominant patterns along the eastward jet region, which is in the focus of this study, as it contains most of the variability, while the last EOF represents high-frequency variabilities across the basin.

Our aim is to emulate the leading 150 PCs of the upper-ocean streamfunction (Figure 3) efficiently and to construct the corresponding reduced-order models.

3. Modeling Frameworks

In this section, we expand on each adopted modeling methodology one-by-one, by providing the mathematical formulations and the relevant parameters and hyper-parameters, and by explaining the model training processes. Before modeling, we normalized the PCs by the respective standard deviations such that each of them follows zero mean and unit standard deviation. However, the results were found to be robust to other normalizations too, for example, division by the standard deviation of the top-most PC (similarly for auxiliary variables, if any). Within each method, we model either tendencies or state of the PCs, denoted by y_i ; $i = 1, 2, 3, \dots, N$, where $N = 150$ is the total number of PCs considered. The tendencies are computed numerically using the finite-difference method with $\Delta t = 10$ days as the time difference between any two successive records. Out of the total 50K records (which covers a time interval of 500K days, as we have recorded the snapshots after every 10 days) of the PCs, we set the initial 40K records as the training data set, and the last 10K records as the test data set. This training length is roughly 400 times the decorrelation time scale of the topmost PC (≈ 2.7 years) and includes nearly 65 periods of the intrinsic low-frequency variability of period 17 years. Note that such a partitioning of training and test data sets may leave the end of the training data set and the beginning of the test data set correlated. But this is inconsequential because the test

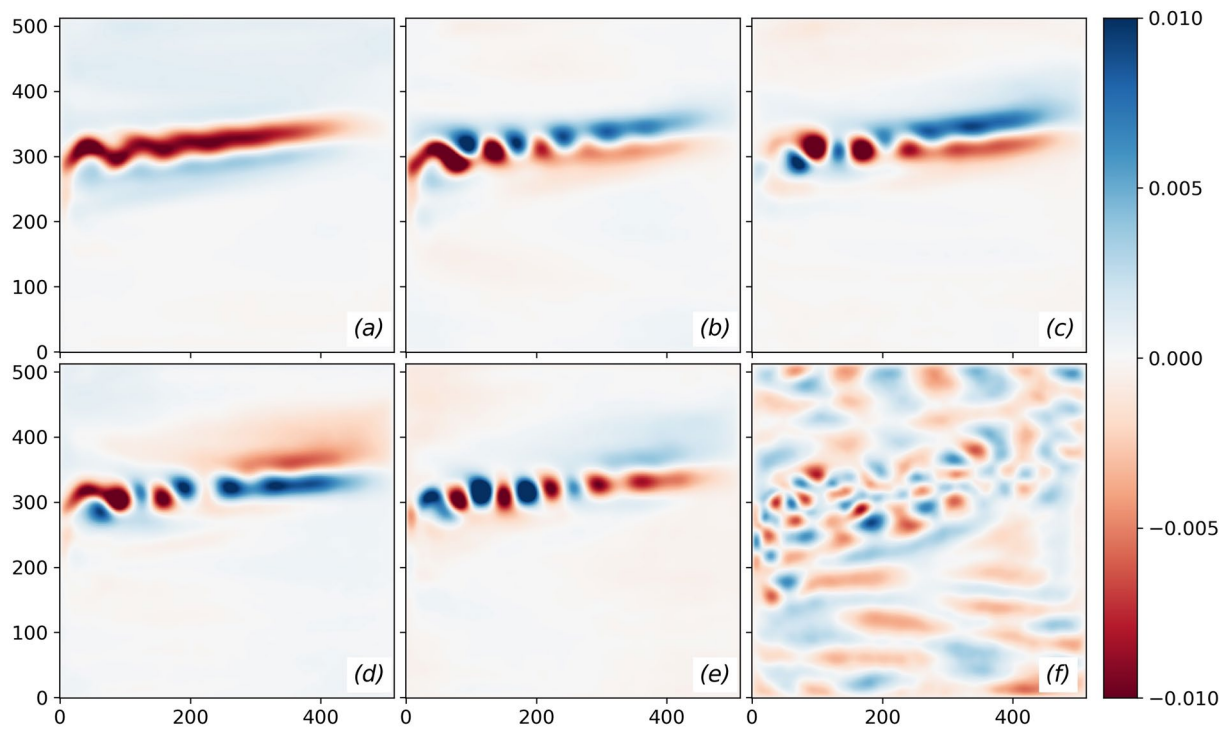


Figure 2. Representation of the reference data in terms of the empirical orthogonal functions (EOFs) of the upper-layer streamfunction. Panels (a)–(e) represent the top-five and panel (f) represents the bottom-most EOFs from the selected suite of top-150 EOFs/ principal components. Note that all top ranked patterns significantly contribute to the multi-scale variability in the eastward jet region, which is the focus of our study, but the rest of the basin is also impacted. The bottom ranked EOF mostly accounts for small-scale variabilities in the entire domain. The color scale values are in the dimensionless units; non-dimensionalization is done using the same scales as mentioned in Figure 1.

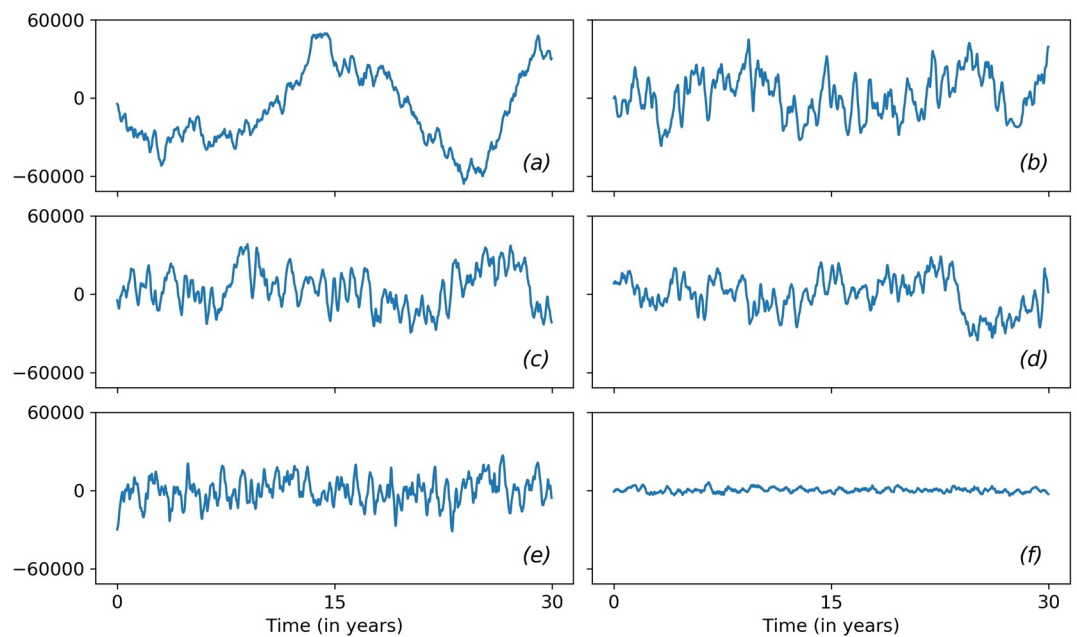


Figure 3. Representation of the reference data in terms of the principal components (PCs) corresponding to the empirical orthogonal functions in Figure 2. Note that PCs are shown only for 30 years of the total record of about 1,400 years. Also, note that the low-ranked PC is much noisier with a smaller variance. The y-axis values are dimensionless.

data set still possesses roughly 16 cycles of the intrinsic low-frequency variability, and we have considered multiple initial conditions for forecasts and reported their average.

We considered the full length of the training data set for all models but also experimented with different lengths ranging from 1K to 40K records as in real-world applications, such a large training data set (over 1,000 years) may not be available. Our analysis showed that at least 10% of the training data (corresponding to ≈ 7 cycles of low-frequency variability) is needed for our models to have similar performance on the test data set as reported for the full length of the training data set. Transfer learning may help when training data-driven models for the real ocean with limited observations. In this approach, the models are trained first using the historical climate data, such as CMIP model outputs, before fine tuning using observations and reanalysis data (Ham et al., 2019; Rasp & Thuerey, 2021).

3.1. Linear Regression (LR)

This model is expressed as follows:

$$\frac{dy}{dt} = \mathbf{L}y; \quad (1)$$

where $\mathbf{y} = [y_1 \ y_2 \ \dots \ y_N]$ is the multivariate PC, and \mathbf{L} is an $N \times N$ matrix of regression coefficients, referred to as the system matrix. LR is the simplest model considered in this work and is used as a baseline for assessing the performance of the others. Linear dynamics is pertinent to the gyres and explains a significant proportion of the low-frequency variabilities in the top PCs. Therefore, we also term LR as the “core dynamics.”

3.2. LR With Additive White Noise (LR-AWN)

This is a stochastic model based on the LR model (1) but additionally equipped with the additive white noise. The model is given as:

$$dy = \mathbf{L}ydt + \mathbf{Q}d\mathbf{W}; \quad (2)$$

where $d\mathbf{W} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ with \mathbf{I} being the identity matrix; \mathbf{Q} is a lower triangular matrix obtained by Cholesky decomposition of the lag-zero covariance matrix $\mathbf{C} = \mathbf{Q}\mathbf{Q}^T$ for the regression residuals $dy - \mathbf{L}ydt$, and $\mathbf{Q}, \mathbf{C} \in \mathbb{R}^{N \times N}$. Here, we use $dt = 1$ for convenience (but without loss of generality) as we deal with discrete time series. We validated the output of this model against a Linear Inverse Model, which is more widespread, but found almost no difference between the two model outputs.

3.3. Multi-Level Linear Regression (ML-LR)

We adopt a linear formulation of multi-level stochastic Empirical Model Reduction approach (Kravtsov et al., 2009; Kondrashov & Berloff, 2015; Kondrashov et al., 2015), where the regression residuals are not immediately replaced by some noise but instead are modeled by using a stack of levels. The top-most level is similar to (Equation 1), and the lower levels regress the higher-level residuals as the additional (hidden) state variables, until the lowest-level residual degenerates into spatially uncorrelated white noise, that is, their autocorrelation approaches zero (technical implementation of the stopping criterion is based on fraction of the explained variance by regression, see Appendix A in Kondrashov et al., 2015). The complete model can be expressed as:

$$\text{Level 1: } d\mathbf{y} = \mathbf{L}ydt + \mathbf{r}(t)dt, \quad (3a)$$

$$\text{Level 2: } d\mathbf{r} = \mathbf{M}^{(1)}[\mathbf{y}; \mathbf{r}]dt + \mathbf{r}^{(1)}(t)dt, \quad (3b)$$

$$\text{Level 3: } d\mathbf{r}^{(1)} = \mathbf{M}^{(2)}[\mathbf{y}; \mathbf{r}; \mathbf{r}^{(1)}]dt + \mathbf{r}^{(2)}(t)dt, \quad (3c)$$

$$\text{Level } L: d\mathbf{r}^{(L-2)} = \mathbf{M}^{(L-1)}[\mathbf{y}; \mathbf{r}; \mathbf{r}^{(1)}; \dots; \mathbf{r}^{(L-2)}]dt + \mathbf{Q}d\mathbf{W}, \quad (3d)$$

where, as in Section 3.2, $d\mathbf{W}$ is an independent Gaussian white noise process; and \mathbf{Q} is the lower triangular matrix obtained by the Cholesky decomposition of the zero-lag covariance between the last-level residuals. In our results, the residuals at the second level $\mathbf{r}^{(1)}(t)$ become sufficiently decorrelated in time (according to the stopping criterion) and are well approximated by the spatially correlated white noise, so we used $L = 2$.

Note that for the prediction experiments in Section 4, the initial conditions for the residuals at various levels need to be determined in strictly “no-look-ahead” procedure, that is, only using the model coefficients and the time history of \mathbf{y} prior to the forecast start time. For example, we would need to begin from time instant $(t - 1)$ to make forecasts from time instant t , when using a model with one extra level (see Appendix B in Kondrashov et al., 2015) and initializing $\mathbf{r}(t - 1) = (\mathbf{y}(t) - \mathbf{y}(t - 1))/dt - \mathbf{L}\mathbf{y}(t)$. To obtain numerical results we have used the publicly available Stochastic Modeling and Prediction Toolbox (see Acknowledgments).

3.4. Artificial Neural Network (ANN)

A neural network is a computational architecture loosely based on the biological networks of neurons in the human brain. Each neuron is an instance of an activation function (e.g., linear, binary, hyperbolic, and sigmoid) that operates on the weighted sum of its inputs with some bias added into it. Multiple neurons can be combined into distinct neural network architectures. A feed-forward artificial neural network (ANN) is among the most common (Nielsen, 2015). It is composed of multiple layers of neurons so that outputs from one layer are the inputs to the next layer, with the ultimate goal to approximate the right functional relationship between the inputs and outputs. In a dense network, each neuron in a layer receives inputs from all the neurons in the previous layer and, thus, exhibits a compact set of connections between the available neurons, inputs, and outputs. Each connection in the network is characterized by its weight, and the goal of an ANN is to optimize them using a suitable loss function and optimization algorithm. An ANN is defined by a few hyper-parameters controlling its performance (e.g., number of hidden layers, number of neurons per layer, activation function, optimizing function, and optimizer). To adjust the hyper-parameters to the optimal values is a non-trivial task (due to the size of hyper-parameter space), and the adjustment is mostly based on trial-and-error testing that adjusts the model complexity to the amount of data available and the complexity of the problem.

We implemented the ANN using Tensorflow, from Keras Google API, which takes the state of the PCs at time t as the input and returns the state at time $(t + 1)$ as the output. For our training data, the best performing ANN contains two hidden layers of neurons, each with 100 neurons, hyperbolic tangent as the activation function, *Adam*—a first-order gradient-based optimizer for stochastic objective functions—as the optimizer, and mean absolute error as the loss function. We tried several combinations of the hyper-parameters—hidden layers, activation function, and optimizer—and picked up their optimal combination based on tracking the loss function and the naked-eye perception of the model results. Here we tested ANNs with up to four hidden layers, each with 100, 150, or 200 neurons; linear, elu, relu, sigmoid, and tanh as the activation functions; and RMSprop, SGD, Adagrad, and Adam as the optimizers. Our intermediate complexity of the model and the training over a prolonged data set (i.e., 40K records) help us avoid over-fitting. As a sanity check, we trained ANN as $dy/dt = ANN(\mathbf{y}(t))$ with the “linear” activation function and compared the outcomes with LR; we found the results to be very similar to each other.

To improve the ANN forecasts further, we added a spatially correlated white noise to the ANN forecasts to account for the residuals, similar to LR in Equation 2. This model is abbreviated as ANN-AWN and is given as:

$$\mathbf{y}(t + 1) = ANN(\mathbf{y}(t)) + \boldsymbol{\zeta};, \quad (4)$$

where $\boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}\mathbf{Q}^T)$, and \mathbf{Q} is the lower triangular matrix obtained by Cholesky decomposition of the covariance of ANN residuals $\mathbf{r}(t) = \mathbf{y}(t + 1) - ANN(\mathbf{y}(t))$. Note that it is also possible to train ANN to predict the perturbation $d\mathbf{y}$ for a given state $\mathbf{y}(t)$ —as done in Chattopadhyay, Hassanzadeh, and Subramanian (2020); Dueben and Bauer (2018)—but this approach resulted in unstable solutions for long integrations.

3.5. Long Short Term Memory (LSTM) Model

LSTM models belong to the class of recursive NNs and function by passing information from previous timesteps to calculate the next timestep when used iteratively. Because these models hold essential dynamical information between the successive time steps, they account for long-time correlations between the model states. This is a significant advantage over ANNs, for an application with significantly autocorrelated time series. Like ANNs, LSTMs can also be upgraded using spatially correlated white noise (abbreviated as LSTM-AWN) with the noise parameters inferred using the LSTM residuals.

We used Keras Google API to implement a two-hidden-layered densely-connected LSTM configuration with 100 neurons in each layer. The model was trained using mean absolute error as the loss function, “Adam” as the optimizer, hyperbolic tangent as the activation function, and the whole 400K days as the training length. Unlike ANNs, the LSTM model takes the state of the PCs at five previous time steps as the input and produces the next state as the output—the so-called “look back” hyperparameter is 5. Higher values of look back did not improve model performance significantly, but the overall computational cost of training/prediction increases many folds (note that LSTM is optimized for taking into account long-time correlation effects by construction). For mini-batches, we used 32, 64, and 128 as its potential values and found 32 to be optimal—amounting to nearly 1 year of observations. For all other hyperparameters, we used the same hyperparameter search space as described in ANN, and the final values were chosen after testing their different combinations and analyzing the resulting model performance on the training data.

3.6. Hybrid Modeling

The hybrid model that combines LR, which conveys the linear dynamics, with the deep-learning models—used as a non-linear correction, state dependence, and memory term—may be more skilled than the standalone implementation of these methods. Such a hybrid model can potentially also preserve some core dynamics of the system and may also benefit from simpler algorithms and architectures in the spirit of the-ory-guided data science (Karpatne, Watkins, et al., 2017).

We trained the deep-learning models (say, f) from the previous sections (ANN and LSTM) to emulate the LR residuals $\mathbf{r}(t) = d\mathbf{y}(t)/dt - \mathbf{L}\mathbf{y}(t)$ and, thus, augment the LR output as:

$$\begin{aligned} d\mathbf{y}(t) &= \mathbf{L}\mathbf{y}(t)dt + \mathbf{r}(t)dt, \\ \mathbf{r}(t) &= f(\mathbf{r}(t-1), \mathbf{y}(t-1); \mathbf{r}(t-2), \mathbf{y}(t-2); \dots; \mathbf{r}(t-l), \mathbf{y}(t-l)) + \xi, \end{aligned} \quad (5)$$

where l is the “look back” hyperparameter for LSTM; for ANN, it is equal to 1 by construction. For LSTM, we set $l = 3$ after checking the LSTM-hybrid model performance on the training data for $l = 1, 2, \dots, 5$ and finding that the model performance does not improve beyond $l = 3$. The model learning proceeds in three successive steps as follows: (a) LR is used to estimate \mathbf{L} , (b) the resulting residual $\mathbf{r}(t) = d\mathbf{y}/dt - \mathbf{L}\mathbf{y}(t)$ is modeled by $f(\mathbf{r}(t-1), \mathbf{y}(t-1); \dots; \mathbf{r}(t-l), \mathbf{y}(t-l))$ that accounts for non-linear correction, state dependence and possibly memory effects, and (c) the final residual from deep-learning minimization is approximated by a spatially correlated white noise process $\xi \sim \mathcal{N}(0, \mathbf{Q}\mathbf{Q}^T)$. This procedure can be interpreted as incorporating state-dependent noise $\mathbf{r}(t)$, and, as said previously, it is implemented for the two deep-learning methods (described in the previous sections), referred to as LR-ANN-AWN and LR-LSTM-AWN. We have also evaluated versions of Equation 5 with no stochastic forcing, that is, without the white noise term, that are abbreviated as LR-ANN and LR-LSTM.

Note that unlike their standalone implementations, the deep-learning models here take both the state $\mathbf{y}(t)$ and residual $\mathbf{r}(t)$ as inputs, and return the residual $\mathbf{r}(t+1)$ as the output. We tested this configuration against several others and found that the current setup performs better than the others on both training and test data sets.

4. Results

In this section, we consider each model assessment metric separately, summarize them, and report the outcome for all models. For the majority of the assessment metrics, we use the reconstructed spatio-temporal streamfunction field (ψ_1^{fcst}), obtained by multiplying the forecasted PCs (y_i ; $i = 1, 2, 3, \dots, 150$) with the

respective EOFs (ϕ_i). We obtain PC forecasts corresponding to a set of initial conditions. The exact number of initial conditions and the lengths of the forecasts differ for the assessment metrics and are provided in the detailed descriptions. Additionally, for stochastic methods, we obtain an ensemble of 100 realizations for each initial condition and calculate the ensemble mean. The reference truth (ψ_1^{ref}) used for assessing the model outputs belongs to the reduced space spanned by the 150 EOFs/PCs. Below, we present the assessment metrics and detailed analyses of the models.

4.1. Root Mean Square Error (RMSE)

The RMSE is given as

$$E(t) = \sqrt{\frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \left(\psi_1^{ref}(x_i, y_j, t) - \psi_1^{cast}(x_i, y_j, t) \right)^2}, \quad (6)$$

and its time series describes the basin-averaged mismatch between the reference and emulated streamfunction snapshots. For perfect forecasts, the RMSE should be zero. However, in practice, RMSE is small for short forecast lead times and grows until it becomes saturated, as the forecast and references truth become uncorrelated at long forecast lead times. We considered the state of the PCs at each of the 10K records of the test data set (10K records correspond to 100K days) as an initial condition and obtained 10 records long (i.e., 100 days) forecasts for all of them; therefore, we used a total of 9,990 initial conditions (as we would not have the reference data for the last 10 initial conditions). Next, we computed RMSEs for each of these forecasts followed by the mean RMSE over all initial conditions; this provides a 10 records long time series for each model (Figure 4a). The motivation for using this forecast length is to study the error growth for short-term forecasts (e.g., on seasonal time scales) as opposed to a change of the mean fields using long-term simulations (e.g., decadal-to-centennial). The short-term predictability for a single initial condition can depend strongly on the underlying flow regime, such as defined by the western boundary current position. It is therefore important to take the average RMSE over a number of initial conditions.

For comparison purposes, we also considered the “persistence” model, where the memory effect is the strongest, and the model state remains constant—equal to the initial condition. The persistence RMSE time series, therefore, characterizes a “constant state” with the absence of a dynamic model.

For 100 days, ML-LR exhibits the best performance with the lowest mean RMSE, followed by the stochastic augmentations of the hybrid models, that is, LR-ANN-AWN and LR-LSTM-AWN (Figure 4a). The deterministic hybrid model RMSEs are similar or worse than LR and its white noise extension; LR-ANN is worse compared to LR-LSTM as the RMSE for the former grows more steeply with the lead time and gets even higher than the Persistence for a lead time beyond 70 days. This suggests that the additive noise is improving the deterministic hybrid methods (clearer from Figure 4b), as the residuals in such models are less correlated and closer to white noise. We also observe that the standalone ANN, both in its deterministic and stochastic version, performs the worst, with its RMSEs being higher than the Persistence at all lead times. The standalone LSTM also performs similarly poorly (yet better than its ANN counterpart), but its stochastic version produces RMSEs lower than the Persistence at high lead times. Nonetheless, a comparison of standalone and hybrid implementations of deep learning methods encourages us to use ANN/LSTM as a nonlinear corrector term rather than using them to represent the complete dynamics.

On the other hand, LR and LR-AWN belong to the middle of the RMSE spectrum and show similar performances (see the box plot, Figure 4b). The similar performances of these models is due to the inability of the noise component to account for the coupled dynamics contained in the LR residuals, which are modeled more efficiently using an extra regression level (as in ML-LR) or deep-learning methods (as in the hybrid methods).

Overall, we conclude that (a) ML-LR and the stochastically augmented hybrid models show better performance than the standalone implementation of LR and deep-learning models, probably, because the former types include all three major components of a reliable model: core dynamics, memory effects, and model errors accounted by stochastic noise and (b) adding simple additive noise to the hybrid models significantly improves their performance.

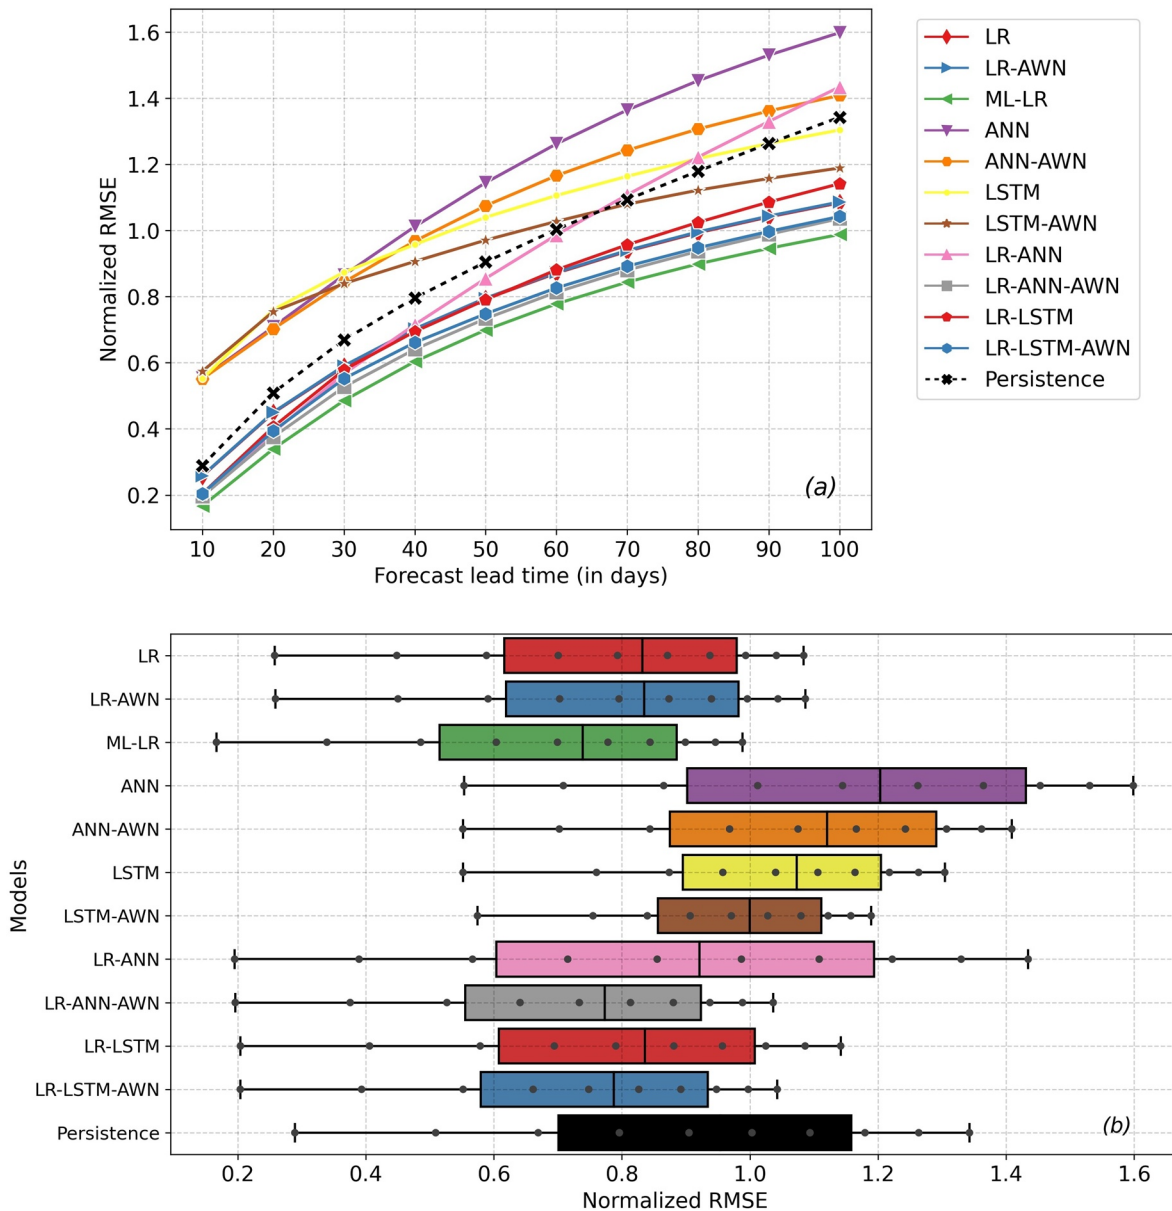


Figure 4. Comparison of the root mean square error (RMSE) curves for all the models used in this work: (a) Line plot, and (b) Box plot. The RMSEs are normalized by the standard deviation of the test data set. The line plot compares the RMSE for all models at each forecast lead time, whereas the box plot provides a visualization of the statistical range and min-max values of the RMSE over 100-days forecast.

4.2. Anomaly Correlation Coefficient (ACC)

Next, we diagnose the correlation between the forecasted (100-days-long forecasts) and the reference truth spatio-temporal streamfunction data sets on spatial and temporal domains, referred to as ASCC and ATCC, respectively. We call this anomaly correlation because we deal with mean-subtracted PCs, and, therefore, the resulting physical fields possess zero mean.

ATCC is a grid-point-wise zero-lag cross-correlation between the forecast and the reference truth over all lead times. It, therefore, gives us a gridded map with the zero-lag cross-correlation value between the forecast and the reference truth for each grid location. Like RMSE, the ATCC map is computed for each of the 9,990 initial conditions followed by their average (Figure 5).

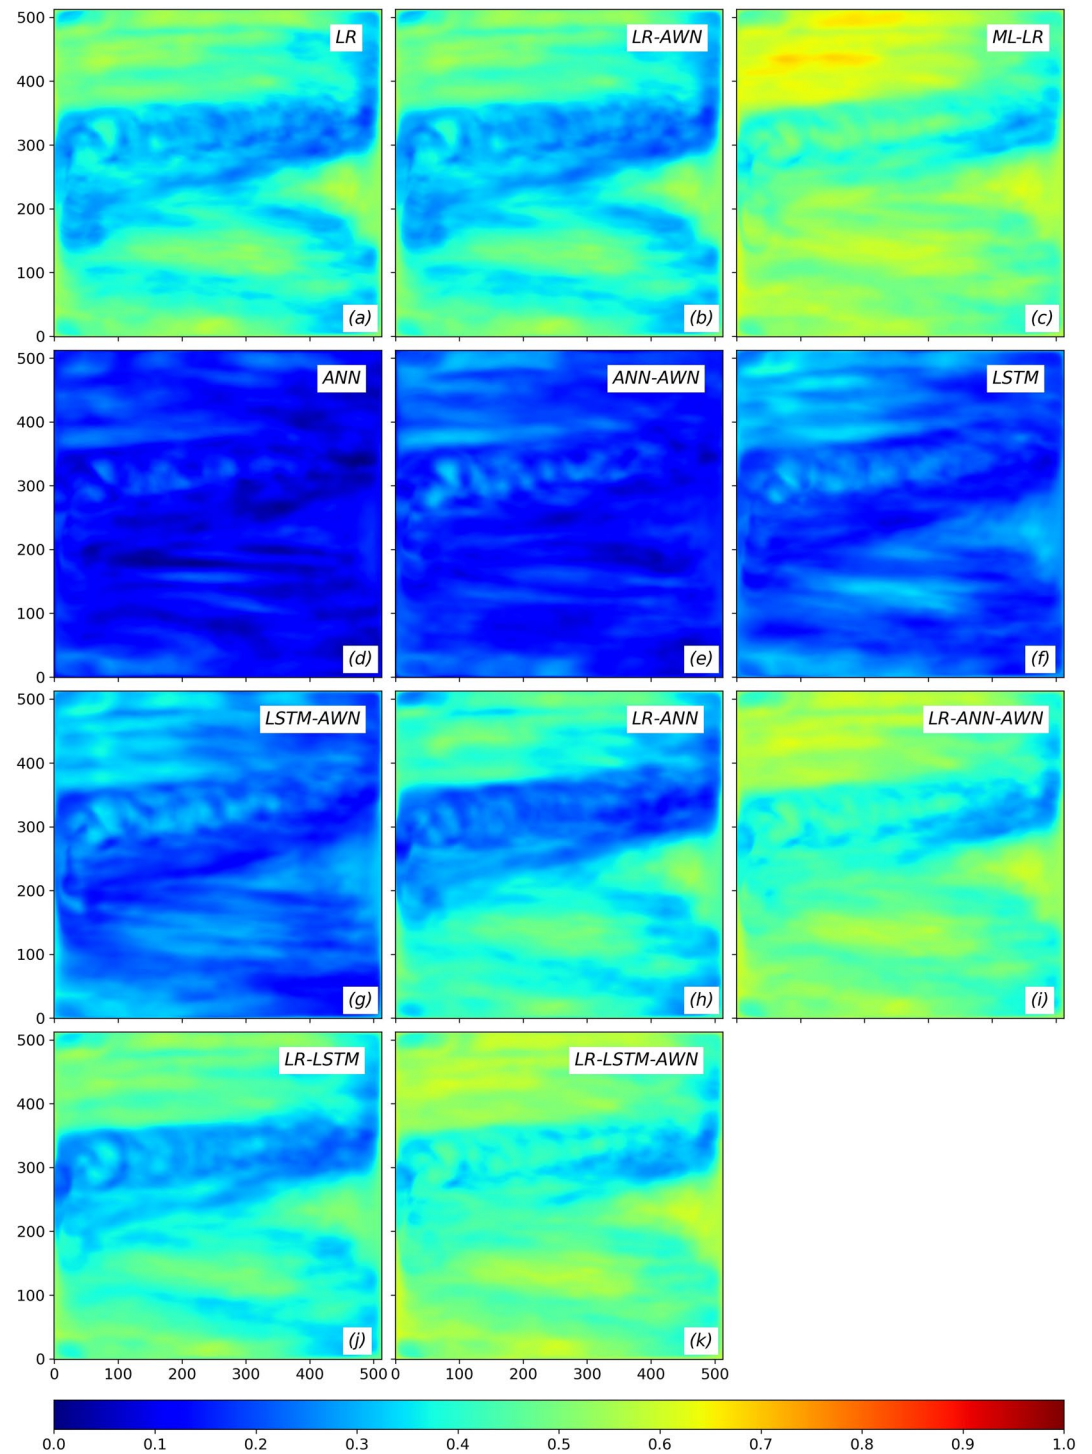


Figure 5. ATCC for 100-days forecasts obtained using different models. The color scale presents the correlation values.

ASCC is the cross-correlation between the spatial snapshots of reference truth and forecast at each lead time, say, t . The snapshots of forecast and the reference truth are reshaped to a vector before cross-correlation. ASCC, therefore, returns a time series of length 100 days (the maximum lead time) for each initial condition, and we report their average (Figure 6).

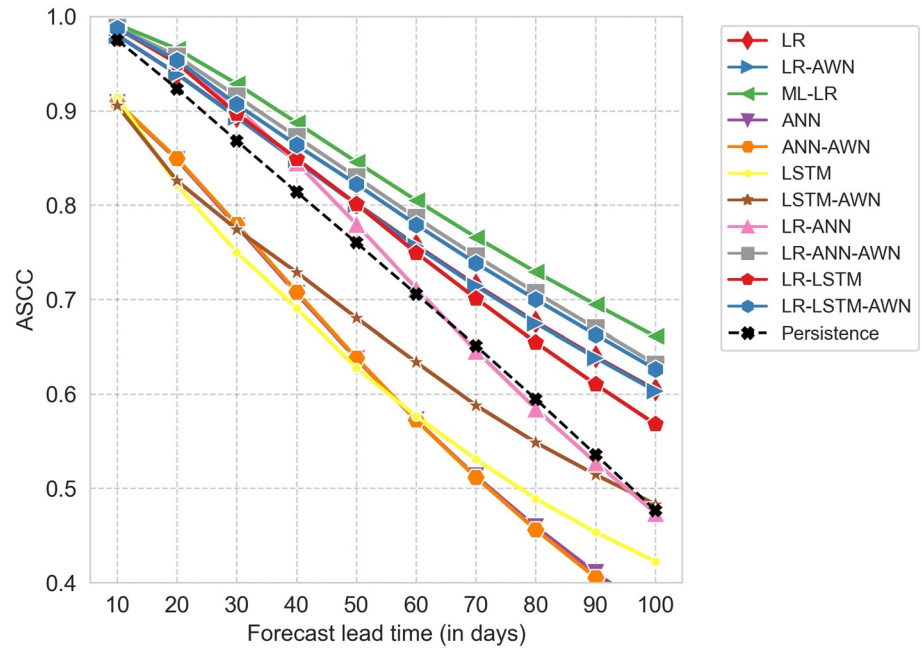


Figure 6. Comparison of ASCC for all models.

$$ATCC(x, y) = \left\langle \left(\frac{\psi_1^{forecast}(x, y, t) - \langle \psi_1^{forecast} \rangle(x, y)}{\sigma(\psi_1^{forecast})(x, y)} \right) \cdot \left(\frac{\psi_1^{ref}(x, y, t) - \langle \psi_1^{ref} \rangle(x, y)}{\sigma(\psi_1^{ref})(x, y)} \right) \right\rangle, \quad (7a)$$

$$ASCC(t) = \left(\frac{\overline{\psi_1^{forecast}(x, y, t) - \overline{\psi_1^{forecast}}(t)}}{\sigma(\psi_1^{forecast})(t)} \right) \cdot \left(\frac{\overline{\psi_1^{ref}(x, y, t) - \overline{\psi_1^{ref}}(t)}}{\sigma(\psi_1^{ref})(t)} \right), \quad (7b)$$

where t is the forecast lead time, (\cdot) and $\langle \cdot \rangle$ indicate the spatial and temporal averages, respectively, and σ refers to the standard deviation. Essentially, ATCC conveys temporal similarity between the forecasted and reference truth over all lead times, as it computes the grid-point-wise temporal correlation between the time series of the two data sets averaged for all initial conditions. In contrast, ASCC exhibits the spatial structure similarity between the two data sets, as it computes the snapshot-wise correlation between the two fields at a given lead time, normalized and then averaged for all initial conditions.

The ATCC maps show that the ML-LR forecasts are the best, followed by the stochastic hybrid models (Figures 5c, 5i and 5k). For most of the models, the correlations are generally higher in the gyre regions than in the eastward jet region, which is justified since the latter area is more turbulent. However, the stochastically improved hybrid methods and ML-LR provide a higher correlation in the jet region than those of the non-stochastic hybrid models (Figures 5h and 5j). The pure deep-learning methods and their stochastic extensions (Figures 5d–5g) fail to reproduce the variabilities entirely, thus, resulting in significant basin-wide dissimilarity with the reference data set. The ATCC maps of LR and its white noise extension (Figures 5a and 5b) show similar basin-wide correlations, and these are similar to the deterministic hybrid models (Figures 5h and 5j). However, small correlations along the jet region are more pronounced in LR-AWN.

A comparison of the ASCC time series (Figure 6) is telling a similar story. The pure deep-learning methods and their stochastic augmentations show the worst structural similarity with the reference truth (lower than the persistence at all lead times), whereas ML-LR provides the highest correlation, followed by the noise-augmented hybrid models. An inspection of the correlation decay rates of the models involving LR suggests that, on average, all of them possess a similar decorrelation rate, but the deterministic hybrid models decay faster than the others on longer lead times with LR-ANN decaying similar to the Persistence

baseline. This suggests that introducing noise in the hybrid models improves both temporal and spatial characteristics.

Overall, we conclude that, on short forecast time scales, both ML-LR and noise-augmented hybrid methods are most realistic regarding the spatial and temporal evolution. The pure deep-learning models and their stochastic extensions perform poorly.

4.3. Climatology and Variance

Here, we diagnose the mean and variance of the forecasted streamfunction field along the temporal domain:

$$\langle \psi_1^{forecast} \rangle = \frac{1}{N} \sum_{n=1}^N \psi_1^{forecast}(x, y, t_n), \quad (8a)$$

$$\sigma^2(\psi_1^{forecast}) = \frac{1}{N} \sum_{n=1}^N \left(\psi_1^{forecast}(x, y, t_n) - \langle \psi_1^{forecast} \rangle \right)^2. \quad (8b)$$

The mean field $\langle \psi_1^{forecast} \rangle$ is also referred to as the “climatology.” However, unlike the operational forecasts, the reference climatology is zero as we have performed an SVD of the mean-subtracted streamfunction field. Therefore, the reconstructed streamfunction $\psi_1^{forecast}$ should show a climatology of zero. We used these diagnostic metrics to characterize long-timescale forecasts, over 20K records, that is, 200K days or nearly 547 years. As results are independent of the initial conditions, we perform simulations from a single initial condition and use only one stochastic realization, wherever applicable. Due to a much longer record of the reference streamfunction compared to the forecast length, we expect a small nonzero value of the predicted time-mean streamfunction field (see Figure 7a for reference data set of the same length), and its value can serve as a reference for the temporal bias introduced by different models. On the other hand, the variance map would characterize the extent of jet reproduction by different methods, as it is the most turbulent and possesses maximum fluctuations in the entire domain (Figure 8a).

The climatology maps suggest that the LR shows the least temporal bias among all methods (Figure 7b). However, this is because the LR output decays to zero after a short lead time, and the PCs exhibit near-zero mean despite the poor forecasts. The multi-level formalism (Figure 7d) shows the second smallest bias, with stable and non-zero forecasts at all lead times. The AWN extension of LR produces the next overall small bias among the LR and its stochastic extensions (Figure 7c).

All standalone deep-learning methods (Figures 7e–7h) produce a relatively large bias in the forecasted streamfunction field. For ANN, the deterministic version produces a higher bias than the stochastic one, whereas LSTM produces a high bias irrespective of the noise. The large bias is primarily because these methods produce a significant drift in the modeled PCs, and this results in a large non-zero temporal mean which also reflects in the reconstructed streamfunction field after multiplying with the EOFs. Among the hybrid models (Figures 7i–7l), the ANN hybrid models generate a smaller basin-wide bias than the LSTM hybrids, and the stochastically improved hybrid models produce a smaller bias than their deterministic counterparts. The latter suggests that the induced drift in the modeled PCs can be contained to some extent by adding stochasticity, which nudges the PCs back toward the reference truth trajectory.

Analyzing the spatial map of grid-point-wise temporal variance of the model outputs (Figure 8), we found that LR-AWN, ML-LR, and the stochastically augmented hybrid models (Figures 8c 8d, 8j and 8l) best reproduce the reference jet variability (Figure 8a). The non-stochastic ANN-hybrid model (Figure 8i) produces correct but extra stretched jet variability in the north-south direction and around the eastward extension, whereas the deterministic LSTM-hybrid model (Figure 8k) produces an irregular and overly narrow region of jet variability. Among the standalone deep-learning models (Figures 8e–8h), the deterministic versions of both ANN and LSTM produce insufficient jet variabilities, with ANN being worse than LSTM. But, when augmented with white noise, both produce more variabilities along the jet region and are nearer to the reference truth, especially ANN-AWN. The LR fails to produce any variability in the basin (Figure 8b) as its outputs decay to zero after a short forecast lead time.

Combining the climatology and variance results, we conclude that (a) on long time scales, LR-AWN, ML-LR, and the stochastically augmented hybrid methods best reproduce the jet variabilities with a small drift

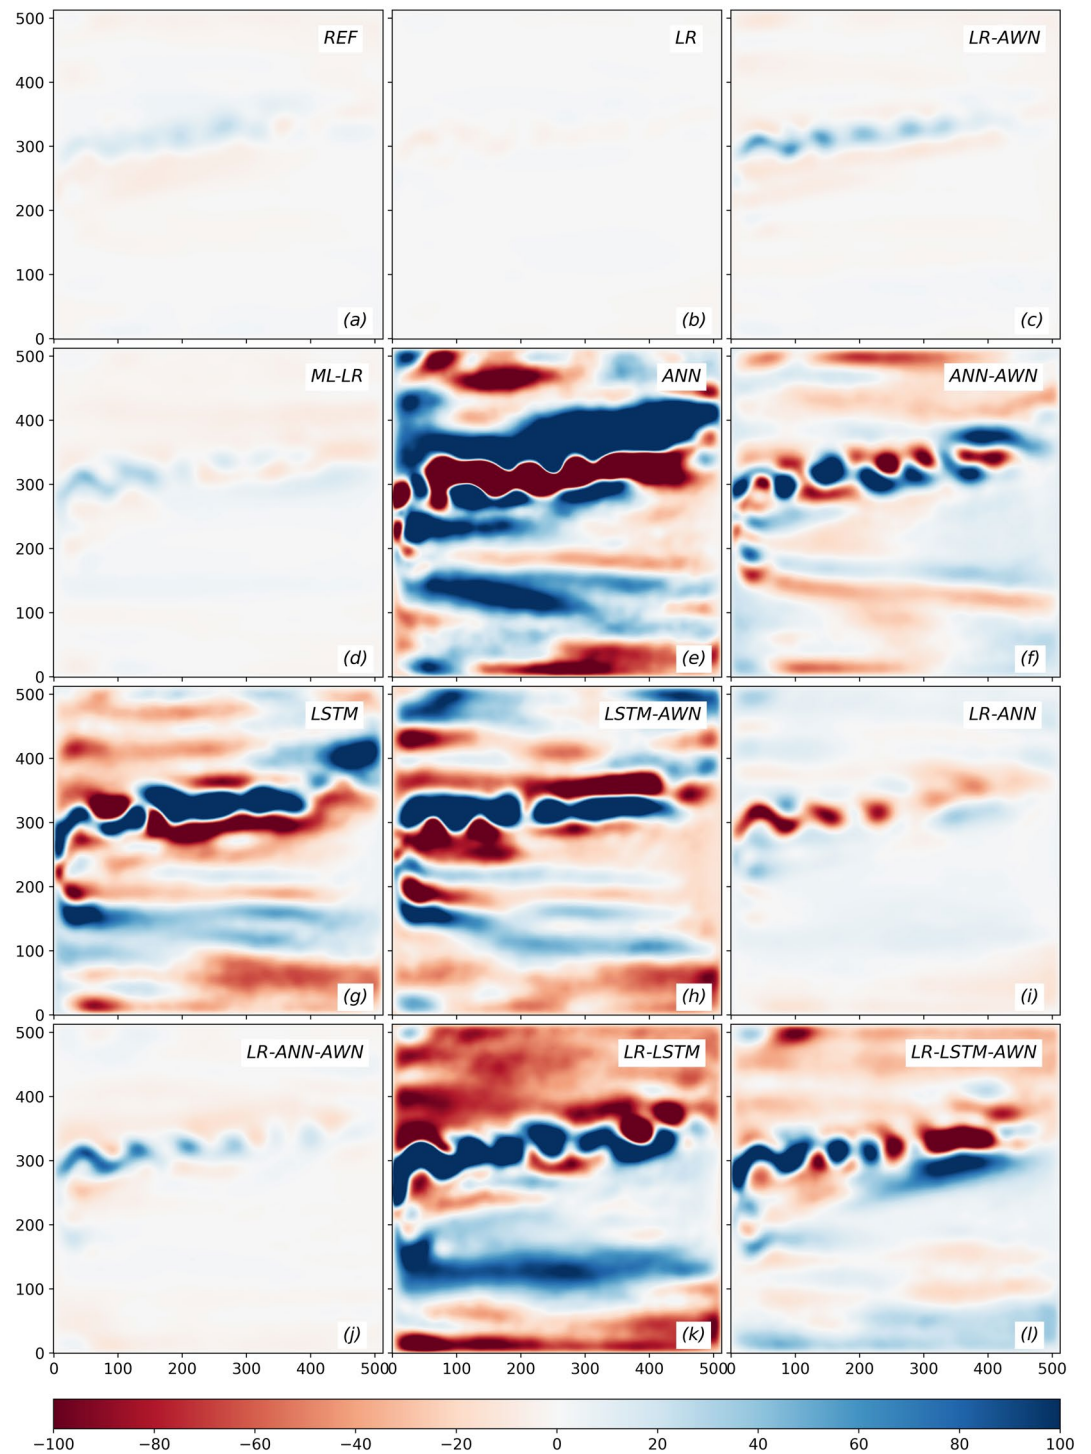


Figure 7. The climatology for 200K days randomly chosen data sample from the 500K days reference data set (REF; top left panel); and the 200K days forecasts produced by different models. The range of the climatological bias $[-100, 100]$ should be compared to the range $[-10^3, 10^3]$ of the time-mean state of streamfunction, see Figure 1b for a streamfunction snapshot; non-dimensional units.

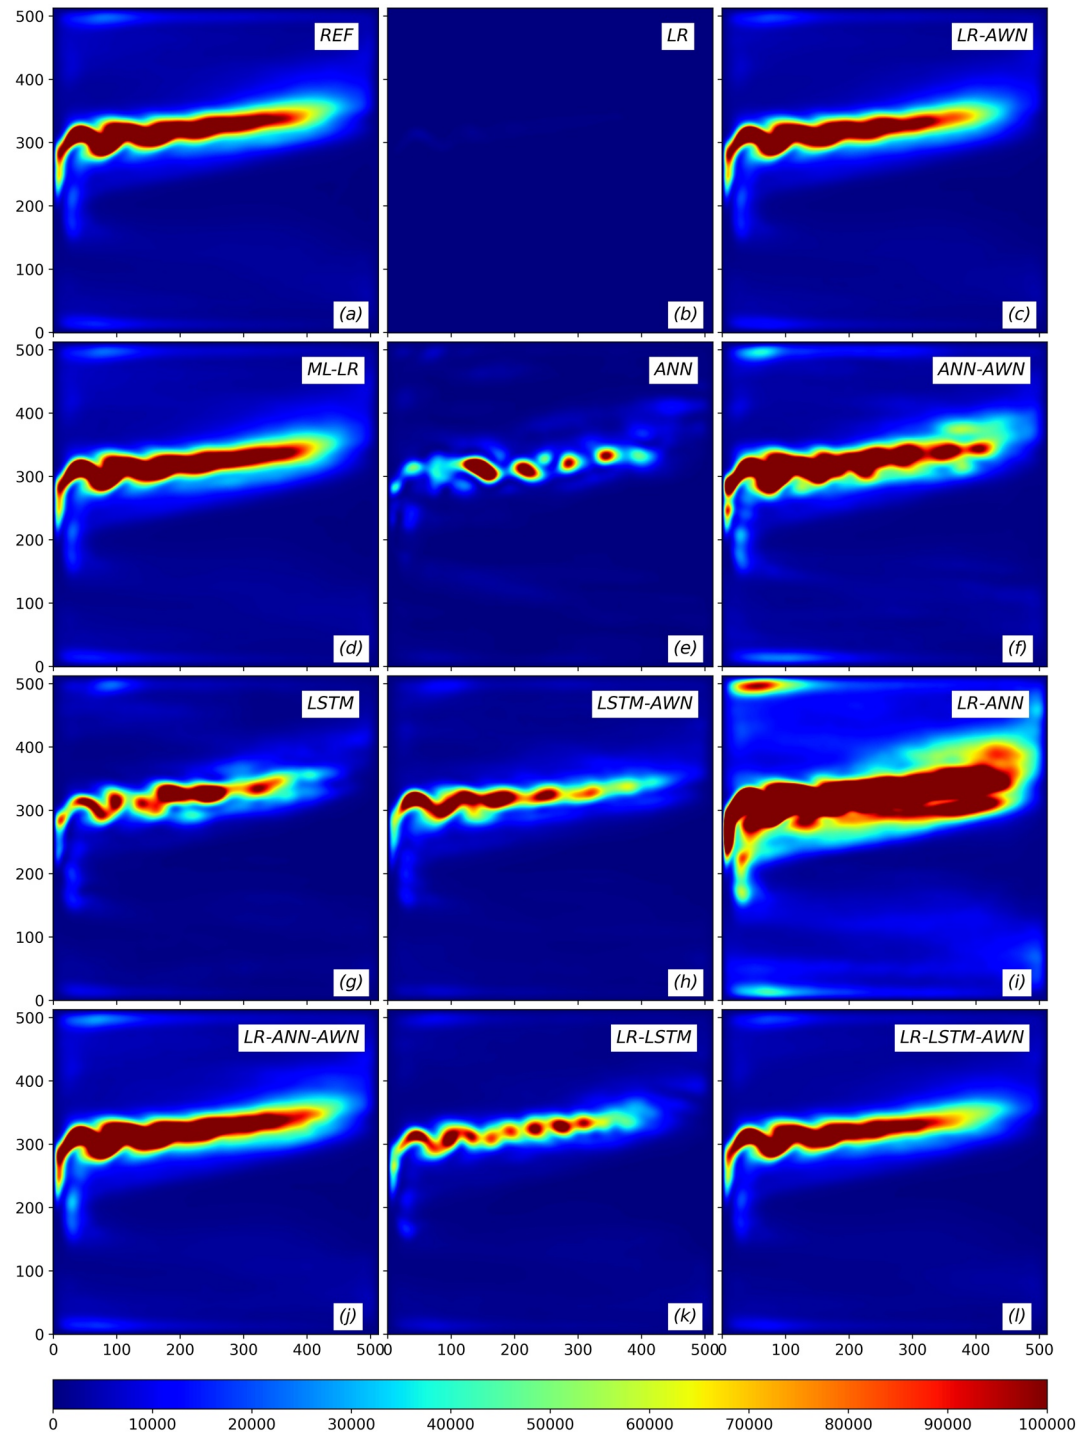


Figure 8. The temporal variance for 200K days randomly chosen data sample from the reference data (REF; top left panel) and 200K days forecasts produced by different models. Like climatology, the variance is also in non-dimensional units.

in the resulting flow field (approx. 2 – 3% of the original mean field); (b) the standalone deep-learning implementations infuse a relatively higher bias in the climatology and are inefficient at reproducing the turbulent jet characteristics on long time scales; (c) the LR model cannot produce climate-like forecasts due to its dissipative nature, and therefore all other models fare better than it on long timescales.

4.4. Frequency Map

Here, we consider *frequency maps* of the emulated long-timescale solutions (same as the one used in the previous section) to quantify their spectral frequency characteristics. For each model output, the frequency map is obtained by diagnosing the frequency value locally (i.e., for each grid cell), as given by the inverse of the decorrelation time scale of the forecasted streamfunction field. The decorrelation time scale is determined as the lag at which the autocorrelation drops by a factor of e from the zero-lag value. Because we repeat this calculation for each spatial location, we get a gridded frequency map of the size $m \times m$. Overall, we expect higher-frequency variability along the eastward jet and in boundary regions, due to the vigorous eddy activities, and low-frequency variabilities elsewhere. This is demonstrated in the reference frequency map (Figure 9a) obtained for a randomly chosen 200K days long data sample from the reference data set.

The results suggest that ML-LR most closely resembles the reference frequency map followed by LR-AWN and the two stochastic hybrid models (Figures 9c, 9d, 9j and 9l). While ML-LR produces the correct frequency patterns in the gyre regions, the magnitude is lower in the jet region when compared to the reference truth. LR-AWN and the stochastic hybrid models also reproduce the frequencies in the gyre regions but are less accurate in the jet region, with the frequency magnitude even lower than ML-LR in this region.

On the other hand, the deterministic hybrid models (Figures 9i and 9k) fail to describe the characteristic frequencies throughout the domain and produce frequency maps mostly dominated by low frequencies—more so for the deterministic ANN hybrid model. The same is also true for the standalone ANN implementations (Figures 9e and 9f) with the stochastic one being better as they show higher frequencies in the domain but not following the correct pattern. Such low-frequency dominated maps suggest that the individual PC outputs lack correct high-frequency contributions. The LSTM-only models also produce incorrect frequency maps irrespective of the noise (Figures 9g and 9h). The deterministic LSTM variant exhibits patterns of low and intermediate frequencies in the entire basin, whereas the stochastic variant produces patches of high frequencies both in the jet and gyre regions and does not resemble their reference truth. This is because the deterministic LSTM outputs (for all 150 PCs) are dominated by low frequencies, and adding noise to them produces somewhat large-amplitude (therefore, higher variance) but much high-frequency-dominated outputs for most of the PCs, which ultimately leads to high frequencies in both the jet and the gyres regions when multiplied by the EOFs.

LR completely misses the high-frequency variability around the jet (Figure 9b), as the solution decays to zero, although it still manages to reproduce the low-frequency variations in the two gyres to a certain extent.

We conclude that (a) as seen for the climatology and variance, LR-AWN, ML-LR, and the stochastic hybrid models perform best regarding frequency characteristics of long-timescale solutions; (b) the deterministic hybrid models fail to correctly reproduce the frequency content despite their low climatological bias and nearly correct variance pattern; (c) standalone deep-learning methods produce the most inaccurate and physically unjustified frequency maps, especially ANN.

4.5. Forecast Horizon

Here, for each model, we estimate the *forecast horizon* which is the time scale for which the model produces stable and non-zero forecasts. This information is vital for deciding on the applicability of a method in short-/long-term forecasts. When the system matrix is available, as in the LR, the forecast horizon is given by the inverse of its maximum eigenvalue. For the other methods, it is computed using the model outputs, which, for long-timescale forecasts, either saturate to a steady-state value or provide non-zero and stable solutions up to the maximum lead time (equal to 20K records, i.e., 200K days or 547 years). In the first case, the forecast horizon for each PC is given by the time beyond which the solution is trapped in a small-amplitude range, that is, a small threshold value. The overall forecast horizon is, then, set equal to the minimum of the individual PC horizons. In the second case, the model forecast horizon is set to ∞ because such models can produce stable non-zero forecasts on any finite lead time.

Only LR belongs to the first category, whereas all the other methods belong to the second (Table 1, column 2). LR is a dissipative deterministic model, and its outputs decay to zero around its forecast horizon. Therefore, this model is not suitable for climate-type applications. It is also worth noting that all stochastically

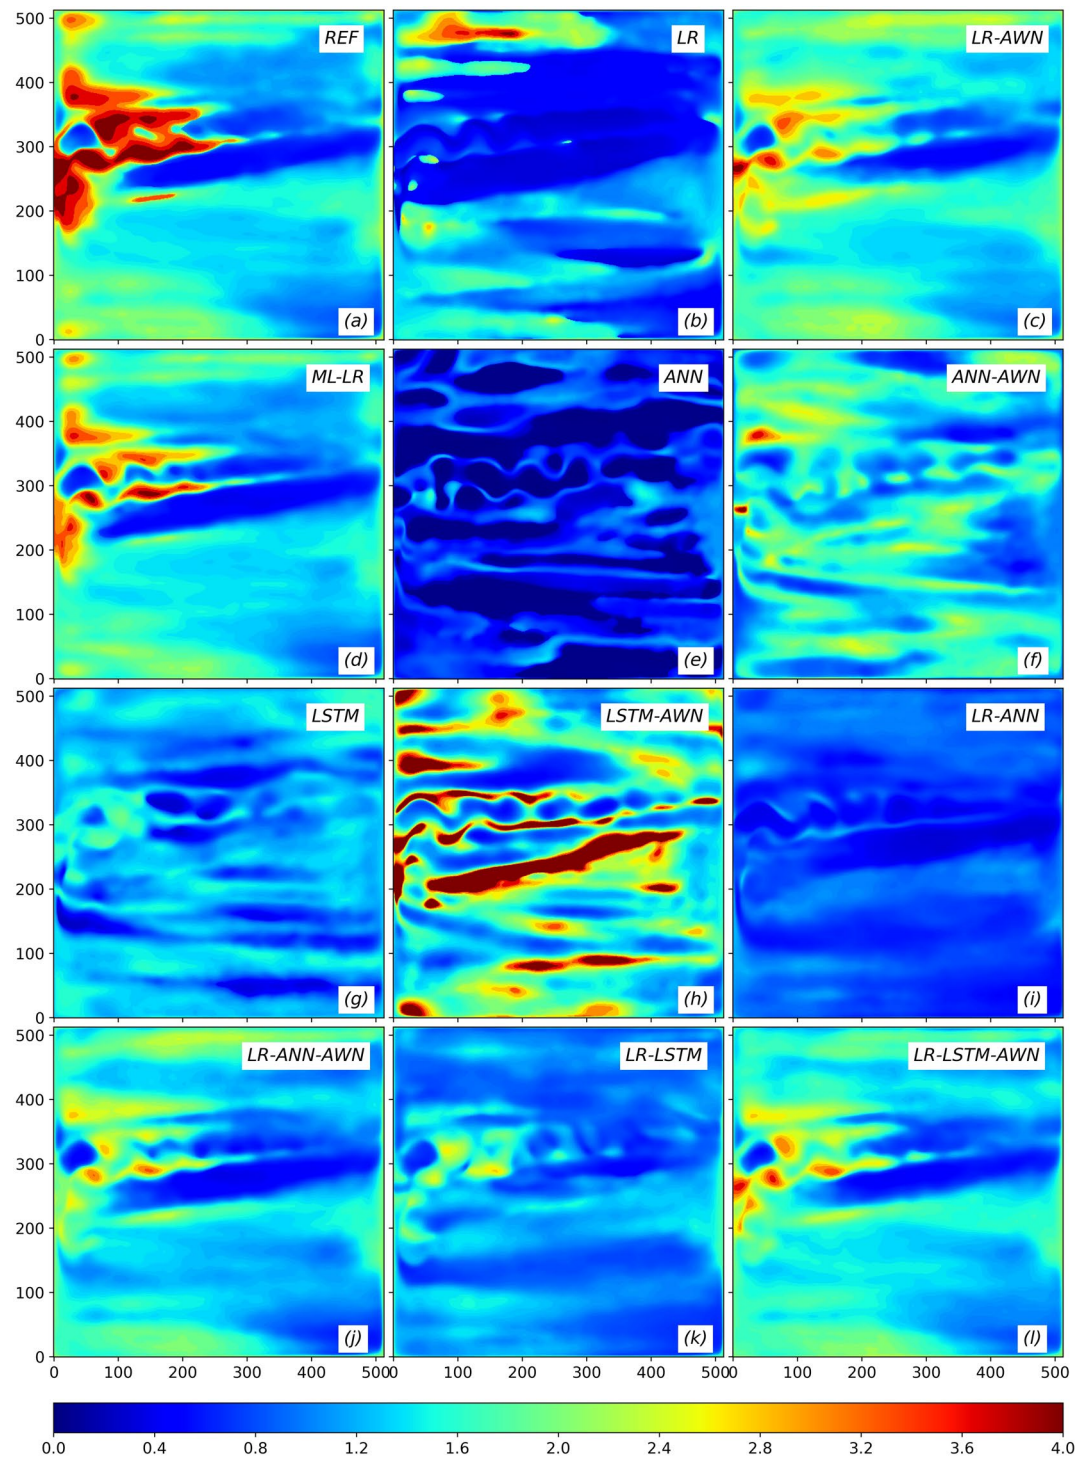


Figure 9. Frequency maps for 200K days long randomly chosen data sample from the reference data (REF; top left panel) and for the 200K days forecasts produced by different models. The units of frequency are in cycles/year.

Table 1
Forecast Horizon and Running Time Complexity Analyses for All the Models

Method	Forecast horizon (in years)	Training time (in seconds)	Prediction time (in seconds)
LR	1.3	1	10^{-1}
LR-AWN	∞	1	10^{-1}
ML-LR	∞	10	1
ANN	∞, ∞	10^3	1
LSTM	∞, ∞	10^4	10^2
LR-ANN hybrid	∞, ∞	10^3	1
LR-LSTM hybrid	∞, ∞	10^3	10^2

Note. The two values of the forecast horizon for artificial neural networks (ANN), long short term memory (LSTM), and the hybrid models represent these values for their deterministic and stochastic variants, in the same order. However, their training and running times are of the same order of magnitude, so we have reported them only once. Both training and prediction times are reported for the codes executed on a 24 core machine with 64 GB memory. For deep-learning-based methods, the training time is sensitive to the number of epochs and the batch size. Here, we have reported them for 200 epochs and the default batch size 32. Note, however, that some of the deep-learning-based methods, for example, ANN, achieved optimal training in a smaller number of epochs.

forced models possess an infinite forecast horizon. This implies that using random variables alongside the PCs ensures the injection of energy (in the form of small-scale variabilities) that prevents decay of the solutions and, at the same time, keeps the solution bounded and stable.

4.6. Training and Prediction Time Complexity

Here, we discuss computational costs and scalability as the number of degrees-of-freedom increases. For each model, we have diagnosed the run times for training and inference and refer to them as the “training time complexity” and “prediction time complexity,” respectively. In practice, the training is done only once, whereas the predictions are made many times, therefore, it makes sense to look at their time complexities separately. Note that the complexity estimates should take into account different levels of possible optimization of the models. Therefore, we wrote all model codes in Python (except ML-LR that is in a publicly available Matlab Toolbox), implemented them on the same hardware, and optimized them to reasonably high levels, including vectorization, function calls, and efficient data structures. The training complexity estimates do not include data processing and variable declaration/initialization and only correspond to the time taken for training the models. The prediction time complexity corresponds to the time taken for producing one realization of a 1,000 records long forecast. Since we can only provide estimates for the optimal performance of different methods, we only report the orders of magnitude of the elapsed time (Table 1, column 3 and 4).

Among the stochastic augmentations of LR, ML-LR is one order of magnitude more expensive to train and forecast compared to LR and its additive white noise extension. The higher complexity of ML-LR is simply due to the extra regression layer. Similar training and prediction times of LR and LR-AWN are due to the same trained core, that is, LR coefficients; the white noise parameters in LR-AWN are inexpensive to train.

Mathematically, for LR and its white noise stochastic extension, as the number of PCs (say, n) increases, we expect the training time complexity to increase as $\mathcal{O}(n^2)$, which is the size of the trained regression matrix, and the prediction time complexity to increase as $\mathcal{O}(n)$. For ML-LR, the training cost of each level l (where $1 \leq l \leq L$) is $\mathcal{O}(n^2l)$, and therefore the total training cost is proportional to $n^2L(L+1)/2$, that is, an $\mathcal{O}(n^2L^2)$ scaling; the prediction cost of ML-LR is expected to scale as $\mathcal{O}(n+l)$, or $\mathcal{O}(n)$ assuming $n \gg l$.

In the standalone deep-learning class, LSTM is an order of magnitude more expensive to train than ANN, but two orders of magnitude costlier to produce forecasts. The higher computational cost of LSTM is due to its more complex architecture and a large number of past states passed as input—set by the “look back” hyperparameter. However, in the hybrid category, both ANN- and LSTM-based methods follow the same order of training time, but the difference in their prediction cost is the same—that is, two orders of magnitude. A modest difference between the training cost of ANN and LSTM in the hybrid design is due to their simpler architectures and a smaller number of past states required by LSTM to predict the next state. Nevertheless, in both categories, both ANN and LSTM are 1–3 orders of magnitude cheaper to produce forecasts than to train them. This makes them suitable for climate-related applications, as the training is done once, but the predictions are obtained numerous times.

We do not compare the time complexity of the deep-learning-based methods against the LR and its stochastic extensions because the time estimates for both ANN and LSTM depends on their hyper-parameter values, mainly *Epochs* and *Batch size* (they were kept the same for all ANN- and LSTM-based models), which are bound to change for different applications. *Epochs* determine the number of passes through the entire training data set needed to optimize the model parameters; *Batch Size* refers to the number of training samples which need to be parsed before updating the model parameters. Due to these hyper-parameters, with a

Table 2
An Overview of All the Methods Considered in This Work

Method	Abbreviation	Section	Input	Output	Cost function	Memory	State-dependent noise
Linear Regression	LR	3.1	State	tendency	OLS	NA	NA
LR + Additive White Noise	LR-AWN	3.2	State	tendency	OLS	×	×
Multi-level Linear Regression	ML-LR	3.3	state, LR residuals	tendency	OLS	✓	✓
Artificial Neural Network (+White Noise)	ANN (- AWN)	3.4	State	State	MAE	×	×
Long Short Term Memory (+White Noise)	LSTM (- AWN)	3.5	State	State	MAE	✓	×
LR + ANN Hybrid (+White Noise)	LR-ANN (-AWN)	3.6	state, LR residuals	tendency	OLS, MAE	×	✓
LR + LSTM Hybrid (+White Noise)	LR-LSTM (-AWN)	3.6	state, LR residuals	tendency	OLS, MAE	✓	✓

Note. Method, abbreviation, the section in which they are explained, input, output, cost function, presence of memory effects (in model definition/construction), and state dependency/independency of noise (only applicable for stochastic models). OLS and MAE stand for ordinary least square and mean absolute error, respectively; NA means not applicable.

few others such as the optimizer, loss function, activation function, etc., and the black-box nature of NNs, it is hard to determine how exactly the running time complexities scale with n .

Overall, we conclude that (a) adding simple stochasticity bears a negligible computational cost, but a more complicated red-type noise addition can increase the training and prediction cost by order of magnitude; (b) Among the deep learning methods, LSTM is equal or more costly to train and forecast than ANN; (c) Both LSTM and ANN can benefit from reduced training time in the hybrid framework compared to its standalone implementation, as a less complex network design is required for optimal performance.

5. Discussion and Conclusion

We have presented a comprehensive inter-comparison of LR, its stochastic extensions, deep-learning models (ANN and LSTM), and their hybrid formulations with additive-noise (see Table 2 for an overview), to obtain a low-cost, reduced-order model for the complex multi-scale spatio-temporal flow of the upper ocean. LR has the simplest form and, thus, provides a baseline for assessing the performance of the other models. The obtained results show that linear models augmented by state-dependent noise and memory effects, either through multi-level regression or deep learning, perform the best across our metrics and tasks to emulate and predict very complex, nonlinear, and multi-scale ocean flow. CNN is another proven methodology in the deep-learning class for emulating image-based data sets but is not considered here as it is well-studied elsewhere and is beyond the scope of this study as it would be most suitable for structured two-dimensional grids rather than for the application in EOF space. It would therefore require a different data set likely with a much larger number of degrees-of-freedom.

A solution of the idealized, eddy-resolving, double-gyre, quasigeostrophic ocean circulation model is used as the reference data set. The testbed for the reduced-order modeling consists of the leading 150 EOFs/PCs of the reference solutions (out of a total 50,000), and these capture about 90% of the total variance. The higher-ranked EOFs show multi-scale variability along the turbulent eastward jet region, while the low-ranked ones show small-scale variations in the entire domain; the corresponding PCs exhibit structured (only for the top few) and noisy patterns, referring to a mixture of low- and high-frequency variability. We have modeled the PCs' dynamics using different methodologies, obtained forecasts, reconstructed the physical spatio-temporal fields using the EOFs, assessed the results using several assessment metrics, and inter-compared model skills. Training of the models is done using a ≈ 1100 years long data set with 10 days as the sampling period. We have studied forecasts with both short- and long-timescales, where the lead times are on seasonal and centennial time scales, respectively. To assess the accuracy of short-time forecasts, we have used RMSE and anomaly cross correlations (ACC), and, for long-timescale forecasts, we used climatology, variance, and frequency maps. Additionally, we have diagnosed the stability and computational costs of the

methods using forecast horizon and training/prediction time complexities, respectively. It is also possible to define a few dynamically inspired performance metrics, such as “product integral” studied by Agarwal et al. (2021), for future eddy emulators, but developing and applying such metrics is beyond the scope of this study.

We have made the following key observations during the assessment:

1. On short forecast lead times (e.g., several months), multi-level LR (ML-LR) delivers the best results, whereas the standalone deep-learning methods, both in the presence and absence of the noise, perform the worst, as evidenced by a higher RMSE and poor ACC's both in space and time. The hybrid methods with additive noise are the next best after ML-LR; the ANN-hybrid model performs marginally better than the LSTM-hybrid. The success of ML-LR and noise-augmented hybrid models reflects the importance of memory effects and stochastic noise when accounting for dynamical interactions with the removed EOFs/PCs (i.e., beyond the rank 150). Including these effects, accounting for the truncated dynamics is rigorously justified in the Mori–Zwanzig theory of statistical mechanics (Kondrashov et al., 2015), which started also to attract attention in deep learning (Wang et al., 2020). The ML-LR conveys memory effects using an additional LR level, leading to a red-type noise, whereas LSTM achieves this by definition. The state dependency of noise is achieved using additional regressions in ML-LR, whereas in hybrid models this is accomplished using state variables as inputs alongside the LR residuals to forecast the residuals for the next time instant. Additionally, both models explicitly resolve the linear dependency, which is vital because pure ANN/LSTM (with or without noise) display performance even worse than persistence for short lead times. The linear dynamics is pertinent to the gyres and must explain a significant proportion of the variance in the top PCs. Representing it using linear regression (a) determines the coefficients accurate up to the machine precision, (b) leaves a lesser number of terms to be learned more, and (c) improves the signal-to-noise ratio in the residual statistics on which ANN/LSTM is later trained to a higher effect. Therefore, we argue that the use of bare ANN/LSTM is not useful in situations where linear terms dominate that can be learned via regression (similar to Pyle et al., 2021). In such cases, it is better to use the ANN/LSTM as a correction (potentially stochastic) term in combination with LR to get an optimal closure model for the dynamics of the retained EOFs.
2. For long-timescale forecasts (e.g., several centuries), the white-noise extension of LR, ML-LR, and stochastic hybrid models perform the best as they correctly resolve both low- and high-frequency variabilities across the domain (with the right frequency magnitude) and as they produce low climatological bias. Stochastic ANN-hybrid produces a lower climatological bias than its LSTM counterpart, but the latter produces a relatively better frequency map; the variance estimates are almost the same for both. LR outputs decay to zero, whereas all standalone deep-learning methods (with or without noise) generate strong climatological drifts and fail to represent the flow variability both in the jet and the gyre regions.
3. All models show a better forecast horizon than the LR. In particular, all the noise-augmented models produce stable, non-zero, and finite forecasts on climate-like lead times, for example, centuries and millennia, meaning that adding noise to deterministic models alongside the system variables boosts its stability while keeping the forecasts bounded and physically relevant.
4. Simple white noise extension of LR shows a similar training and prediction time complexity as the LR model, but they become more expensive for more complex architecture, such as the ML-LR. Deep-learning-based methods are the most expensive to train but relatively cheaper during forecasts—a potentially useful property from a climate forecasting viewpoint. Note that LSTM models are more expensive to train and forecast than the others, so it is better to use them in the hybrid configuration, as they benefit from a simpler model configuration, and thus fewer trainable parameters and lesser training (and forecast) time.

Overall, our results prove the superiority of ML-LR and stochastically augmented hybrid models for building simple, stable, and low-cost reduced-order ocean emulators (within the EOFs/PCs framework) for producing short-/long-timescale forecasts. The success of these methods highlights the importance of including core dynamics, memory effects, and model errors for building reliable emulators. In this application, we have considered the core dynamics as linear and concentrated on the latter two components to prove that state-dependent additive noise is an excellent way to account for memory and unknown dynamical model

errors in emulators. ML-LR allows for only linear additive state-dependent noise (and memory), but hybrid deep-learning models can potentially learn very general forms of non-linear and multiplicative state-dependent noise. Similar outcomes of ML-LR and hybrid deep-learning models prove that ML-LR produces the most optimal noise configuration which deep learning learned successfully. Another evidence of the importance of state-dependent noise is the poor performance of a purely red noise augmentation of LR (not shown), which has memory but no state-dependency for the noise.

However, as for any data-driven method, the results presented here are valid for the current training length, and we admit that a longer training set may improve the models' performance and even change how these models compare. However, using a longer training length can be computationally prohibitive, and too-long ocean observations may not be available in real life. Therefore the emphasis here is also on identifying the models that perform relatively well even on shorter training data. The use of orthogonal bases, that is, EOFs/PCs, puts another constraint on the current study, as many fluid dynamical systems may not follow this assumption, but using and comparing different bases transformations, for example, dynamical mode decomposition (DMD) modes, is beyond the scope of the study and is left as a future exercise.

The current research can be extended along the following lines. The first and also straightforward direction is to test the performance of the best-performing stochastic models on a more complicated testbed, for example, on the output of a comprehensive ocean general circulation model or coupled ocean-atmosphere models for emulating, say, ENSO or Madden-Julien Oscillation, which incorporate significant delay time response. Such an implementation would further ascertain our findings for building reliable emulators for complicated ocean/atmospheric processes. It is worth looking for the ways of imposing physical constraints, such as energy/mass conservation, into the emulators, for example, using an appropriate penalizing term in the loss function. Second, the results obtained here are directly relevant for emulation of various complex and multi-scale fields in the context of eddy parameterizations and test the alternative definitions of eddies investigated recently (Agarwal et al., 2021; Berloff et al., 2021). Finally, a possible sequel to this work is including more stochastic and deep-learning methods, or a mixture of both, for example, the Stochastic Neural Networks (Guillaumin & Zanna, 2021). We started to develop a rigorous testbed for data-driven models and used this for several model setups, but we will expand this to more complex testbeds, models, and data sets in the future and check if the conclusions still hold.

Acknowledgments

The authors are thankful to S. Juricke, P. Hassanzadeh, and two other anonymous reviewers for their feedback and suggestions, which improved the manuscript significantly. Niraj Agarwal is grateful to the Research Computing Service (RCS) team of Imperial College London for the help and assistance with HPC, and to MPE CDT for providing the financial and technical support for conducting this research. P. Dueben gratefully acknowledges funding from the Royal Society for his University Research Fellowship as well as the ES-iWACE, MAELSTROM, and A14Copernicus under Horizon 2020 and the European High-Performance Computing Joint Undertaking (JU; grant agreement No 823988, 955513 and 101016798). The JU receives support from the European Union's Horizon 2020 research and innovation programme and the United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway. P. Berloff and E. Ryzhov are supported by NERC grant NE/R011567/1 and Royal Society Exchange Grant IEC/R2/181033. D. Kondrashov is supported by the National Science Foundation grant OCE-1658357. Analyses in Sections 3.3 and 3.6 were supported by the Russian Science Foundation (Grant No. 18-12-00231). The part of the study conducted by E. Ryzhov that involved producing data with varying spatial resolutions of the baroclinic ocean flow model for further analysis was partly supported by the Russian Science Foundation (Project 19-17-00006). P. Berloff gratefully acknowledges the support by NERC grant NE/T002220/1 and Leverhulme grant RPG-2019-024. P. Berloff was also supported by the Moscow Center of Fundamental and Applied Mathematics (supported by the Agreement 075-15-2019-1624 with the Ministry of Education and Science of the Russian Federation).

Data Availability Statement

ML-LR Toolbox is available at research.atmos.ucla.edu/tcd/dkondras/Software.html. The source codes for the other methods and the data set used can be found in the figshare repository: <https://doi.org/10.6084/m9.figshare.14188349.v4>.

References

- Agarwal, N., Ryzhov, E., Kondrashov, D., & Berloff, P. (2021). Correlation-based flow decomposition and statistical analysis of the eddy forcing. *Journal of Fluid Mechanics*, 924. <https://doi.org/10.1017/jfm.2021.604>
- Berloff, P. (2015). Dynamically consistent parameterization of mesoscale eddies. Part I: Simple model. *Ocean Modelling*, 87, 1–19. <https://doi.org/10.1016/j.oceomod.2014.12.008>
- Berloff, P., Dewar, W., Kravtsov, S., & McWilliams, J. (2007). Ocean eddy dynamics in a coupled ocean-atmosphere model. *Journal of Physical Oceanography*, 37(5), 1103–1121. <https://doi.org/10.1175/jpo3041.1>
- Berloff, P., Ryzhov, E., & Shevchenko, I. (2021). On dynamically unresolved oceanic mesoscale motions. *Journal of Fluid Mechanics*, 920. <https://doi.org/10.1017/jfm.2021.477>
- Berloff, P. S. (2005). Random-forcing model of the mesoscale oceanic eddies. *Journal of Fluid Mechanics*, 529, 71–95. <https://doi.org/10.1017/s0022112005003393>
- Berloff, P. S., & McWilliams, J. C. (2003). Material transport in oceanic gyres. Part III: Randomized stochastic models. *Journal of Physical Oceanography*, 33(7), 1416–1445. [https://doi.org/10.1175/1520-0485\(2003\)033<1416:mtiogg>2.0.co;2](https://doi.org/10.1175/1520-0485(2003)033<1416:mtiogg>2.0.co;2)
- Berner, J., Achatz, U., Batte, L., Bengtsson, L., Cámara, A. D. L., Christensen, H. M., et al. (2017). Stochastic parameterization: Toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3), 565–588. <https://doi.org/10.1175/bams-d-15-00268.1>
- Bolton, T., & Zanna, L. (2019). Applications of deep learning to ocean data inference and subgrid parameterization. *Journal of Advances in Modeling Earth Systems*, 11(1), 376–399. <https://doi.org/10.1029/2018ms001472>
- Brenner, M., Eldredge, J., & Freund, J. (2019). Perspective on machine learning for advancing fluid mechanics. *Physical Review Fluids*, 4(10), 100501. <https://doi.org/10.1103/physrevfluids.4.100501>
- Buizza, R., Milleer, M., & Palmer, T. N. (1999). Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Quarterly Journal of the Royal Meteorological Society*, 125(560), 2887–2908.

- Chattopadhyay, A., Hassanzadeh, P., & Subramanian, D. (2020). Data-driven predictions of a multiscale Lorenz 96 chaotic system using machine-learning methods: Reservoir computing, artificial neural network, and long short-term memory network. *Nonlinear Processes in Geophysics*, 27(3), 373–389. <https://doi.org/10.5194/npg-27-373-2020>
- Chattopadhyay, A., Mustafa, M., Hassanzadeh, P., & Kashinath, K. (2020). Deep spatial transformers for autoregressive data-driven forecasting of geophysical turbulence. In *Proceedings of the 10th International Conference on Climate Informatics* (pp. 106–112). <https://doi.org/10.1145/3429309.3429325>
- Chattopadhyay, A., Subel, A., & Hassanzadeh, P. (2020). Data-driven super-parameterization using deep learning: Experimentation with multiscale Lorenz 96 systems and transfer learning. *Journal of Advances in Modeling Earth Systems*, 12(11), e2020MS002084. <https://doi.org/10.1029/2020ms002084>
- Chen, C., Cane, M. A., Henderson, N., Lee, D. E., Chapman, D., Kondrashov, D., & Chekroun, M. D. (2016). Diversity, nonlinearity, seasonality, and memory effect in ENSO simulation and prediction using empirical model reduction. *Journal of Climate*, 29(5), 1809–1830. <https://doi.org/10.1175/JCLI-D-15-0372.1>
- D'Andrea, F., & Vautard, R. (2001). Extratropical low-frequency variability as a low-dimensional problem I: A simplified model. *Quarterly Journal of the Royal Meteorological Society*, 127(574), 1357–1374.
- Debussche, A., Glatt-Holtz, N., Temam, R., & Ziane, M. (2012). Global existence and regularity for the 3D stochastic primitive equations of the ocean and atmosphere with multiplicative white noise. *Nonlinearity*, 25(7), 2093–2118. <https://doi.org/10.1088/0951-7715/25/7/2093>
- DelSole, T. (2004). Stochastic models of quasigeostrophic turbulence. *Surveys in Geophysics*, 25(2), 107–149. <https://doi.org/10.1023/b:geop.0000028164.58516.b2>
- DelSole, T., & Hou, A. Y. (1999). Empirical stochastic models for the dominant climate statistics of a general circulation model. *Journal of the Atmospheric Sciences*, 56(19), 3436–3456. [https://doi.org/10.1175/1520-0469\(1999\)056<3436:esmftd>2.0.co;2](https://doi.org/10.1175/1520-0469(1999)056<3436:esmftd>2.0.co;2)
- Dueben, P. D., & Bauer, P. (2018). Challenges and design choices for global weather and climate models based on machine learning. *Geoscientific Model Development*, 11(10), 3999–4009. <https://doi.org/10.5194/gmd-11-3999-2018>
- Ewald, B., Petcu, M., & Temam, R. (2007). Stochastic solutions of the two-dimensional primitive equations of the ocean and atmosphere with an additive noise. *Analysis and Applications*, 5(02), 183–198. <https://doi.org/10.1142/s0219530507000948>
- Farrell, B. F., & Ioannou, P. J. (1993). Stochastic forcing of the linearized Navier-Stokes equations. *Physics of Fluids A: Fluid Dynamics*, 5(11), 2600–2609. <https://doi.org/10.1063/1.858894>
- Farrell, B. F., & Ioannou, P. J. (1995). Stochastic dynamics of the midlatitude atmospheric jet. *Journal of the Atmospheric Sciences*, 52(10), 1642–1656. [https://doi.org/10.1175/1520-0469\(1995\)052<1642:sdotma>2.0.co;2](https://doi.org/10.1175/1520-0469(1995)052<1642:sdotma>2.0.co;2)
- Franzke, C. (2012). Predictability of extreme events in a nonlinear stochastic-dynamical model. *Physical Review E*, 85(3), 031134. <https://doi.org/10.1103/physreve.85.031134>
- Franzke, C., Majda, A. J., & Vanden-Eijnden, E. (2005). Low-order stochastic mode reduction for a realistic barotropic model climate. *Journal of the Atmospheric Sciences*, 62(6), 1722–1745. <https://doi.org/10.1175/jas3438.1>
- Franzke, C. L. (2013). Predictions of critical transitions with non-stationary reduced order models. *Physica D: Nonlinear Phenomena*, 262, 35–47. <https://doi.org/10.1016/j.physd.2013.07.013>
- Ghil, M., Allen, M. R., Dettinger, M. D., Ide, K., Kondrashov, D., Mann, M. E., et al. (2002). Advanced spectral methods for climatic time series. *Review of Geophysics*, 40(1), 1–41. <https://doi.org/10.1029/2000rg000092>
- Ghil, M., Groth, A., Kondrashov, D., & Robertson, A. (2018). Extratropical sub-seasonal-to-seasonal oscillations and multiple regimes: The dynamical systems view. In A. W. Robertson, & F. Vitart (Eds.), *The gap between weather and climate forecasting: Sub-seasonal to seasonal prediction* (pp. 119–142). Elsevier.
- Glatt-Holtz, N., Ziane, M., & Ziane, M. (2008). The stochastic primitive equations in two space dimensions with multiplicative noise. *Discrete and Continuous Dynamical Systems Series B*, 10(4), 801–822. <https://doi.org/10.3934/dcdsb.2008.10.801>
- Guillaumin, A. P., & Zanna, L. (2021). Stochastic-deep learning parameterization of ocean momentum forcing. *Journal of Advances in Modeling Earth Systems*, e2021MS002534.
- Ham, Y.-G., Kim, J.-H., & Luo, J.-J. (2019). Deep learning for multi-year ENSO forecasts. *Nature*, 573(7775), 568–572. <https://doi.org/10.1038/s41586-019-1559-7>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Holden, P., Edwards, N., Garthwaite, P., Fraedrich, K., Lunkeit, F., Kirk, E., & Babonneau, F. (2013). PLASIM-ENTSem: A spatio-temporal emulator of future climate change for impacts assessment. *Geoscientific Model Development Discussions*, 6(2), 3349–3380.
- Holden, P. B., Edwards, N. R., Garthwaite, P. H., & Wilkinson, R. D. (2015). Emulation and interpretation of high-dimensional climate model outputs. *Journal of Applied Statistics*, 42(9), 2038–2055. <https://doi.org/10.1080/02664763.2015.1016412>
- Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science*, 304(5667), 78–80. <https://doi.org/10.1126/science.1091277>
- Jia, X., Willard, J., Karpatne, A., Read, J., Zwart, J., Steinbach, M., & Kumar, V. (2019). Physics guided RNNs for modeling dynamical systems: A case study in simulating lake temperature profiles. In *Proceedings of the 2019 Siam International Conference on Data Mining* (pp. 558–566). Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611975673.63>
- Juricke, S., Lemke, P., Timmermann, R., & Rackow, T. (2013). Effects of stochastic ice strength perturbation on arctic finite element sea ice modeling. *Journal of Climate*, 26(11), 3785–3802. <https://doi.org/10.1175/jcli-d-12-00388.1>
- Juricke, S., Palmer, T. N., & Zanna, L. (2017). Stochastic subgrid-scale ocean mixing: Impacts on low-frequency variability. *Journal of Climate*, 30(13), 4997–5019. <https://doi.org/10.1175/jcli-d-16-0539.1>
- Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., et al. (2017). Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10), 2318–2331. <https://doi.org/10.1109/tkde.2017.2720168>
- Karpatne, A., Watkins, W., Read, J., & Kumar, V. (2017). *Physics-guided neural networks (PGNN): An application in lake temperature modeling*.
- Karunasinghe, D. S., & Liang, S.-Y. (2006). Chaotic time series prediction with a global model: Artificial neural network. *Journal of Hydrology*, 323(1–4), 92–105. <https://doi.org/10.1016/j.jhydrol.2005.07.048>
- Kondrashov, D., & Berloff, P. (2015). Stochastic modeling of decadal variability in ocean gyres. *Geophysical Research Letters*, 42, 1543–1553. <https://doi.org/10.1002/2014GL02871>
- Kondrashov, D., Chekroun, M., & Berloff, P. (2018). Multiscale Stuart-Landau emulators: Application to wind-driven ocean gyres. *Fluids*, 3(1), 21. <https://doi.org/10.3390/fluids3010021>

- Kondrashov, D., Chekroun, M. D., & Ghil, M. (2015). Data-driven non-Markovian closure models. *Physica D*, 297, 33–55. <https://doi.org/10.1016/j.physd.2014.12.005>
- Kondrashov, D., Chekroun, M. D., Robertson, A. W., & Ghil, M. (2013). Low-order stochastic model and “past-noise forecasting” of the Madden-Julian oscillation. *Geophysical Research Letters*, 40, 5305–5310. <https://doi.org/10.1002/grl.50991>
- Kondrashov, D., Kravtsov, S., Robertson, A. W., & Ghil, M. (2005). A hierarchy of data-based ENSO models. *Journal of Climate*, 18(21), 4425–4444. <https://doi.org/10.1175/jcli3567.1>
- Kondrashov, D., Ryzhov, E., & Berloff, P. (2020). Data-adaptive harmonic analysis of oceanic waves and turbulent flows. *Chaos*, 30, 061105. <https://doi.org/10.1063/5.0012077>
- Krasnopolsky, V. M., & Fox-Rabinovitz, M. S. (2006). Complex hybrid models combining deterministic and machine learning components for numerical climate modeling and weather prediction. *Neural Networks*, 19(2), 122–134. <https://doi.org/10.1016/j.neunet.2006.01.002>
- Kravtsov, S., Kondrashov, D., & Ghil, M. (2005). Multi-level regression modeling of nonlinear processes: Derivation and applications to climatic variability. *Journal of Climate*, 18(21), 4404–4424. <https://doi.org/10.1175/jcli3544.1>
- Kravtsov, S., Kondrashov, D., & Ghil, M. (2009). Empirical model reduction and the modeling hierarchy in climate dynamics and the geosciences. In T. N. Palmer, & P. Williams (Eds.), *Stochastic physics and climate modeling* (pp. 35–72). Cambridge University Press.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097–1105).
- Liu, B., Tang, J., Huang, H., & Lu, X.-Y. (2020). Deep learning methods for super-resolution reconstruction of turbulent flows. *Physics of Fluids*, 32(2), 025105. <https://doi.org/10.1063/1.5140772>
- Liu, Y., Racah, E., Correa, J., Khosrowshahi, A., Lavers, D., Kunkel, K., et al. (2016). *Application of deep convolutional neural networks for detecting extreme weather in climate datasets*.
- Lorenz, E. N. (1956). *Empirical orthogonal functions and statistical weather prediction*.
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130–141. [https://doi.org/10.1175/1520-0469\(1963\)020<0130:dnf>2.0.co;2](https://doi.org/10.1175/1520-0469(1963)020<0130:dnf>2.0.co;2)
- Lorenz, E. N. (1996). Predictability: A problem partly solved. *Proceedings on the seminar on predictability*. ECMWF.
- Majda, A. J., Timofeyev, I., & Eijnden, E. V. (1999). Models for stochastic climate prediction. *Proceedings of the National Academy of Sciences*, 96(26), 14687–14691. <https://doi.org/10.1073/pnas.96.26.14687>
- Maulik, R., San, O., Rasheed, A., & Vedula, P. (2019). Subgrid modelling for two-dimensional turbulence using neural networks. *Journal of Fluid Mechanics*, 858, 122–144. <https://doi.org/10.1017/jfm.2018.770>
- McGovern, A., Lagerquist, R., Gagne, D. J., Jergensen, G. E., Elmore, K. L., Homeyer, C. R., & Smith, T. (2019). Making the black box more transparent: Understanding the physical implications of machine learning. *Bulletin of the American Meteorological Society*, 100(11), 2175–2199. <https://doi.org/10.1175/bams-d-18-0195.1>
- Mukhin, D., Kondrashov, D., Loskutov, E., Gavrilov, A., Feigin, A., & Ghil, M. (2015). Predicting critical transitions in ENSO models. Part II: Spatially dependent models. *Journal of Climate*, 28(5), 1962–1976. <https://doi.org/10.1175/jcli-d-14-00240.1>
- Nadiga, B. (2021). Reservoir computing as a tool for climate predictability studies. *Journal of Advances in Modeling Earth Systems*, 13, e2020MS002290. <https://doi.org/10.1029/2020MS002290>
- Nielsen, M. A. (2015). *Neural networks and deep learning* (Vol. 2018). Determination press.
- Ollinaho, P., Lock, S.-J., Leutbecher, M., Bechtold, P., Beljaars, A., Bozzo, A., et al. (2017). Towards process-level representation of model uncertainties: Stochastically perturbed parametrizations in the ECMWF ensemble. *Quarterly Journal of the Royal Meteorological Society*, 143(702), 408–422. <https://doi.org/10.1002/qj.2931>
- Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018). Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical Review Letters*, 120, 024102. <https://doi.org/10.1103/PhysRevLett.120.024102>
- Pawar, S., Ahmed, S. E., San, O., & Rasheed, A. (2020). Data-driven recovery of hidden physics in reduced order modeling of fluid flows. *Physics of Fluids*, 32(3), 036602. <https://doi.org/10.1063/5.0002051>
- Penland, C. (1989). Random forcing and forecasting using principal oscillation pattern analysis. *Monthly Weather Review*, 117, 2165–2185. [https://doi.org/10.1175/1520-0493\(1989\)117<2165:rfaful>2.0.co;2](https://doi.org/10.1175/1520-0493(1989)117<2165:rfaful>2.0.co;2)
- Penland, C. (1996). A stochastic model of Indo-Pacific sea surface temperature anomalies. *Physica D*, 98, 534–558. [https://doi.org/10.1016/0167-2789\(96\)00124-8](https://doi.org/10.1016/0167-2789(96)00124-8)
- Penland, C., & Sardeshmukh, P. D. (1995). The optimal growth of tropical sea surface temperature anomalies. *Journal of Climate*, 8, 1999–2024. [https://doi.org/10.1175/1520-0442\(1995\)008<1999:togots>2.0.co;2](https://doi.org/10.1175/1520-0442(1995)008<1999:togots>2.0.co;2)
- Portwood, G. D., Nadiga, B. T., Saenz, J. A., & Livescu, D. (2021). Interpreting neural network models of residual scalar flux. *Journal of Fluid Mechanics*, 907. <https://doi.org/10.1017/jfm.2020.861>
- Pyle, R., Jovanovic, N., Subramanian, D., Palem, K. V., & Patel, A. B. (2021). Domain-driven models yield better predictions at lower cost than reservoir computers in Lorenz systems. *Philosophical Transactions of the Royal Society A*, 379, 20200246. <https://doi.org/10.1098/rsta.2020.0246>
- Rahman, S., San, O., Rasheed, A. (2018). A hybrid approach for model order reduction of barotropic quasi-geostrophic turbulence. *Fluids*, 3(4), 86. <https://doi.org/10.3390/fluids3040086>
- Rasp, S., Pritchard, M. S., & Gentile, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684–9689. <https://doi.org/10.1073/pnas.1810286115>
- Rasp, S., & Thuerey, N. (2021). Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: A new model for weatherbench. *Journal of Advances in Modeling Earth Systems*, 13(2), e2020MS002405. <https://doi.org/10.1029/2020ms002405>
- Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., & Carvalhais, N. (2019). Deep learning and process understanding for data-driven earth system science. *Nature*, 566(7743), 195–204. <https://doi.org/10.1038/s41586-019-0912-1>
- Rowley, C. W., & Dawson, S. T. (2017). Model reduction for flow analysis and control. *Annual Review of Fluid Mechanics*, 49, 387–417. <https://doi.org/10.1146/annurev-fluid-010816-060042>
- Ryzhov, E., Kondrashov, D., Agarwal, N., & Berloff, P. (2019). On data-driven augmentation of low-resolution ocean model dynamics. *Ocean Modelling*, 142, 101464. <https://doi.org/10.1016/j.ocemod.2019.101464>
- Ryzhov, E., Kondrashov, D., Agarwal, N., McWilliams, J., & Berloff, P. (2020). On data-driven induction of the low-frequency variability in a coarse-resolution ocean model. *Ocean Modelling*, 153, 101664. <https://doi.org/10.1016/j.ocemod.2020.101664>
- Salman, A. G., Heryadi, Y., Abdurahman, E., & Suparta, W. (2018). Weather forecasting using merged long short-term memory model (LSTM) and autoregressive integrated moving average (ARIMA) model. *Journal of Computer Science*, 14(7), 930–938. <https://doi.org/10.3844/jcssp.2018.930.938>

- Sardeshmukh, P., Penland, C., & Newman, M. (2001). Rossby waves in a stochastically fluctuating medium. In P. Imkeller, J.-S. von Storch (Eds.), *Stochastic climate models* (pp. 369–384). Springer. https://doi.org/10.1007/978-3-0348-8287-3_17
- Sardeshmukh, P., Penland, C., & Newman, M. (2003). Drifts induced by multiplicative red noise with application to climate. *EPL (Europhysics Letters)*, 63(4), 498–504. <https://doi.org/10.1209/epl/i2003-00550-y>
- Scher, S., & Messori, G. (2019). Generalization properties of feed-forward neural networks trained on Lorenz systems. *Nonlinear Processes in Geophysics*, 26(4), 381–399. <https://doi.org/10.5194/npg-26-381-2019>
- Seiffert, R., & Von Storch, J.-S. (2008). Impact of atmospheric small-scale fluctuations on climate sensitivity. *Geophysical Research Letters*, 35(10). <https://doi.org/10.1029/2008gl033483>
- Seiffert, R., & von Storch, J.-S. (2010). A stochastic analysis of the impact of small-scale fluctuations on the tropospheric temperature response to CO₂ doubling. *Journal of Climate*, 23(9), 2307–2319. <https://doi.org/10.1175/2009jcli3043.1>
- Seleznnev, A., Mukhin, D., Gavrilov, A., Loskutov, E., & Feigin, A. (2019). Bayesian framework for simulation of dynamical systems from multidimensional data using recurrent neural network. *Chaos*, 29(12), 123115. <https://doi.org/10.1063/1.5128372>
- Sexton, D. M., Murphy, J. M., Collins, M., & Webb, M. J. (2012). Multivariate probabilistic projections using imperfect climate models part I: Outline of methodology. *Climate Dynamics*, 38(11–12), 2513–2542. <https://doi.org/10.1007/s00382-011-1208-9>
- Strouine, K., Kravtsov, S., Kondrashov, D., & Ghil, M. (2010). Reduced models of atmospheric low-frequency variability: Parameter estimation and comparative performance. *Physica D: Nonlinear Phenomena*, 239(3), 145–166. <https://doi.org/10.1016/j.physd.2009.10.013>
- Sura, P. (2011). A general perspective of extreme events in weather and climate. *Atmospheric Research*, 101(1–2), 1–21. <https://doi.org/10.1016/j.atmosres.2011.01.012>
- Sura, P. (2013). Stochastic models of climate extremes: Theory and observations. In A. AghaKouchak, D. Easterling, K. Hsu, S. Schubert, & S. Sorooshian (Eds.), *Extremes in a changing climate* (pp. 181–222). Springer. https://doi.org/10.1007/978-94-007-4479-0_7
- Sura, P., Newman, M., Penland, C., & Sardeshmukh, P. (2005). Multiplicative noise and non-gaussianity: A paradigm for atmospheric regimes? *Journal of the Atmospheric Sciences*, 62(5), 1391–1409. <https://doi.org/10.1175/jas3408.1>
- Vlachas, P. R., Byeon, W., Wan, Z. Y., Sapsis, T. P., & Koumoutsakos, P. (2018). Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks. *Proceedings of the Royal Society A: Mathematical, Physical & Engineering Sciences*, 474(2213), 20170844. <https://doi.org/10.1098/rspa.2017.0844>
- Wang, Q., Ripamonti, N., & Hesthaven, J. S. (2020). Recurrent neural network closure of parametric POD-Galerkin reduced-order models based on the Mori-Zwanzig formalism. *Journal of Computational Physics*, 410, 109402. <https://doi.org/10.1016/j.jcp.2020.109402>
- Watson, P. A. (2019). Applying machine learning to improve simulations of a chaotic dynamical system using empirical error correction. *Journal of Advances in Modeling Earth Systems*, 11(5), 1402–1417. <https://doi.org/10.1029/2018ms001597>
- Weyn, J. A., Durran, D. R., & Caruana, R. (2020). Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere. *Journal of Advances in Modeling Earth Systems*, 12(9), e2020MS002109. <https://doi.org/10.1029/2020ms002109>
- Williamson, D. (2015). Exploratory ensemble designs for environmental models using k-extended Latin hypercubes. *Environmetrics*, 26(4), 268–283. <https://doi.org/10.1002/env.2335>
- Williams, P. D., Howe, N. J., Gregory, J. M., Smith, R. S., & Joshi, M. M. (2016). Improved climate simulations through a stochastic parameterization of ocean eddies. *Journal of Climate*, 29(24), 8763–8781. <https://doi.org/10.1175/jcli-d-15-0746.1>
- Zanna, L., Mana, P. P., Anstey, J., David, T., & Bolton, T. (2017). Scale-aware deterministic and stochastic parametrizations of eddy-mean flow interaction. *Ocean Modelling*, 111, 66–80. <https://doi.org/10.1016/j.ocemod.2017.01.004>
- Zhang, Q., Wang, H., Dong, J., Zhong, G., & Sun, X. (2017). Prediction of sea surface temperature using long short-term memory. *IEEE Geoscience and Remote Sensing Letters*, 14(10), 1745–1749. <https://doi.org/10.1109/lgrs.2017.2733548>
- Zhang, Y., & Held, I. M. (1999). A linear stochastic model of a GCM's midlatitude storm tracks. *Journal of the Atmospheric Sciences*, 56(19), 3416–3435. [https://doi.org/10.1175/1520-0469\(1999\)056<3416:alsmoa>2.0.co;2](https://doi.org/10.1175/1520-0469(1999)056<3416:alsmoa>2.0.co;2)