

Probing forced responses and causality in data-driven climate emulators: Conceptual limitations and the role of reduced-order models

Fabrizio Falasca ^{*}

Courant Institute of Mathematical Sciences, New York University, New York, New York 10012-1185, USA



(Received 14 July 2025; accepted 13 November 2025; published 19 December 2025; corrected 20 January 2026 and 5 March 2026)

This article is part of a Physical Review Collection on the [Physics of Changing Climate](#).

A central challenge in climate science and applied mathematics is developing data-driven models of multiscale systems that capture both stationary statistics and responses to external perturbations. Current neural climate emulators aim to resolve the atmosphere-ocean system in all its complexity but often struggle to reproduce forced responses, limiting their use in causal studies such as Green's function experiments. To explore the origin of these limitations, we first examine a simplified dynamical system that retains key features of climate variability. We interpret the results through linear response theory, providing a rigorous framework to evaluate neural models beyond stationary statistics and probe causal mechanisms. We argue that the ability of emulators of multiscale systems to reproduce perturbed statistics depends critically on (1) the choice of an appropriate coarse-grained representation and (2) careful parameterizations of unresolved processes. These insights highlight reduced-order models, tailored to specific goals, processes, and scales, as valuable alternatives to general-purpose emulators. We next consider a real-world application, developing a neural model to investigate the joint variability of the surface temperature field and radiative fluxes. The model infers a multiplicative noise process directly from data, largely reproduces the system's probability distribution, and enables causal studies through forced responses. We discuss its limitations and outline directions for future work. These results expose key challenges in data-driven modeling of multiscale physical systems and underscore the value of coarse-grained, stochastic approaches, with response theory as a principled framework to guide model design and enhance causal understanding.

DOI: [10.1103/2f4r-k8lr](https://doi.org/10.1103/2f4r-k8lr)

I. INTRODUCTION

A key challenge in current applied mathematics is to build data-driven models from partial observations of multiscale physical systems. One of the most complex examples is the climate system: a forced, dissipative, and chaotic system whose dynamics spans at least 15 orders of magnitude across spatial and temporal scales [1]. This problem has been extensively studied over the past three decades, mainly in the realm of reduced-order models [2]. Seminal contributions include the work of Penland *et al.* [3–5], the strategy of Kravtsov *et al.* [6–9], and the contributions of Majda *et al.* [10,11]. Recent proposals include Markov chain and generative approaches [12–16], as well as autoencoder-based, latent-space formulations [17]. Most prior work has focused on evaluating reduced-order models on equilibrium statistics and forecasting, with limited attention to responses to *external* perturbations. Notable exceptions include the contributions of

Majda and Qi [18,19]. Capturing such responses is key for building models that not only mimic the system's dynamics but also faithfully capture causal relations across degrees of freedom [20,21].

Recently, autoregressive models based on deep learning architectures have been proposed as potential *emulators* of complex turbulent systems such as the weather [22–26] and climate [27–36]. Unlike traditional reduced-order models, these neural architectures aim to emulate the full hierarchy of spatial and temporal scales present in the system. For example, Refs. [33,34] emulate the full ocean and atmosphere components of state-of-the-art climate models. Such emulators often reproduce stationary statistics and generate forecasts across a range of horizons. They can potentially be used for tuning and sensitivity analysis [37], and to improve the study of rare events [38], offering a promising direction for fast and efficient tools in climate modeling. However, several limitations remain: For example, long rollouts can suffer from error accumulation that can lead to instabilities [39,40]; rare events are not captured if absent from the training set, limiting their applicability for rare-event analysis in a changing climate [41]; key atmospheric modes, such as the quasibiennial oscillation and the southern annular mode, are not well represented [42]. A limitation particularly relevant for climate science, as opposed to weather forecasting, is that high skill in reproducing stationary statistics does not ensure accurate

^{*}Contact author: fabri.falasca@nyu.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](#) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

responses in perturbation experiments, thereby limiting the use of such models for causal studies of climate dynamics. This has two consequences: (1) Neural emulators often fail to generalize in out-of-distribution climate change scenarios and (2) they may not capture the correct causal relations across variables, two issues that are likely related. The work of Van Loon *et al.* [43] illustrates this limitation in the context of the “pattern effect” problem [44,45], i.e., the influence of sea surface temperature (SST) warming patterns on the Earth’s energy imbalance. Using the ACE2 emulator [29] trained on the ERA5 reanalysis [46], they constructed a Green’s function by perturbing SSTs regionally and measuring the radiative response. The resulting Green’s function failed to capture the causal relation between SST and radiative fluxes in a quantitative way. This is just one example of limitations of neural emulators in capturing responses to perturbations with additional examples in Refs. [33,34,47]. Even if such issues are addressed in specific cases, as shown by Wu *et al.* [48] in the context of the pattern effect, they highlight a broader challenge: Building general-purpose emulators that reliably capture forced responses remains inherently difficult and context dependent. Such limitations are also likely to become more pronounced in coupled settings. Additionally, emulators’ skill in capturing forced responses is often assessed using single trajectories. In climate applications, however, the relevant quantities are probability distributions, particularly the responses of the mean and variance. We will show that focusing on the skill in capturing variance responses can reveal important limitations otherwise hidden when examining only the mean response. The limitations examined here are especially relevant for emulators built from observational data and intended to model “real-world” dynamics.

While the field of neural emulation is advancing rapidly and yielding impressive results, it is useful to step back and examine its fundamental limitations in capturing forced responses, both from a theoretical standpoint and in controlled settings. In the first part of this paper, we build on a series of previous works by Vulpiani *et al.* on the limitations of data-driven modeling [49–51]. We consider a simplified stochastic dynamical system and ask the following question: Can neural autoregressive models, trained on (practically) infinitely long time series, accurately capture responses to general external perturbations? In addressing this question, we focus on two relevant scenarios in data-driven-only modeling:

(1) Unknown equations, fully observed state vector: Data-driven modeling when the entire state vector is available. Here, we assume no knowledge of the underlying equations, but we have access to very long time series of the complete system state vector.

(2) Unknown equations, partially observed state vector: Data-driven modeling in the presence of partial observations. This reflects common constraints in modeling high-dimensional systems from data, where the full set of relevant variables is typically unobserved [52], for example because of unresolved fast and small-scale processes, or even unknown [51]. In this more realistic setting, we again assume no knowledge of the underlying equations and only partial access to the system’s state vector.

To tackle such questions, we proceed pedagogically by focusing on a multiscale “triad model” first proposed by

Majda *et al.* in Refs. [53,54]. The model mimics the abstract structure of dynamical cores of large-scale climate models with a linear operator and energy conserving quadratic nonlinearities. It provides a simplified yet relevant mathematical framework to unambiguously explore fundamental (rather than technical) limitations in capturing forced responses in climate emulators. We address the questions above by training neural stochastic emulators on a long integration of the triad model and test the models’ performance on stationary and perturbed cases. In doing so, we face the following issue: quantifying how an emulator respond to general forcing is an ill-posed question, as there are infinite possible forcings. Nonetheless, the linear response can be computed by convolving the forcing with an impulse response operator [55,56]. This operator can be inferred from the model and is a *building block* that an emulator must capture to reproduce the response to general external perturbations. Therefore, linear response theory is here proposed as a rigorous framework to evaluate neural models beyond stationary statistics and probe causal mechanisms [20]. In the case of a fully observed state vector, we show that the inferred neural model can reproduce stationary and perturbed statistics of the system. In the case of partial observations, as in realistic settings, the performance of a data-driven model necessarily depends on (1) the choice of proper variables to model and (2) suitable parametrizations of unobserved processes. We then note that the lack of clear scale separation in realistic applications further complicates the problem, leading to non-unique coarse-graining choices tailored to specific goals. These issues are illustrated through a real-world example in a later section. The field of reduced-order modeling, introduced earlier, has long sought to confront these challenges by formulating effective Langevin equations capturing the coarse-grained, slow dynamics of the system [5,57,58], though typically without addressing responses to external perturbations. Therefore, we will argue that, a promising perspective is to build on such established approaches, emphasizing coarse-grained dynamics and stochastic parametrizations, while using neural networks as powerful tools within this framework.

We emphasize that the focus on a simplified dynamical system is not a limitation in this context. In fact, the goal is not to demonstrate the success of a method, but to illustrate, by analogy, fundamental issues of data-driven modeling that are already evident in low dimensions and are likely to become more pronounced in higher-dimensional settings [59]. The first part of this paper is pedagogical in nature, aiming to bridge perspectives across researchers in statistical physics, climate science, and machine learning, while also providing practical methodological tools for the design and evaluation of data-driven models.

Building on these ideas, we conclude this study by considering a real-world example. Throughout this paper, we emphasize, when possible, the distinction between “real-world examples/systems” and merely high-dimensional dynamics. By real-world examples, we refer to situations where (1) we observe only the temporal evolution of a subset of variables without access to the full state vector (i.e., a projection of the full dynamics) or (2) when even in the presence of all variables, the system encompasses a huge number of spatial and temporal scales. Both cases (1) and (2) lead to the need

to formulate *effective* rather than *primitive* equations for both practical and conceptual reasons [57]. For an early instance of case (2), see the work of Charney [60], Charney *et al.* [61], and the discussion in Dalmedico [58]. The example we consider is the “pattern effect” introduced above, which requires Green’s function experiments to quantify the response of top-of-atmosphere (TOA) radiative fluxes to perturbations in the SST field. To this end, we develop a simplified stochastic neural model that emphasizes a large-scale representation of SST and global-mean net radiative flux, rather than attempting to resolve all spatial and temporal scales of variability as in current emulators. We discuss the coarse-graining choices relevant for model construction, as well as a practical and efficient way to fit multiplicative noise from data. The model largely reproduces stationary statistics, especially the most energetic SST mode, corresponding to the El Niño Southern Oscillation (ENSO) [62], and it enables Green’s functionlike experiments to infer direct causal relations between SST and TOA radiative flux. The discussion also addresses the model’s limitations and potential methodological developments.

The paper is organized as follows: In Sec. II, we introduce the theoretical ideas underpinning this work. In Sec. III, we discuss key conceptual limitations in data-driven modeling of real-world dynamical systems. In Sec. IV, we present a real-world application. A summary of the main findings and proposals follow in Secs. V and VI.

II. THEORETICAL FRAMEWORK: DATA-DRIVEN STOCHASTIC CLIMATE MODELING AND RESPONSE THEORY

We begin by outlining the main theoretical concepts that underpin this work. First, we briefly introduce general ideas for reduced-order stochastic modeling relevant for climate dynamics. We then review key aspects of response theory, which offers a rigorous framework to move beyond the evaluation of stationary dynamics, probe causal mechanisms, and ultimately guide model design.

A. Reduced-order stochastic climate models: theoretical and practical considerations

A general dynamical system can be represented in the form

$$d\mathbf{u}/dt = \mathbf{F} + \mathbf{L}\mathbf{u} + \mathbf{n}(\mathbf{u}), \quad (1)$$

where $\mathbf{F}$ is an external forcing, and $\mathbf{L}$ and $\mathbf{n}(\mathbf{u})$ represent linear and nonlinear operators, respectively. $\mathbf{u}$ is not necessarily finite dimensional, and in the context of coupled atmosphere-ocean variability, it could in principle represent the deviation of all fields, at all scales, from their climatology [4]. However, for obvious practical reasons, Eq. (1) needs to be simulated on finite spatial and temporal resolutions. In contrast to numerical weather prediction, which demands high spatial and temporal resolution, low-frequency climate variability is governed by a few dominant large-scale modes [4,7,21,54,55,63]. The goal of reduced-order stochastic climate models is to capture the evolution of these modes and their interactions. This modeling choice allows us to move away from a grid-based representation of the system, emphasizing resolution-independent modes as fundamental

components of the system. The cumulative effect of unresolved fast variables, present at smaller spatial scales, is parametrized as stochastic terms. More formally, we decompose $\mathbf{u}$ into slow and fast components, $\mathbf{u} = (\mathbf{x}, \mathbf{y})$, leading to an *effective* Langevin equation of the form

$$d\mathbf{x} = \mathbf{f}_{\text{eff}}(\mathbf{x})dt + \mathbf{\Sigma}(\mathbf{x})d\mathbf{W}, \quad (2)$$

where $\mathbf{f}_{\text{eff}}$ represents the deterministic “drift” dynamics, $\mathbf{x}$ are the slow variables, and $\mathbf{W}$ is the standard Wiener process. The noise covariance matrix $\mathbf{\Sigma}$ can be independent of $\mathbf{x}$, corresponding to additive noise, or state-dependent $\mathbf{\Sigma}(\mathbf{x})$, corresponding to multiplicative noise. In the latter case, the equation is here interpreted in the Itô sense. Equation (2) is mathematically justified in the limit of infinite timescale separation between fast and slow modes [1]. In the most general (and often intractable) setting, the effective dynamics of the slow variables also involve memory effects induced by the fast degrees of freedom, leading to non-Markovian dynamics. Kondrashov *et al.* [7] addressed this problem in a practical and efficient way by introducing the framework of multilevel stochastic models. Nonetheless, appropriate coarse-graining procedures have been shown to effectively average out non-Markovian contributions in many climate applications, yielding useful Markovian approximations [3,4]. We refer the reader to Ref. [2] for a review on the role of non-Markovian contributions in effective climate equations and to Refs. [64,65] for specific examples.

Reduced-order stochastic climate models, as in Eq. (2), provide an efficient alternative to state-of-the-art climate models for studying low-frequency climate variability. They are much more computationally efficient and allow us for mechanistic understanding of climate dynamics by isolating core processes and examining their qualitative behavior. This helps bridging the gap between simulations and understanding [66]. This has fueled the development of low-dimensional stochastic models learned either from data [2–4,7,10,18,67,68] or from known equations [53,69–72], as well as inference-based techniques to infer interactions across modes of variability [21,73–75].

In practice, given a phenomenon of interest, we typically begin with spatiotemporal data consisting of N time series of length T . We then often proceed as follows.

1. Coarse-grained dynamics

In spatiotemporal dynamical systems like climate, coarse graining involves: (1) choosing a limited set of climate fields (e.g., temperature, velocities, etc.), (2) retaining a few relevant modes or large-scale averages, and (3) restricting the range of temporal scales. Throughout the paper, we will often refer to the so-called *proper* variables as the coarse-grained representation obtained via (1)–(3). The relevant modes are often identified through empirical orthogonal functions (EOFs) [8,76], by retaining the first $n < N$ of such bases. The correspondent principal components serve as the reduced time series for model training. Although EOFs are convenient, they may not be always optimal, and alternatives such as principal interaction patterns [77], community detection [21], and autoencoders [17,78] may be more suitable depending on the application. The performance of the final model is sensitive

to all coarse-graining choices [(1)–(3)]; see, e.g., [79,80]. In climate dynamics, there is no clear space- and time-scale separation. In this case, identifying the appropriate coarse-grained representation depends on the specific goals of the researcher and remains one of the most difficult challenges in data-driven modeling [20,49].

2. Inferring a stochastic model

Given the coarse-grained data $\mathbf{x}$, the goal is to infer a stochastic Markovian model as in Eq. (2). We summarize several traditional approaches that motivated our work, while acknowledging that many other methods exist beyond those discussed here. The simplest example is the Linear Inverse Model (LIM) proposed by Penland [3], where the drift is linear and the noise is additive, both estimated directly from data. LIMs have shown forecasting skill for sea surface temperature anomalies comparable to climate models [4,5,81,82]. Recent developments extended the LIM framework to include colored noise [83] and nonlinearities [84]. Kravtsov *et al.* and Kondrashov *et al.* [6,7] generalized this framework through multilevel regression, incorporating quadratic drift terms and non-Markovian effects. These models perform well for stationary statistics [8] and forecasting [85]. The work of Majda and Harlim extended the multilevel regression framework by proposing stability constraints [11]. A theoretically grounded alternative is the approach of Majda, Timofeyev, and Vanden-Eijnden (MTV) [69–72], dubbed MTV by the authors in [72], which derives effective equations for the coarse-grained variables $\mathbf{x}$ rather than assuming them. MTV has been successful for simplified models, though less so for complex climate systems, likely due to choices of coarse-graining steps (1)–(3) shown above rather than intrinsic limitations; see also discussion in Ref. [8]. We will expand on this approach in Sec. III and Appendix B.

Overall, these works demonstrated that robust reduced-order stochastic climate models can be built through relatively simple regression strategies. In this paper, we build on these insights and investigate how simple neural architectures can infer both the drift and state-dependent noise in Eq. (2) with minimal assumptions. Our neural models retain a decomposition into linear and nonlinear terms, $\mathbf{L}$ and $\mathbf{n}(\mathbf{x})$, as in the general representation of dynamical systems in Eq. (1).

B. Response theory: A (very) brief overview

In this paper, we argue that data-driven models should be tested on their ability to capture both stationary statistics and responses to external perturbations. Assessing responses to arbitrary forcings is ill-posed, as the space of possible perturbations is essentially infinite. A way forward is to consider small impulse perturbations, which allow us to define a response operator and derive general conclusions about causality and forced responses within the linear response regime.

Impulse response operator

Given a n -dimensional coarse-grained state vector $\mathbf{x}(t) = [x_1, x_2, \dots, x_n](t)$, we perturb a degree of freedom x_j with

a small, impulse perturbation $\delta x_j(0)$, $x_j(0) \rightarrow x_j(0) + \delta x_j(0)$ [20,86]. We then consider a general observable $A(x_k(t))$ (i.e., a general functional of the system) and quantify its response, in terms of ensemble average, as

$$\delta \langle A(x_k(t)) \rangle = \langle A(x_k(t)) \rangle_p - \langle A(x_k(t)) \rangle, \quad (3)$$

where $\langle \cdot \rangle$ refers to ensemble average and the subscript “p” refers to the perturbed system. In this paper, we are going to focus on two observables: (1) $A(x_k(t)) = x_k(t)$ and (2) $A(x_k(t)) = (x_k(t) - \mu_k(t))^2$, with $\mu_k(t)$ representing the time-dependent mean of the distribution. The two observables allow us to capture time-dependent responses in the mean and variance of the perturbed distribution, respectively. The impulse response operator $\mathbf{R}(t) \in \mathbb{R}^{n,n}$ is then defined as

$$R_{k,j}(t) = \lim_{\delta x_j(0) \rightarrow 0} \frac{\delta \langle A(x_k(t)) \rangle}{\delta x_j(0)}. \quad (4)$$

We refer to Appendix A for details on how to infer $R_{k,j}(t)$ in practice. The utility of the inferred impulse response operator is two-fold:

(1) Causal inference: Understanding causal relations across degrees of freedoms in physical systems requires to study the effect of interventions on the system [87]. The inferred response operator quantifies time-dependent cause-effect interactions across degrees of freedom x_j and x_k [20,21,88–90]. We refer to Appendix A for details on the relation between the response operator and causality.

(2) Response to generic time-dependent forcing: The response operator $\mathbf{R}(t)$ can be used to predict the response of the system to small, but general, time-dependent external forcing $\mathbf{f}(t) \in \mathbb{R}^n$ as

$$\delta \langle A(x_k(t)) \rangle = \sum_j \int_0^t R_{k,j}(\tau) f_j(t - \tau) d\tau. \quad (5)$$

Importantly, while we assume linearity in the system’s response, the formulas above are valid for both linear and nonlinear systems.

A data-driven model capable of accurately capturing the impulse response operator will also capture the time-dependent causal interactions across degrees of freedom. Response theory can then be used to assess the skills of emulators by going beyond purely statistical metrics such as r -squared, assessing the causal sensitivity of the system to external perturbation [21]. Thus, this approach opens rigorous avenues for evaluating and studying data-driven models. Moreover, recovering the impulse response operator ensures a faithful approximation of the system’s linear response to small general forcings. For this reason, we view the impulse response operator as a fundamental *building block* that a data-driven model must reproduce to capture the responses to general perturbations. Importantly, the response operator can be inferred from data alone through the fluctuation-dissipation theorem (FDT) [56], as discussed in Sec. V and in Appendix A.

III. LIMITATIONS OF DATA-DRIVEN MODELING BY ANALOGY: THE CASE OF PARTIAL OBSERVATIONS

A fundamental limitation of purely data-driven modeling of real-world, multiscale climate systems is that many degrees of freedom remain unobserved. Even when the governing equations are known, as in meteorology, climate fields are only sampled at finite temporal and spatial resolutions leading to unresolved fast and small-scale processes. The challenge is even greater in settings where governing equations are unknown and provide no guidance for emulator design. In both cases, data-driven approaches can reproduce the variability of the retained variables, including their probability distributions, autocorrelation functions, and short-term forecasts. However, the response of probability distributions to perturbations, particularly the mean and variance in realistic climate applications, can uncover structural deficiencies stemming from unresolved dynamics that would otherwise remain hidden from stationary statistics alone. The goal of this section is to illustrate such limitations within a controlled mathematical setting, where they can be made explicit and unambiguous.

A. A triad model as a simplified prototype for coarse-grained, stochastic climate modeling

We focus on a simplified yet relevant dynamical system that retains essential features of high-dimensional geophysical flows, making the conclusions of this section meaningful for climate applications. The model was introduced by Majda *et al.* [53] and later used in Ref. [54] in the context of climate dynamics. In what follows, we present the model and refer to Appendix B for an in-depth discussion on its relevance as a building block for complex turbulent flows. The triad model considered in this study takes the following form:

$$\begin{aligned} dx_1 &= (L_{11}x_1 + L_{12}x_2 + L_{13}x_3 + Ix_1x_2)dt, \\ dx_2 &= (-L_{12}x_1 - Ix_1^2 - \gamma_2\epsilon^{-1}x_2)dt + \sigma_2\epsilon^{-1/2}dW_2, \\ dx_3 &= (-L_{13}x_1 - \gamma_3\epsilon^{-1}x_3)dt + \sigma_3\epsilon^{-1/2}dW_3. \end{aligned} \quad (6)$$

This idealized system mimics the functional structure of turbulent geophysical flows, featuring linear couplings L_{11} , L_{12} , L_{13} and energy-conserving quadratic nonlinearities. The parameter I controls the strength of the nonlinearity: When $I = 0$, the model reduces to purely linear dynamics and Gaussian statistics. The quadratic interactions typically arise when geophysical turbulent flows are expressed in terms of orthogonal bases such as EOFs [18]. The parameter ϵ sets the timescale separation between the slow mode x_1 and the fast modes (x_2, x_3) , allowing exploration of multiscale dynamics in a controlled setting. The terms W_2 and W_3 represent independent Wiener processes. The model generates smooth, unimodal distributions with relatively small (but significant) deviations from Gaussianity, and it exhibits fast decay of correlations. Both features are relevant in the study of low-frequency climate variability [91] and found in both general circulation models and observational data (see, e.g., Refs. [92–94]).

B. The case of unknown equations but a fully observed state vector

We consider the relatively simple case of a low-dimensional dynamical system with only one characteristic timescale (i.e., $\epsilon = 1$). The parameters chosen for this first test are $L_{11} = -2$, $L_{12} = 1$, $L_{13} = 1$, $\gamma_2 = 0.6$, $\gamma_3 = 0.4$, $\sigma_2 = 1.2$, $\sigma_3 = 0.8$, $I = 1$, and $\epsilon = 1$. We now integrate the system in Eq. (6) with an Euler-Maruyama scheme for a period of length $T = 10^7$ with time step $dt = 0.01$ (i.e., 10^5 model time units). Given the system's decorrelation timescale (shown later), the resulting time series can be effectively treated as infinitely long. Given this long trajectory of the full state vector $\mathbf{x}(t) = [x_1, x_2, x_3](t)$, we learn data-driven models directly from the time series. We construct both linear and nonlinear stochastic models. We specify the functional forms of these models and assess their ability to reproduce both the stationary and perturbed statistics of the original triad system.

1. Functional form of the data-driven models considered

We consider two data-driven models. To keep the main text focused, we leave the technical details of the fitting procedures, along with details of the neural networks, in Appendix C.

(1) Linear model: We consider the following linear model:

$$d\mathbf{x} = \mathbf{L}xdt + \Sigma d\mathbf{W}, \quad (7)$$

where $\mathbf{L}$ is a linear operator, $\Sigma\Sigma^T$ is the noise covariance matrix, and $\mathbf{W}$ is the standard Wiener process.

(2) Nonlinear model: The nonlinear model takes the following form:

$$d\mathbf{x} = (\mathbf{L}x + \mathbf{n}(\mathbf{x}))dt + \Sigma d\mathbf{W}. \quad (8)$$

$\mathbf{L}$ is a linear operator and $\mathbf{n}(\mathbf{x})$ is parametrized by a multilayer perceptron (MLP). Crucially, the term $(\mathbf{L} + \mathbf{n}(\mathbf{x}))$ is fitted jointly. In what follows, we are going to often refer to this model as “neural model” or “neural emulator.”

The two models above are trained on a long simulation of the triad system in Eq. (6).

Note that this is the only section in this work where we compare linear and nonlinear autoregressive models. The linear method is considered for two reasons. First, we show that linear inverse modeling, as originally proposed in Ref. [3], serves as a strong baseline for neural autoregressive models, successfully capturing both the mean and variance of the system's stationary distribution, the autocorrelation functions, and the ensemble mean response to external perturbations. Second, and consistent with theoretical expectations [54], we demonstrate that linear models are fundamentally limited in capturing changes in the full probability distribution: For instance, they cannot represent any response in the variance to external forcings. This limitation restricts their usefulness for studying forced changes in distributions beyond the mean.

2. Stationary statistics

We integrate the two data-driven models in Eqs. (7) and (8) for $T = 10^7$, yielding time series of the same length as those obtained from the original triad system. We then compare the modeled and original probability density functions (PDFs)




FIG. 1. First row: stationary distributions of the variables $x_1(t)$, $x_2(t)$, and $x_3(t)$ as modeled by the original triad system in Eq. (6), the linear emulator in Eq. (7), and the neural emulator in Eq. (8). Stationary distributions are obtained after running the three systems for 10^5 model time units. Second row: same as the first row but for the autocorrelation functions.

and autocorrelation functions (ACFs), as shown in Fig. 1. The neural model described in Eq. (8) successfully captures the stationary distribution of the triad system. In particular, it outperforms the linear model by accurately reproducing the non-Gaussian features of x_1 . Both models also reproduce the autocorrelation functions well, as seen in the second row of Fig. 1. In summary, the neural emulator is capable of capturing both the PDF and autocorrelations of the full state vector $\mathbf{x}(t)$. At the same time, the linear model serves as a strong baseline for evaluating more sophisticated data-driven approaches.

3. Response of probability distributions to external perturbations

a. Impulse response functions. We now examine the impulse response of the original triad model [Eq. (6)], the linear model [Eq. (7)], and the nonlinear neural model [Eq. (8)]. Specifically, we analyze the time-dependent response of observables $A(x_k(t))$, with $k = 1, 2, 3$, to an impulse perturbation imposed on $x_1(0)$ at time $t = 0$. We choose as observables the mean and variance of the distribution. The results, shown in Fig. 2, reveal that the nonlinear neural model captures the mean response of the distribution quantitatively, while the linear model captures it only qualitatively. The most significant differences arise in the response of the variance: The linear model yields exactly zero variance response, as expected [54]. In contrast, the neural model captures the variance response with relatively high fidelity. We therefore consider the neural model presented here as a skillful model both in terms of reproducing the stationary distribution as well as the response operator.

b. Responses to a specific logarithmic forcing. We now consider a more targeted test and introduce a logarithmic forcing term to the right-hand side of both the original and data-driven triad model. The forcing is applied only to x_1 and takes the form $(0.01\log(1+t), 0, 0)$ per time step. Each model is integrated for 2500 time steps (25 model time units), and we compute the difference between the forced and unforced simulations in terms of ensemble mean and variance responses. Results are presented in Fig. 3. The neural model [Eq. (8)] accurately reproduces the forced responses across

all three components. The linear model [Eq. (7)] performs qualitatively well but exhibits significant deviations in the ensemble mean response of the second mode x_2 as shown in Fig. 3(b). This is consistent with the earlier section: The largest error in the linear model's impulse response also occurred in x_2 [see Fig. 2(b)]. This supports the interpretation that accurately capturing the impulse response operator is critical for a model's reliability in predicting responses to general, time-dependent forcings. As expected, the linear model is not able to capture the response in ensemble variance. The neural model shows skillful response, capturing the time-dependent changes in ensemble variance relatively well, with the largest error in x_2 . In general, considering that the changes in variance are very small, we conclude that the neural emulator is skillful in responding to the imposed logarithmic forcing.

4. Main takeaways

Given a (practically) infinitely long trajectory of a stochastic dynamical system, and assuming the ideal condition where the full state vector is observed, it is possible to construct a neural stochastic model that is stable, reproduces the stationary statistics, and responds accurately to small but general external perturbations. The analysis does not rule out that other systems may struggle to capture responses even with full observations. These issues, while important, are considered here as technical, rather than fundamental, and point to the need for more sophisticated architectures, improved algorithms, or the inclusion of additional physical constraints within existing frameworks.

While the results shown in this Section are encouraging, they remain far from real-world scenarios. Emulators of climate models would require full access to all climate fields at all temporal and spatial scales, which is often impractical due to the enormous data demands. However, if that requirement were met, then the main challenge would reduce to the need for very long time series for model training [49]. More fundamentally, emulators built from observational data and aimed at capturing “real-world” dynamics typically have access only to partial observations of a complex, interconnected system.




FIG. 2. Impulse response functions as modeled by the original triad system in Eq. (6), the linear emulator in Eq. (7), and the neural emulator in Eq. (8). First row: response of the ensemble mean to *impulse* perturbations imposed on $x_1(0)$ at time $t = 0$. Second row: response of the ensemble variance to *impulse* perturbations imposed on $x_1(0)$ at time $t = 0$. Impulse response functions are computed using an ensemble size of $N_e = 10^5$ members.

C. The case of unknown equations and a partially observed state vector

We now return to the triad model defined in Eq. (6) and consider the more realistic scenario of model inference under partial observations and multiple timescales. The parameters chosen for this test are $L_{11} = -2$, $L_{12} = 0.2$, $L_{13} = 0.1$, $\gamma_2 = 0.6$, $\gamma_3 = 0.4$, $\sigma_2 = 1.2$, $\sigma_3 = 0.8$, $I = 1$, and $\epsilon = 0.01$. In this

setting, the variable x_1 clearly is the slow mode, forcing the much faster modes x_2 and x_3 . We again integrate the system in Eq. (6) with an Euler-Maruyama scheme for a period of length $T = 10^7$ with time step $dt = 0.01$ (i.e., 10^5 model time units).

Unlike the case considered in the previous section, where the entire state vector $\mathbf{x}(t) = [x_1, x_2, x_3](t)$ was available, we now consider a more challenging scenario: Only one




FIG. 3. Response to a small logarithmic forcing $F = (0.01\log(1+t), 0, 0)$ imposed on the right-hand side in Eq. (6), the linear emulator in Eq. (7), and the neural emulator in Eq. (8). First row: response of the ensemble mean to external forcing. Second row: response of the ensemble variance to external forcing. Responses are computed using an ensemble of $N_e = 10^5$ members.

component of the system is observed (or available or chosen), either x_1 , x_2 , or x_3 . The modeling task is thus reduced to constructing a scalar stochastic model for the observed component alone. We require the scalar stochastic model to satisfy two criteria:

- (1) It must reproduce the stationary statistics of the observed component.
- (2) It must correctly capture the mean and variance response to perturbations as in the original triad system.

Both criteria are nontrivial, with the second posing the greater difficulty. Together they mirror core challenges in data-driven modeling built from partial observations. Satisfying these criteria leads to two central challenges:

(1) Choice of variable to model: The reduced-order model depends critically on the choice of the modeled variable. Not all components can be selected to satisfy criteria (1) and (2) discussed above. The only viable option is to model the slow mode x_1 , since the cumulative effect of the much faster x_2 and x_3 can be represented stochastically. Identifying such a *proper* variable is a necessary prerequisite for constructing a meaningful model under partial observations [49,55,95,96]. In realistic spatiotemporal climate dynamics, such choice translates into selecting appropriate climate fields together with a narrow range of spatiotemporal scales. We will give a practical example in Sec. IV.

(2) The need for suitable parametrizations: Having identified the proper variable x_1 to model, the next challenge is to appropriately parametrize the cumulative effects of the unresolved components x_2 or x_3 onto the slow mode x_1 . This leads to the classic problem of stochastic parametrization [1,2,76,97–99]. As we will show, the success of the resulting model, particularly in capturing forced responses in the variance, hinges critically on the quality of this parametrization.

Two brief remarks. (1) In partially observed settings, one might be tempted to appeal to Takens' embedding theorem [153]. This is often not a practical option in real-world systems: the theorem does not extend to stochastic systems, and it becomes infeasible in high-dimensional settings [20,100,135]. (2) For the specific case of the triad model in Eq. (6), Majda and collaborators showed that an analytical reduced-order equation for the slow variable x_1 can be rigorously derived in the limit $\epsilon \rightarrow 0$ [53,54]; see discussion in Appendix B. Our aim here is not to test such analytical strategies, which require full knowledge of the governing equations and are often limited in real-world settings [8]. Rather, we use the triad model as a simple conceptual tool to highlight key challenges of data-driven modeling of multiscale systems. In realistic settings, both the governing equations and the full set of relevant variables are typically unknown [100]: even before specifying a model structure, identifying the slow variables is nontrivial, non-unique and highly problem-dependent.

In what follows, we assume that the proper slow variable, x_1 , has been correctly identified. Given this, we examine how the choice of stochastic parametrization and temporal coarse graining influences the ability of neural models to emulate both stationary statistics and responses to perturbations. This discussion serves as a precursor to the more challenging real-world application in Sec. IV, where identifying the relevant variables requires a series of nontrivial coarse-graining decisions.

1. Functional form of the reduced-order models considered

Given the correct choice of variable x_1 to model, we now proceed in proposing two data-driven scalar models. Technical details of the fitting procedures are left to Appendix C. We consider two neural models:

(1) Neural model with additive noise: We consider the following neural model fitted using an artificial neural network, i.e., a multilayer perceptron. The model takes the following form:

$$dx = (Lx + n(x))dt + \sigma dW, \quad (9)$$

where L is a linear operator (here just a scalar), $n(x)$ is a nonlinear operator, and σ^2 is the noise variance. W is the standard Wiener process.

(2) Neural model with multiplicative noise: We consider a neural model with multiplicative

$$dx = (Lx + n(x))dt + \sigma(x)dW. \quad (10)$$

We use the same drift $(Lx + n(x))$ fitted for the additive model in Eq. (9), so that differences can only arise from the stochastic parametrization. The multiplicative (state-dependent) noise $\sigma(x)$ is fitted with an additional neural network; see Appendix C.

To train the two models above, we consider a long simulation of the triad system in Eq. (6) using the set of parameters reported in Sec. IIIC. We then extract only the time series of x_1 and train the models above on this long, scalar time series $x_1(t)$.

2. Stationary and perturbed statistics of the scalar model

We integrate the two data-driven scalar models in Eqs. (9) and (10) for $T = 10^7$ time steps, yielding time series of the same length as the time series of $x_1(t)$ used for training. To evaluate the models' ability to capture both stationary and perturbed statistics, we examine the stationary probability distribution and the impulse response operator in ensemble mean and variance.

The neural model with additive noise [Fig. 4(a)] fails to accurately capture the stationary statistics in a quantitative way. Moreover, the ensemble mean response to an impulse perturbation decays too slowly, and notably, the model fails to reproduce significant responses in the ensemble variance. In contrast, incorporating multiplicative noise [Eq. (10)] leads to a substantially improved representation of the stationary distribution [Fig. 4(d)] and the model now captures a qualitatively correct signal in the ensemble variance, albeit with a similarly slow decay.

These results indicate that the inclusion of multiplicative noise improves the reduced-order model's ability to represent both stationary statistics and variance responses to perturbations. Moreover, as discussed in Appendix B, the use of multiplicative noise in reduced-order modeling is not an *ad hoc* addition but is theoretically justified and explicitly predicted by rigorous mode reduction strategies [69]. We emphasize two key points:

(1) The analysis in Fig. 4 highlights the importance of evaluating a model's ability in predicting responses in higher-order moments, such as the variance. In the simple example discussed here, focusing solely on the ensemble mean




FIG. 4. First row: stationary probability distribution [panel (a)]; ensemble mean response to an impulse perturbation [panel (b)]; and ensemble variance response to an impulse perturbation [panel (c)] for the variable x_1 in the triad system in Eq. (6) and the variable x in the neural scalar emulator with additive noise in Eq. (9). Second row: same as the first row but we are now considering the neural scalar emulator with multiplicative noise in Eq. (10). Impulse response functions are computed using an ensemble of $N_e = 10^6$ members.

response, which is captured by the deterministic drift, would lead to the incorrect conclusion that the two models respond to perturbations in the same way [see panels (b) and (e) in Fig. 4].

(2) Matching the stationary distribution quantitatively does not guarantee equally accurate responses to perturbations, as shown in Figs. 4(d)–4(f).

The presence of slowly decaying responses in the ensemble mean and variance to perturbations is indicative of a system that is no longer Markovian. The reduced-order stochastic scalar model is trained solely on the time series of $x_1(t)$. However, the time series $x_1(t)$ is generated by the full triad system in Eq. (6) and implicitly reflects the influence of the unobserved variables x_2 and x_3 on it. The variability of the unobserved variables x_2 and x_3 manifests as memory effects in the reduced-order scalar model for x_1 . One strategy for addressing such memory effects is the multilevel regression approach first introduced by Kratsov *et al.* [6] and later generalized by Kondrashov *et al.* [7]. The multilevel regression strategy augments the reduced-order model with a hierarchy of linear equations solving for the residuals from the deterministic dynamics and it can be thought as a practical implementation of the Mori-Zwanzig formalism in statistical mechanics [2,7]. An alternative, pragmatic approach for handling non-Markovian effects is to coarse-grain the data in the temporal dimension, for instance, by averaging over a time window. This suppresses fast fluctuations within that window leading to an effective “Markovianization” of the dynamics. Due to its simplicity and practicality, this method is commonly used in the literature (see, e.g., Ref. [4]) and is also adopted in the present study.

We consider again the time series $x_1(t)$ from the long simulation of the triad system in Eq. (6). As an additional step, we average $x_1(t)$ over a short temporal window of ten time steps (equivalent to 0.1 model time units). This averaging smooths out fast variability and effectively defines a coarse-grained system that is “more Markovian.” We then retrain the neural models with additive and multiplicative noise on the coarse-

grained data. Results are shown in Fig. 5. Now, the model with multiplicative noise not only reproduces the stationary distribution but also closely matches the variance response, capturing its peak magnitude and deviating only slightly in its temporal decay. One could, in principle, tune the averaging window to match the autocorrelation of the emulated $x(t)$ with that of the original $x_1(t)$, thereby ensuring closer agreement in the decay of responses. Furthermore, training on multiple time steps should also result in a more accurate model. We view these as methodological refinements for future work, and leave it outside the scope of this study.

3. Main takeaways

Even with access to very long time series, modeling real-world dynamical systems from data alone remains fundamentally constrained by observing only a subset of the full state vector. The most severe limitation arises from the choice of variables: (1) the model’s skill depends critically on which variables are observed and included. Once these variables are fixed, (2) the influence of unobserved degrees of freedom must be adequately parametrized to capture responses in probability distributions. Although these issues can be illustrated with low-dimensional examples, they become much more severe in realistic spatiotemporal climate systems, where there is little or no scale separation and variable selection involves nontrivial, non-unique decisions about which fields to include and which spatial and temporal scales to retain—that is, how to construct an appropriate coarse-grained representation of the system.

These considerations highlight intrinsic limitations of general-purpose emulators of coupled climate systems in capturing responses in mean and variance of the distribution to *external* perturbations. In contrast, reduced-order neural models are often tailored to specific problems, allowing one to leverage prior literature, insights, and results to guide variables and coarse-graining choices. Such prior knowledge can be essential for constructing data-driven models and for clearly delineating their domains of validity.




FIG. 5. Same as in Fig. 4; however, here the scalar emulators in Eqs. (9) and (10) have been trained on a long trajectory of $x_1(t)$ after coarse graining by averaging every ten time steps. No additional preprocessing was performed for the original triad system in Eq. (6); i.e., the blue curves in this figure and in Fig. 4 are the same.

IV. A REAL-WORLD EXAMPLE: RADIATIVE RESPONSES TO TEMPERATURE PERTURBATIONS IN A CLIMATE MODEL

Building on the insights discussed above, we present a stochastic reduced-order model tailored to a specific real-world, high-dimensional test case. We introduce the scientific problem at hand and the proposed emulator to tackle this specific question. We then highlight the coarse-graining choices defining the proper variables to model. Finally, we showcase the skill of this emulator in studying causal mechanisms in large-scale climate dynamics, as well as its limitations and future work.

A. Scientific problem at hand

1. The “pattern effect” problem

Radiative forcing (e.g., from changes in CO_2) drives surface warming, which alters sea surface temperatures and associated radiative feedbacks. Recent studies show that even for the same global-mean warming, different SST warming patterns can yield different feedbacks. For example, eastern Pacific warming tends to amplify warming through positive feedbacks, while western Pacific warming has a stabilizing effect [44,45]. This sensitivity of Earth’s energy imbalance to spatial SST patterns is known as the “pattern effect.” The pattern effect describes the fast response of TOA radiative fluxes to spatial temperature variations and can be formalized in energy balance models via a Taylor expansion around the steady-state temperature [101].

A common approach to study the pattern effect is to relate SST patterns to TOA radiative fluxes using Green’s function experiments [102], where the radiative response to localized SST perturbations is quantified with atmosphere-only models [45,103–106]. These experiments are valuable for testing model sensitivities but are computationally costly and typically require large SST perturbations (1–4 K) to obtain a measurable response within short simulations. Such large perturbations can trigger nonlinear behavior, undermining the linearity assumptions of the Green’s function

framework [107]. This motivates the need for fast, interpretable data-driven models that can simulate responses of large ensembles at low cost and with small perturbation amplitudes. Furthermore, since climate models are known to be biased, data-driven emulators offer, in principle, a complementary way to learn climate dynamics directly from observations.

2. Possible emulator strategies

Two broad strategies exist for building emulators. One is to construct a general-purpose emulator of the climate system and then apply it to specific problems. However, as shown by Van Loon *et al.* [43] in the context of the ACE2 emulator and the pattern effect, this strategy may not always succeed. A second approach for large-scale climate dynamics is to design an emulator tailored to a given problem; see, e.g., Ref. [5]. While this sacrifices generality, it can also offer key advantages by clarifying the choice of variables and scales to model, grounding the design on prior literature, and isolating only the processes most relevant to the problem.

B. A reduced-order stochastic model of surface temperature and radiative fluxes

Goals and model formulation: Our goal is to develop a skillful emulator of large-scale sea surface temperature and radiative fluxes that captures both stationary variability and responses to perturbations, enabling the study of their causal interactions. In this work, we focus on ensemble-mean responses, which allows for direct comparison with previous studies. We adopt the following functional form for the neural model:

$$d\mathbf{x} = (\mathbf{L}\mathbf{x} + \mathbf{n}(\mathbf{x}))dt + \mathbf{\Sigma}(\mathbf{x})d\mathbf{W}, \quad (11)$$

where $\mathbf{W}$ is the standard Wiener process. In Appendix C, we propose a practical algorithm to fit the multivariate multiplicative noise $\mathbf{\Sigma}(\mathbf{x})$. Differences with the case of additive noise mainly arise in the reconstruction of the stationary density as further discussed in Appendix D.

As outlined throughout the paper, a central challenge in data-driven modeling lies in selecting preprocessing steps to define the coarse-grained dynamics $\mathbf{x}$ to be modeled. These choices, guided by the study's objectives, are critical for identifying the proper variables, since the climate system involves not only SST and radiative fluxes but also many other interacting variables. This is thus a practical instance of a partially observed state vector. Below, we describe the dataset and preprocessing procedure, which largely follows Ref. [108], with further details provided in that work.

1. Data, preprocessing, and the choice of proper variables

a. Data. We consider outputs from the state-of-the-art coupled climate model GFDL-CM4 (the Geophysical Fluid Dynamics Laboratory's CM4 model [5]) [109], which provides a 600-year preindustrial control run with constant CO_2 . The ocean component is MOM6 [110] with 0.25° horizontal resolution and 75 vertical layers, and the atmosphere is AM4 [111,112] with 100-km resolution and 33 vertical layers. The model in Eq. (11) is trained on this control run. We consider two climate fields: sea surface temperature and the global-mean net radiative flux at the top of the atmosphere. The SST field is restricted to latitudes $[-45^\circ, 45^\circ]$, following earlier studies that highlight the importance of radiative feedbacks in the tropics and subtropics [45,104].

b. A possible choice of proper variables. We coarse-grain the GFDL-CM4 control run in space and time, averaging daily data to monthly means and removing the seasonal cycle to focus on anomalies rather than periodic signals. Considering 1 month averages in TOA fluxes is not arbitrary: The immediate response of TOA fluxes to perturbations in SST is fast (e.g., ~ 20 days in tropical large-scale convective overturning motion in the GFDL-CM4 model) [108]. This averaging filters out higher-frequency processes that mediate surface-TOA coupling, leading to a “more Markovian” system. Given so, *direct* causal links between TOA fluxes at SST patterns can be detected as one time step responses; see Appendix A. We then high-pass filter the data with a cutoff of 10^{-1} yr, removing (1) low-frequency (multidecadal) oscillations, which are poorly sampled even in 600-yr runs, and (2) slow drifts from incomplete ocean equilibration. This choice sacrifices the ability to capture multidecadal variability but ensures robust sampling of subdecadal statistics, yielding a larger effective sample size. Implicitly, the choice assumes that the response of the net radiative flux to SST perturbations occurs within a 10-yr window. Next, we project the SST field $\mathbf{X}(t) \in \mathbb{R}^N$ onto an EOF basis:

$$\mathbf{X}(t) = \sum_{i=1}^n x_i(t) \mathbf{e}_i + \sum_{j=n+1}^N y_j(t) \mathbf{e}_j, \quad (12)$$

and retain the first $n = 20$ modes. The combined steps of monthly averaging, high-pass filtering, and projection over EOF bases define a coarse-grained representation of the SST variability. This choice follows earlier work showing the skill of linear Markov models for SST built on similar preprocessing [4,5]. For radiative fluxes, we focus on the global-mean TOA net flux, as the TOA flux exhibits little spatial structure and we are mainly interested in its average value. The state vector is then defined as $\mathbf{x}(t) \in \mathbb{R}^{n+1}$, consisting of the n

SST principal components concatenated with the global-mean TOA flux.

Crucially, these choices yield an imperfect model that is skillful in some respects—such as capturing probability distributions and ensemble-mean responses to temperature perturbations—but inadequate in others, including autocorrelation structure and responses in ensemble variance. We discuss these successes and limitations in the following sections.

c. Forced simulations. We will test the model's ability to reconstruct changes in global-mean TOA radiative flux given prescribed SST trajectories from two 150-yr-long forced runs: 1pct CO_2 and $4 \times \text{CO}_2$. The 1pct CO_2 experiment imposes a 1% CO_2 increase per year starting from preindustrial conditions, while the $4 \times \text{CO}_2$ experiment abruptly quadruples CO_2 at $t = 0$. A complication is that TOA fluxes respond directly to rising CO_2 , whereas our reduced-order model only uses SST and fluxes and does not explicitly include CO_2 . To address this, we remove the portion of change of the net TOA radiative flux that is attributable solely to rising CO_2 . This allows us to isolate and study the component of radiative feedback that is driven by changes in SST alone. This correction exploits the logarithmic dependence of global-mean forcing on CO_2 concentration [113,114]. As is standard, we subtract 8 W m^{-2} in the $4 \times \text{CO}_2$ run [115], and apply a time-dependent correction $\alpha \log(C_t/C_0)$ in the 1pct CO_2 run, where C_t and C_0 denote CO_2 concentrations in the forced and control simulations.

2. Stationary statistics and forced responses

a. Stationary statistics. Given the $(n+1)$ -dimensional state vector $\mathbf{x}(t) \in \mathbb{R}^{n+1}$ defined above, we train the emulator in Eq. (11) on the long control run of GFDL-CM4. The trained model is then used to simulate a trajectory of equal length (i.e., 600 yr at monthly time step). The stationary probability distributions of the most energetic SST mode, corresponding to the El Niño Southern Oscillation [62], and the net radiative TOA flux are well reproduced [Figs. 6(a) and 6(b)]. Statistics and comparisons with the case of additive noise are reported in Appendix D. The main discrepancies arise in the autocorrelation functions, shown for the first SST mode and the TOA flux in Figs. 6(c) and 6(d). This limitation is further discussed in section “Limitations” below.

b. Forced responses. The pattern-effect problem (Sec. IV A) concerns identifying causal links between SST warming patterns and TOA flux changes. Before using our emulator for Green's functionlike studies to investigate the pattern effect, we first assess its skill as a reconstruction tool. Specifically, we reconstruct the global-mean TOA flux from two prescribed, forced SST trajectories, i.e., from the 1pct CO_2 and $4 \times \text{CO}_2$ simulations. Forced SSTs are projected onto the first $n = 20$ EOFs from the control run, yielding coarse-grained trajectories $\mathbf{x}_{\text{SST}}^f(t)$, f denotes “forced.” Given $\mathbf{x}_{\text{SST}}^f(t)$, the emulator returns the TOA variable $x_{\text{TOA}}(t+1)$ at the next time step, initialized with $x_{\text{TOA}}(0) = 0$. Results [Figs. 6(e) and 6(f)] show that the emulator largely reproduces TOA flux changes in both scenarios. This reconstruction is possible because of the focus on 1-month averages, reducing the time-dependent TOA response, driven by an atmospheric reorganization




FIG. 6. (a) Stationary distributions obtained for the first principal component of the SST field as in the original GFDL-CM4 climate model and as emulated by the neural emulator with multiplicative noise in Eq. (11). Units: nondimensional. (b) Same as panel (a) but for the global-mean radiative flux at the top of the atmosphere. Units: $W m^{-2}$. The stationary distributions have been obtained after simulating the neural emulator for 600 yr. Second row: same as the first row but for the autocorrelation functions. Third row: reconstruction of the net radiative flux at the TOA by the neural emulator in Eq. (11) after constraining the SST degrees of freedom to follow a precomputed trajectory in two forced experiments in the GFDL-CM4 model. The forced experiments considered are the 1pctCO₂ and $4 \times CO_2$; see Sec. IV B 1 for details. Note: the net flux at the TOA variable is negative upward (i.e., leaving the planet).

at submonthly scales, into a one-step prediction. For the 1pctCO₂ experiment, the linear trend in net TOA flux is $-0.043 W m^{-2} yr^{-1}$ in GFDL-CM4 and $-0.034 W m^{-2} yr^{-1}$ in the emulator. We then compute correlations between the reconstructed and simulated radiative flux at the TOA, after removing a linear trend in 1pctCO₂ and an initial 50-yr transient in $4 \times CO_2$. At monthly resolution, correlations are $r = 0.37$ (1pctCO₂) and $r = 0.43$ ($4 \times CO_2$). If yearly averages are considered (as in the current literature, see, e.g., Ref. [44]), these increase to $r = 0.70$ (1pctCO₂) and $r = 0.74$ ($4 \times CO_2$). All correlations are significant at the 5σ level under a null hypothesis of zero correlation, tested via bootstrap resampling with replacement.

3. Interpretability and causal understanding through ensemble mean responses

In the previous section, we showed that the reduced-order neural emulator can reconstruct the net radiative flux at the top of the atmosphere from prescribed changes in the SST field. This amounts to the mapping $\mathbf{x}_{SST}^f(t) \rightarrow x_{TOA}(t+1)$, where $\mathbf{x}_{SST}^f(t) \in \mathbb{R}^n$, $x_{TOA}(t) \in \mathbb{R}$, and $x_{TOA}(0) = 0$. We now ask the following question: *What are the SST patterns that contribute the most to this skillful prediction?* Addressing this question requires moving beyond prediction tasks and identifying the causal influence of individual SST components on radiative fluxes. We frame the problem in terms of linear response

theory, focusing on the direct link $x_{j,SST}(t) \rightarrow x_{TOA}(t+1)$, where $x_{j,SST}(t)$ denotes the j th SST component. We refer the reader to Appendix A for details on the causal interpretation of the response operator.

We apply an impulse to the j th SST component at $t = 0$ and compute the ensemble-mean response of the net TOA flux at $t = 1$, denoted $\delta \langle x_{TOA}(1) \rangle$ (see Appendix A for details). This quantity, referred to as $R_{TOA,j}(1)$, represents the one-step ensemble-mean response and is estimated using 10^6 initial conditions to ensure robust statistics. By focusing on the one-step response, we isolate fast atmospheric adjustments while neglecting the slower oceanic feedbacks (see Sec. 7 in Ref. [108] for details). This setup is thus conceptually closer to atmospheric-only Green's function experiments, where the ocean acts as a prescribed boundary condition rather than a dynamically coupled component of the system.

The vector $\mathbf{R}_{TOA}(1) \in \mathbb{R}^n$ is projected back to the full N -dimensional SST field using the EOFs, $\mathbf{e} \in \mathbb{R}^{n,N}$:

$$\mathbf{S} = \mathbf{R}_{TOA}(1) \cdot \mathbf{e} = \sum_j^n R_{TOA,j}(1) \mathbf{e}_j. \quad (13)$$

At each ocean grid point i , the sensitivity map $\mathbf{S} \in \mathbb{R}^N$ approximates the sensitivity $\frac{\partial x_{TOA}}{\partial x_{SST,i}}$, where x_{TOA} is the global-mean radiative flux at the TOA and $x_{SST,i}$ is the SST at location i .

The sensitivity map $\mathbf{S}$ is shown in Fig. 7. The obtained sensitivity map $\mathbf{S}$ is in qualitative agreement with previous results in response theory (Fig. 2(a) in Ref. [108]), with other




FIG. 7. Sensitivity map obtained using the neural emulator in Eq. (11). At each SST grid point i , we show the one-time-step response in the ensemble mean (S) of the global-mean net radiative flux at the TOA following an impulse perturbation in SST at that location. The computation is carried out in a 21-dimensional reduced space spanned by empirical orthogonal functions, and the maps represent the projection of these low-dimensional results back to the full grid.

data-driven approaches (see Fig. 7(b) in Ref. [116] and Fig. 2 in Ref. [117]), and with the one obtained by directly perturbing the SST boundary condition of atmosphere-only climate models (see third row of Fig. 2 in Ref. [44] and Fig. 5(a) in Ref. [45]). In the latter case, the maps exhibit fewer spatial features due to the coarse resolution of the perturbation patches (spanning tens of degrees) in numerical models experiments. For additional statistical reconstruction approaches to the pattern effect that yield structures similar to those in Fig. 7, see the recent analysis by Fredericks *et al.* [118].

The dominant feature of S is a dipole in the tropical Pacific: negative sensitivity in the west and positive sensitivity in the east [45]. This implies that eastern Pacific warming produces a positive radiative feedback, while western Pacific warming produces the opposite. The mechanisms are well established [107,119]: In the east, high climatological low-level cloud cover decreases under warming, reducing reflected shortwave radiation and yielding positive sensitivity. In the west, deep convection transports surface warming vertically and then horizontally via gravity waves, strengthening the inversion over the eastern Pacific and other stratocumulus regions, increasing low-level cloud fraction and reflected shortwave radiation, thus producing a negative feedback. This analysis confirms that the accurate reconstructions in the forced regime [Figs. 6(e) and 6(f)] arise from physically consistent mechanisms.

4. Limitations and a few ways forward

Our coarse-grained representation enabled the development of a data-driven model to study the direct causal interaction $x_{\text{SST}} \rightarrow x_{\text{TOA}}$ through perturbation experiments, while also capturing the stationary distribution well. However, these coarse-graining choices also introduce some drawbacks. We discuss these limitations below, further delineating the model's domain of validity and outlining possible ways forward:

(1) The model does not fully capture the decay of auto-correlations. For example, the first SST mode in Fig. 6(a) suggests residual non-Markovian behavior. This is not an issue for our focus on *direct* causal links (responses at the shortest time step), but it would matter for responses over longer time horizons. Remedies include: (a) expanding the

state vector by including slower oceanic variables and more EOF components, (b) incorporating memory terms in the stochastic parametrization [7], or (c) further coarse graining the temporal resolution (e.g., 3-month averaging). Such choices will depend on the intended scientific application.

(2) We found that responses in ensemble variance are sensitive to retraining the multiplicative noise term, so we do not present these results. The fitting algorithm in Appendix C requires far more data than available here (only 7200 samples for training and validation), to achieve high accuracy. In this data-scarce setting, exploring simple analytical forms for the multiplicative noise may be beneficial and it is left for future work. Importantly, data-driven approaches, such as the fluctuation-dissipation theorem [55,56], can provide an approximate benchmark for responses. This allows calibration of models with different stochastic terms, particularly to test variance responses where direct benchmarks are often unavailable. This strategy aligns with recent proposals such as Ref. [18] and will be explored in future work.

(3) While our emulators are empirically stable, stability is not guaranteed, which may matter if tested against large perturbations. Future work could impose additional constraints on the neural architectures, similar to previous approaches for physically constrained regression models as in Refs. [7,11].

5. Main takeaways

In this section, we developed a stochastic neural emulator that captures the joint evolution of the sea surface temperature field and the global-mean net radiative flux at the top of the atmosphere in the GFDL-CM4 climate model. Addressing a specific, well-posed problem, proved crucial in guiding the construction of the model: It enabled us to build on previous literature, restrict attention to just two variables and a limited range of spatial and temporal scales. These simplifications led to a stable and computationally efficient emulator that responds to external forcing and can be used to investigate causal relationships via perturbation experiments. Such experiments focused on very large ensembles (10^6 trajectories), currently impossible to perform with climate models.

V. DISCUSSION

In summary, the first part of this work discussed important limitations of purely data-driven models of complex climate systems. We highlighted the potential of combining traditional reduced-order modeling approaches, including state-dependent (multiplicative) stochastic parametrizations, with recent advances in neural networks to study large-scale, low-frequency climate dynamics. We emphasized the need to evaluate such models against both stationary and perturbed statistics. In this context, we discussed how response theory provides rigorous guidance to move beyond stationary evaluation and probe causal mechanisms. These ideas were then illustrated through a climate-dynamics problem focusing on inferring causal relations between surface warming and changes in net radiative fluxes. Given careful variable selection, reduced-order neural stochastic models can capture both stationary and perturbed statistics in a computationally efficient way. Below, we expand on these points and outline directions for future work.

A. Limits of data-driven-only climate emulators and the role of stochastic approaches

In the first part of this paper, we distinguished two possible scenarios in data-driven modeling:

(1) Unknown equations, fully observed state vector: In this idealized setting, it is possible to infer a model from data alone that reproduces the system's stationary statistics and responds accurately to small external perturbations.

(2) Unknown equations, partially observed state vector: This more realistic case presents fundamental difficulties, even with access to high-quality data and powerful neural model architectures. The most critical limitation is (a) identifying an appropriate coarse-grained representation; once this is determined, (b) the cumulative effect of unobserved degrees of freedom needs to be accurately parametrized. Learning such parametrizations from data remains a central challenge, and we showed that multiplicative noise can be advantageous over additive noise, though it requires a larger dataset for reliable training.

These results highlight intrinsic limitations of purely data-driven modeling of high-dimensional, multiscale climate systems, where the full set of relevant variables is unobserved or even unknown. This is especially true for emulators built to capture real-world dynamics from observational data, where observations are inherently partial. In this context, the “correct” choice of variables, scales, and stochastic parametrizations depends ultimately on the specific scientific goals [120]. This perspective helps explain the continued success of linear inverse models proposed by Penland [3], which rely on careful selection of climate fields and preprocessing choices to capture targeted aspects of climate variability; see, for example, Refs. [4,5,121] among others.

We emphasize that the analysis presented here applies to data-driven models of turbulent geophysical systems that must capture both stationary statistics and responses to external perturbations, as in (causal) Green's function experiments [44]. For Earth system models, these issues mainly concern autoregressive climate emulators, particularly coupled atmosphere-ocean systems, and should not be confused with short-term weather forecasting using neural networks (e.g., Refs. [26,122]), where long-term stability and forced responses are less relevant. Our study also does not address statistics-based emulators for climate change projections, which have proven effective for investigating responses of key climate variables to CO₂ forcing [123–125].

Future work: The discussion above suggests two possible directions for future research in autoregressive emulators of climate. First, a stochastic formulation of neural autoregressive climate emulators is preferable to a purely deterministic one. Probabilistic formulations align more closely with the inherently stochastic and uncertain nature of climate dynamics [126]. Recent advances in weather emulation already point in this direction such as FourecastNet3 [26], GenCast [26], and the work of Alet *et al.* [127]. Interesting directions also include probabilistic emulation in the framework of stochastic interpolants [128–130]. Second, the ability of climate emulators in capturing forced responses in the full probability distribution, akin to Green's function experiments, remains largely unexplored, especially for responses in higher-order

moments. A promising way forward is to evaluate architectures on their ability to capture the mean and variance response to external perturbations in a hierarchy of simplified models relevant for large scale climate variability, such as the two-scale Lorenz '96 system [131]. These controlled test beds enable rigorous quantification of forced responses and could inform the design of future architectures for data-driven climate modeling.

B. Response theory as a theoretical framework for model evaluation

We argued that the impulse response operator is a building block that a model should capture in order to reproduce causal interactions across degrees of freedom and to accurately respond to external perturbations (Sec. II B and Appendix A). This perspective enables us to move beyond evaluation metrics based solely on stationary statistics and instead consider the effect of external interventions. Assessing impulse responses is straightforward in the case of data-driven emulators built from simple numerical models, i.e., models that allow us to simulate thousands of ensemble members, as in the first part of this study. In such cases, one can actively perturb the reference model to generate a ground truth for the response operator. However, for models constructed from observational data, such direct perturbations are not possible. In this context, one can leverage the fluctuation-dissipation theorem to predict the response operator from data alone [20,91]. Traditionally, the quasi-Gaussian (qG) approximation have yield reliable estimates of impulse responses in the ensemble mean [54,132]. Recent advances in score-based generative models have shown the potential to extend these predictions to responses in higher-order moments [133].

Future work: In the absence of direct perturbation experiments, FDT provides a rigorous bridge between stationary and forced dynamics [55,56]. With sufficiently long integrations of current climate emulators, one can readily focus on representative large-scale modes and estimate their impulse responses from the stationary record, thereby extending evaluation beyond stationary statistics at minimal computational cost. A recent example illustrating this approach in the climate context can be found in Ref. [21]. Importantly, the analysis reveals the need to test emulators in their ability to respond to higher-order moments. Authors of Ref. [133] have shown that coupling the FDT with generative models provides a flexible framework to estimate responses in the full probability distribution. Moreover, the response operator at the shortest time step encodes *direct* causal links, i.e., the input-output relationships across variables in a reduced-order emulator (see Appendix A). Future work will leverage this information to suppress spurious, noncausal input-output dependencies during training, similarly to the recent work of Chen and Liu [68], leading to causally constrained, sparse reduced-order models.

C. A realistic application: Radiative responses to temperature perturbations

Building on the theoretical insights introduced in the first part, the second part of the paper presented a stochastic neural emulator for the coarse-grained dynamics of sea surface

temperature and net radiative flux at the top of atmosphere. The model was trained on a long unforced simulations of a coupled climate model and designed with a specific physical task in mind: capturing radiative feedbacks to SST anomalies. The emulator focuses on the evolution of 20 principal components of the SST field and on the scalar global-mean, net radiative flux; unobserved degrees of freedom are parametrized as multiplicative noise through an additional neural network. The emulator largely reproduces the stationary statistics of the system and allows us to reconstruct TOA radiative fluxes in two forced simulations. We inferred the direct causal links between temperature modes and TOA fluxes through a perturbation approach, akin to Green's functionlike experiments with climate models [44], targeting the so-called "pattern effect" [45]. Such responses are then projected back onto the EOFs to obtain maps of sensitivity in mean and variance change of the radiative flux at the TOA to SST warming perturbations. We discussed similarities of the sensitivity map for the ensemble mean with recent studies, indicating that the emulator's skill in reconstructing TOA fluxes under forcing arises from physically meaningful mechanisms. We also identified limitations. In particular, the variance response is sensitive to the training of the multiplicative noise component, restricting the emulator's reliability in perturbation experiments beyond those targeting the ensemble mean. We outline possible paths forward to resolve this issue. These findings highlight the potential of reduced-order neural emulators to perform thousands of perturbation experiments, eventually shifting the focus of climate feedback analysis from mean responses to the responses of probability distributions.

VI. CONCLUSION

Modeling the climate system from data alone is constrained by observing only a subset of variables of a multiscale, high-dimensional dynamical system. This raises the fundamental challenge of identifying the proper variables (both fields and spatiotemporal scales), a fundamental problem in physics emphasized in classical work [134] and revisited in recent studies [135]. Such choices are inherently non-unique, especially because the climate system lacks clear space- and time-scale separation, and are therefore largely determined by the study's scientific goal rather than by the pursuit of an unattainable perfect model [120]. In contrast to numerical weather prediction, which demands high spatial and temporal resolution, low-frequency climate variability is dominated by a few recurrent spatiotemporal patterns, or modes of variability, and their interactions. Thus, at least for certain tasks, the multiscale structure of the climate system can be an advantage, allowing us to focus on the coarse-grained dynamics while treating unresolved, fast processes stochastically. Traditional reduced-order models stand out from general-purpose emulators by targeting specific climate problems from the outset, guided by prior literature, physical insight, and a clear understanding of the relevant variables and scales at play. Such data-driven modeling approach has a long history in climate science [2], though most efforts have emphasized reproducing stationary statistics or short-term forecasting [3,4,85] rather than forced responses. A promising perspective is to build on such established reduced-order ap-

proaches, including those developed by Majda and Qi [18,19], while leveraging flexible neural networks architectures to discover low-dimensional latent representations [17,78] and fit deterministic and stochastic components. Furthermore, tools from response theory such as the FDT [21,133,136] can be used to guide model development and eventually lead to reduced-order models that efficiently capture responses in the mean and variance to external perturbations, advancing causal inference and attribution analyses in climate science.

ACKNOWLEDGMENTS

Most of the ideas discussed in this work stem from collaborations and exchanges with colleagues and friends. I am grateful for the supportive and engaging environment of the Zanna Group at Courant, as well as for interactions within the broader M²Lines project. I thank Laure Zanna and Carlos Fernandez-Granda for their feedback during early stages of the project. I am especially grateful to Pasha Perezhgin, Rory Basinski, Chris Pedersen, and Andre Souza for their thoughtful feedback and for many enriching discussions on data-driven modeling of the Climate system. I acknowledge an insightful email exchange with Prof. Tony Roberts, who drew my attention to recent advances on stochastic model reduction. I thank two anonymous reviewers for their very useful and constructive comments. This research received support through Schmidt Sciences, LLC, under the M²Lines project. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: RESPONSE OPERATOR: PRACTICAL ESTIMATION AND CAUSAL INTERPRETATION

1. Inferring the response operator in practice

Given a model $d\mathbf{x} = \mathbf{f}(\mathbf{x})dt + \Sigma(\mathbf{x})d\mathbf{W}$, we build the impulse response operator as follows:

(1) We simulate a very long trajectory of the model starting from a random initial condition and remove an initial transient. We, then, sample N_e random points from the simulation to define an ensemble of N_e initial conditions on the model's attractor. N_e should be very large (i.e., $N_e \gg 1$) in order to sample the whole attractor and approximate averages over the invariant distribution.

(2) For each one of the N_e initial conditions, we impose an impulse perturbation Δ_j to the degree of freedom x_j at time $t = 0$. The amplitude of the perturbation should be theoretically infinitesimally small, ensuring linearity of the response even in nonlinear systems. In practice we do as follows: we consider the long time series of $x_j(t)$ from the long control integration above and define $\Delta_j = 10^{-2}\sigma_j$, where σ_j is the standard deviation of time series $x_j(t)$.

(3) Therefore, for a given initial condition $\mathbf{x}(0)$, we simulate two trajectories: a control trajectory without perturbation and a perturbed one, where an impulse perturbation Δ_j has been imposed on the j th degree of freedom at time $t = 0$. This procedure is repeated in parallel for all initial conditions, resulting in an ensemble of N_e paired of control and perturbed trajectories.

(4) At each time t , we then estimate the time-dependent ensemble average $\langle A(x_k(t)) \rangle$ of an observable $A(x_k(t))$ for both the perturbed and unperturbed run. We refer to the perturbed and control ensemble average, respectively, as $\langle A(x_k(t)) \rangle_p$ and $\langle A(x_k(t)) \rangle$. The observables considered in this study are going to be (a) $A(x_k(t)) = x_k(t)$ and (b) $A(x_k(t)) = (x_k(t) - \mu_k(t))^2$, with $\mu_k(t)$ representing the time-dependent mean of the distribution, therefore quantifying the response in ensemble mean and variance.

(5) We define the impulse response operator for observable $A(x_k(t))$ as

$$R_{k,j}(t) = \frac{\langle A(x_k(t)) \rangle_p - \langle A(x_k(t)) \rangle}{\Delta_j}.$$

2. Causal interpretation of the response operator

Causality in natural systems is inferred through external interventions: An agent perturbs the system and the response is measured [87,88]. Climate modeling follows the same principle through Green's function experiments, where one variable is perturbed and the response of another is evaluated [44,90]. Linear response theory formalizes this principle, providing a rigorous framework for causal analysis both in models and directly from data [20,21,133].

Because responses are inherently time dependent, causal links through responses can be distinguished in three forms, all encoded in the time-dependent response operator: (1) direct links, (2) indirect links, and (3) cumulative links (or causal sensitivity). In what follows, we use the triad system in Eq. (6) as a simple example to illustrate these ideas in a clear and unambiguous way. Finally, we summarize recent approaches for estimating responses directly from data and refer to Baldovin *et al.* [20] for details on the theoretical foundations of this approach.

Consider a high-dimensional, stochastic dynamical system of the kind

$$\frac{dx_i}{dt} = f_i(x_1, x_2, \dots, x_N) + \text{stochastic forcing};$$

with $i = 1, 2, \dots, N$. (A1)

In this setting, the existence of a *direct* causal link $x_j \rightarrow x_k$ is determined by whether x_j appears in the arguments of f_k . More formally, the *absence* of a direct causal link $x_j \rightarrow x_k$, $j \neq k$ is expressed by

$$\left\langle \frac{\partial f_k(\mathbf{x})}{\partial x_j} \right\rangle = 0.$$

We adopt the view that causal estimates in real-world, high-dimensional stochastic systems are most naturally, and most practically, formulated as ensemble averages, making the ensemble-averaged Jacobian the appropriate structural object for direct causal links. As a concrete example, we return




FIG. 8. (a) Functional form of the triad system in Eq. (6). (b) Graph structure of the triad system. Note from panel (b) that causality in dynamical systems can involve cycles, in contrast with the acyclic structures often assumed in applications outside physics; see, e.g., Ref. [137].

to the triad system in Eq. (6), for which the underlying direct causal links are known exactly. This example provides a straightforward, intuitive illustration of how perturbation experiments reveal causal interactions. Figure 8 reports both [panel (a)] its functional form and [panel (b)] its causal graph, which encode the same information. Causal relations can then be probed by adding independent impulse perturbations to each degree of freedom and record the system's response. Responses are computed up to a long time horizon $\tau_{\max}$, with discrete times $t = 0, dt, 2dt, \dots, \tau_{\max}$. Since we are interested in causal relations, it suffices to consider the identity observable $A(x_k(t)) = x_k(t)$, i.e., the response of the ensemble mean [20]. We now clarify the three different types of causal links mentioned at the start of this Appendix.

a. Direct causal links

The direct causal links are encoded in the response operator $\mathbf{R}(dt)$ evaluated at the shortest time step $t = dt$. The intuition follows directly from the graph representation in Fig. 8(b) and is rather straightforward.

(1) Any perturbation applied to x_2 will affect, on average, x_2 itself and x_1 after one time step, but not x_3 . Hence, we infer $\langle \partial f_3(\mathbf{x}) / \partial x_2 \rangle = 0$.

(2) A perturbation applied to x_3 will affect, on average, x_3 and x_1 , but not x_2 , so $\langle \partial f_2(\mathbf{x}) / \partial x_3 \rangle = 0$.

(3) By contrast, a perturbation to x_1 propagates to all three variables within a single step, indicating direct causal links from x_1 to both x_2 and x_3 .

b. Indirect causal links

Causal relations in physical systems extend beyond direct interactions. A variable x_j may affect another variable x_k indirectly through an intermediate variable x_i , forming a path $x_j \rightarrow x_i \rightarrow x_k$. Such indirect links are central in climate dynamics [138] and in complex systems more broadly [139]. Within the response framework, these paths are revealed by the time-dependent operator $R_{k,j}(t)$: The first nonzero response identifies the shortest causal path in the graph connecting x_j to x_k . Crucially, if $R_{k,j}(t) = 0$ for all time t , then no external intervention on x_j can ever influence the behavior of x_k .

c. Cumulative causal links or causal sensitivity

Finally, Baldovin *et al.* [20] introduced a cumulative (in time) measure of causation between two variables, defined as $D_{k,j} = \int_0^{t_{\max}} R_{k,j}(t)dt$. Such cumulative link can be also referred to as “causal sensitivity.” Such a measure is especially relevant in practical experiments with high-dimensional models where interactions unfold across multiple spatial and temporal scales. A concrete example using climate models is provided by the approach of Bloch-Johnson *et al.* [44], where causal links are inferred between patterns of surface warming and the net radiative flux at the top of the atmosphere. Their method reconstructs responses to perturbations by (1) applying surface warming perturbations, (2) integrating an atmospheric-only model with and without the perturbation, and (3) computing the average (in time) difference between the two runs after removing an initial transient. In this case, the causal link emerges precisely as a cumulative (in time) atmospheric response to the surface perturbations.

We emphasize the crucial role of ensemble averages in establishing the general validity of response theory. Evaluating responses along single trajectories of nonlinear systems is strongly affected by chaos, restricting the applicability of these ideas to short times and very small perturbations. In contrast, focusing on observables, i.e., quantities averaged over the invariant density, ensures the robustness of the theory and largely removes the dependence on chaos. For a detailed discussion of this point, we refer the reader to Ref. [140]. Future efforts to estimate responses in climate science, which often rely on perturbing single trajectories, would therefore benefit from focusing on ensemble-averaged observables.

d. Response operator from data alone

Given the assumption of Markovianity it is possible to infer the system’s response operator from a long stationary trajectory using the fluctuation-dissipation theorem [20]. The FDT formula provides a link between stationary and perturbed dynamics by predicting forced responses from stationary statistics alone [141]. Specifically,

$$R_{k,j}(t) = -\left\langle A(x_k(t)) \frac{\partial \ln \rho(\mathbf{x})}{\partial x_j} \right\rangle_{\mathbf{x}(0)}, \quad (\text{A2})$$

where $A(x_k(t))$ is a general observable and $\rho(\mathbf{x})$ is the stationary distribution. $\langle \cdot \rangle$ refers to ensemble average, approximated through temporal averages. $\frac{\partial \ln \rho(\mathbf{x})}{\partial x_j} \Big|_{\mathbf{x}(0)}$ represents the so-called “score function” evaluated at the initial state, i.e., on the attractor. Therefore, in principle the response operator can be computed from data if the score $\nabla \ln \rho(\mathbf{x})$ is available. In practice, the score is often inaccessible, and most applications rely on the qG approximation, by replacing $\rho(\mathbf{x})$ with a Gaussian distribution $\rho^G(\mathbf{x})$ with same mean and covariances [55,142]. For responses in the ensemble mean, the qG approximation achieves high skill even for nonlinear systems [20,132,133]. Recent work [133] has demonstrated that generative modeling techniques can be used to efficiently estimate the score directly from data, leading to high skill in computation of higher-order responses in strongly nonlinear regimes. Other interesting proposals have been recently formulated in the

context of Koopman theory and can be used for future exploration of forced responses from data alone [136,143,144].

It is worth emphasizing that the central limitation discussed in the main text, i.e., the need to choose the proper variables for the system to be Markovian, remains relevant for the inference of response operators via FDT as well (and for any causal inference methodology); see the discussion in Ref. [20].

APPENDIX B: RELEVANCE OF THE TRIAD SYSTEM FOR GEOPHYSICAL TURBULENCE

Here, we motivate the choice of the triad model analyzed in Sec. III as a building block for geophysical turbulent flows. We further present and test the associated analytical stochastic model for the slow variable x_1 (as discussed in Sec. III C 1). For detailed mathematical derivations and additional examples, we refer to the original work by Majda *et al.* [18,53,55,69,70].

1. Link to geophysical turbulence

Projecting many geophysical turbulent flows, such as quasigeostrophic models, onto a complete set of orthogonal basis functions (spectral or empirical modes such as EOFs), gives rise to the following generic form for the modes’ amplitudes:

$$d\mathbf{u}/dt = \mathbf{F} + \mathbf{L}\mathbf{u} + \mathbf{B}(\mathbf{u}, \mathbf{u}), \quad (\text{B1})$$

where $\mathbf{F}$ is an external forcing, $\mathbf{L}$ is a linear operator, and $\mathbf{B}(\mathbf{u}, \mathbf{u})$ is a nonlinear, quadratic operator. The linear operator $\mathbf{L} = \mathbf{L}' + \mathbf{D}$, is formed by a skew-symmetric operator $\mathbf{L}'$ and by a diagonal dissipation matrix $\mathbf{D}$ with negative real entries. The quadratic nonlinearity should be energy conserving, satisfying $\mathbf{u} \cdot \mathbf{B}(\mathbf{u}, \mathbf{u}) = 0$ [55,69,145]. Eq. (B1) is the typical abstract form of dry cores of atmospheric models [76]. Galerkin truncations of turbulent systems often reduce to this structure, as in the Lorenz ’63 and Charney-DeVore models [77,145,146]. $\mathbf{u} \in \mathbb{R}^N$ represents the full turbulent geophysical system, with N infinite in principle and finite but very large in applications ($N \gg 1$). The state $\mathbf{u}$ is then decomposed into slow and fast components, $\mathbf{u} = (\mathbf{x}, \mathbf{y})$, i.e., $\mathbf{x}$ represents the first n components of $\mathbf{u}$ and $\mathbf{y}$ the next $N - n$.

We consider the simple case of $N = 3$: one large-scale resolved mode x_1 and two fast modes $\mathbf{y} = (x_2, x_3)$. We then project the three modes onto the dynamics in Eq. (B1). This gives rise to three equations, defined by a forcing term, linear and quadratic self- and crossinteraction across slow and fast timescales. The triad model in Eq. (6) is then obtained by approximating the quadratic self-interactions of the fast modes as a damping and stochastic forcing, as $-\gamma_i \epsilon^{-1} x_i + \sigma_i \epsilon^{-1/2} \xi_i$, with $i = 2, 3$. ξ_i is a standard Gaussian white noise process. ϵ sets the timescale separation between the slow and fast modes. The key assumption is that the detailed nonlinear self-interactions among fast modes are not important if the goal is to capture the coarse-grained slow dynamics. This approximation transforms the original deterministic equations into a stochastic system. The specific version used in Eq. (6) is further simplified by setting to zero a few interactions, such as the linear cross and self-interactions of the fast modes as well as the external forcing terms.

Therefore, the relevance of the triad system [Eq. (6)] to this study is that it provides a minimal multiscale setting with linear and quadratic nonlinear interaction of the kind that naturally arise when geophysical turbulent flows are expressed in terms of orthogonal bases such as EOFs. Furthermore, the model is characterized by fast decays of correlations and generates PDFs that are regular and quasi-Gaussian, as often seen in the analysis of large-scale modes of climate variability [91]. For additional details, see Ref. [53] and Sec. 4 of Ref. [18].

2. Analytical stochastic model for the slow variable x_1

Starting from the triad model in Eq. (6), Majda *et al.* [69] showed that an analytical reduced-order equation for the slow variable x_1 can be rigorously derived in the limit $\epsilon \rightarrow 0$ using standard projection techniques [53–55]. The mode reduction predicts the functional form of the dynamics governing the resolved, slow variable [69–72]. The resulting stochastic equation is

$$dx = (\tilde{F} + ax + bx^2 - cx^3)dt + (A - Bx)dW_2 + \sigma dW_3, \quad (\text{B2})$$

where the coefficients are expressed in terms of parameters of the triad system in Eq. (6) as $a = L_{11} + \epsilon(\frac{L_{12}^2}{2\gamma_2^2} - \frac{L_{12}^2}{\gamma_2} - \frac{L_{13}^2}{\gamma_3})$, $b = -2\epsilon\frac{L_{12}L}{\gamma_2}$, $c = \epsilon\frac{L^2}{\gamma_2}$, $B = -I\epsilon^{1/2}\frac{\sigma_2}{\gamma_2}$, $L = -B\frac{L_{12}}{I}$, $\tilde{F} = -\frac{AB}{2}$, $\sigma = \epsilon^{1/2}L_{13}\frac{\sigma_3}{\gamma_3}$. W_2 and W_3 are independent Wiener processes. Interestingly, the equation for the coarse-grained variable x_1 displays a structure richer than that of the full triad system in Eq. (6). In addition to the original linear and quadratic interactions, the reduced equation includes following terms: an effective forcing $\tilde{F}$, a cubic nonlinearity, and both additive and multiplicative noise. The additive noise results from the linear interactions across fast modes, while the multiplicative noise reflects the advection of the slow mode by the fast modes [76]. Thus, Eq. (B2) demonstrates that multiplicative noise arises from theoretical considerations, rather than being introduced ad hoc, providing a firm basis for its use in reduced-order stochastic climate models. The scalar model above has been used to test the fluctuation-dissipation theorem in climate dynamics [91,133] as well as for inverse modeling of the leading principal component of a quasigeostrophic model and the observed variability associated with the North Atlantic Oscillation, demonstrating its relevance for theoretical studies in climate dynamics.

We evaluate the model in Eq. (B2) as done in Sec. III, so to reproduce the stationary statistics of the slow variable x_1 as well as its response to perturbations as in the original triad system in Eq. (6). The results are reported in Fig. 9. The analysis shows that, given $\epsilon \ll 1$, it is indeed possible to build a Markovian model of the slow variable x_1 reproducing both stationary and perturbed statistics as in the original system.

Finally, as shown in this Appendix, rigorous model reduction frameworks, such as those developed by Majda *et al.* [70] and by Roberts [147,148], among others [2], provide invaluable theoretical guidance for multiscale climate modeling. While these strategies were traditionally developed for systems with known equations of motion, they may also in-




FIG. 9. In each panel: “ x_1 Triad” represents statistics of x_1 in the original triad model in Eq. (6); “ x_1 scalar” represent statistics of x_1 in the scalar reduced model in Eq. (B2). (a) Stationary probability distribution. (b) Response in ensemble mean. (c) Response in ensemble variance. Impulse response functions are computed using an ensemble of $N_e = 10^7$ members.

spire future work aimed at incorporating such principles into data-driven-only emulators.

APPENDIX C: TECHNICAL AND PRACTICAL DETAILS ON THE MODEL FORMULATIONS

1. Multivariate data-driven stochastic models

Consider an N -dimensional dynamical system described by a stochastic differential equation $d\mathbf{x} = \mathbf{f}(\mathbf{x})dt + \Sigma(\mathbf{x})d\mathbf{W}$. This equation can be discretized by the forward Euler method as follows:

$$\Delta\mathbf{x} = \mathbf{x}(t + \Delta t) - \mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t))\Delta t + \Sigma(\mathbf{x})\sqrt{\Delta t}\xi(t), \quad (\text{C1})$$

where $\xi(t)$ is a Gaussian white noise with zero mean, unit variance, and $\sqrt{\Delta t}\xi(t)$ is the increment of the Wiener process $\Delta\mathbf{W} = \mathbf{W}(t + \Delta t) - \mathbf{W}(t)$. Without loss of generality, we rescale time units by the time sampling interval such that a nondimensional time increment is $\Delta t = 1$ [6–8,149]. The strategy is then to fit a rescaled version of Eq. (C1) from observational data alone. Such rescaled equation takes the following form:

$$\Delta\mathbf{x} = \mathbf{x}(t + 1) - \mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t)) + \Sigma(\mathbf{x})\xi(t). \quad (\text{C2})$$

We proceed in two steps:

(1) Fitting the drift: We compute $\Delta\mathbf{x} = \mathbf{x}(t + 1) - \mathbf{x}(t)$ and fit the deterministic dynamics such that $\Delta\mathbf{x} \approx \mathbf{f}(\mathbf{x}(t))$. In the main paper, we considered two cases:

(a) Linear drift: In this case, given a trajectory $\mathbf{x}(t) \in \mathbb{R}^N$, we need to fit a matrix $\mathbf{L} \in \mathbb{R}^{N,N}$ such that $\Delta\mathbf{x} \approx \mathbf{L}\mathbf{x}(t)$

as close as possible. In practice, this is solved as a linear regression problem by minimizing the mean squared error (MSE). We use the linear regression module from the scikit-learn library [150].

(b) Nonlinear drift: We consider the following form: $\Delta \mathbf{x} \approx \mathbf{f}(\mathbf{x}(t)) = \mathbf{L}\mathbf{x}(t) + \mathbf{n}(\mathbf{x}(t))$, where $\mathbf{L} \in \mathbb{R}^{N,N}$ is the linear operator and $\mathbf{n}(\mathbf{x}(t))$ is the nonlinear operator parametrized by a multilayer perceptron. Both operators are fitted jointly by minimizing the mean squared error using a version of the gradient descent algorithm (Adam, [151]) with weight decay. The MLP considered here always consisted of one hidden layer. The number of neurons in the hidden layer depended on the application: we considered 30 neurons in the experiments with the triad model in Sec. III B 1 and 1000 neurons in the real-world application in Sec. IV. We played with more hidden layers but found a comparable performance, and sometimes a larger network also led to instabilities. The activation function used is the swish function (SiLU), which we found yields smoother responses to external perturbations. In contrast, the commonly used ReLU activation produced noisier responses. Furthermore, to ensure high accuracy in perturbation-response experiments, all data were processed using double-precision floating-point format.

(2) Fitting the stochastic term: Having fitted the deterministic dynamics $\mathbf{f}$, we compute the residuals $\mathbf{r}(t) \in \mathbb{R}^N$ as $\mathbf{r}(t) = \Delta \mathbf{x} - \mathbf{f}(\mathbf{x}(t))$. We then construct a suitable stochastic parametrization for this residual term. In the paper, we focus on two cases: additive and multiplicative noise.

(a) Additive noise: Given the residuals $\mathbf{r}(t)$, we proceed by computing their covariance matrix $\mathbf{C} \in \mathbb{R}^{N,N}$ as $\mathbf{C} = \langle \mathbf{r}\mathbf{r}^T \rangle$, where brackets denote time averaging. We then consider the Cholesky decomposition of $\mathbf{C}$ to derive a lower triangular matrix $\mathbf{\Sigma} \in \mathbb{R}^{N,N}$ such that $\mathbf{C} = \mathbf{\Sigma}\mathbf{\Sigma}^T$. The full fitted equation solved by the emulator can then be written as

$$\mathbf{x}(t+1) = \mathbf{x}(t) + \mathbf{f}(\mathbf{x}(t)) + \mathbf{\Sigma}\xi(t), \quad (\text{C3})$$

and we remind the reader that $\xi(t) \in \mathbb{R}^N$ is a Gaussian white noise with zero mean, unit variance.

(b) Multiplicative noise: Here, we consider state-dependent noise. To approximate it in a computationally efficient manner, we proceed as follows:

(i) Given the residuals $\mathbf{r}(t)$ at each time t , we compute their outer product $\mathbf{C}(t) = \mathbf{r}(t)\mathbf{r}(t)^T$, which is a matrix of rank one.

(ii) We train a neural network predicting a lower triangular matrix $\mathbf{\Sigma}(\mathbf{x}(t))$ with positive diagonal, which can be interpreted as a Cholesky factor, such that $\mathbf{C}(t) \approx \mathbf{\Sigma}(\mathbf{x}(t))\mathbf{\Sigma}(\mathbf{x}(t))^T$ as close as possible.

In practice, the loss function minimizes the mean squared error between $\mathbf{C}(t)$ and $\mathbf{\Sigma}(\mathbf{x}(t))\mathbf{\Sigma}(\mathbf{x}(t))^T$ at each batch. A minimizer of the MSE loss approximates the conditional mathematical expectation $\mathbb{E}[\mathbf{C}(t)|\mathbf{x}(t)]$, that is the covariance matrix averaged over all residuals corresponding to the same state vector, as a function of this state vector, and can be seen as a generalization of a method presented in Ref. [152] to multiple dimensions. The positivity of the diagonal terms of the lower triangular matrix $\mathbf{\Sigma}(\mathbf{x}(t))$ is enforced through a

softplus activation function. Note that this (high-dimensional) multiplicative noise parametrization is only applied in the real-world-like case, where data are scarce: To prevent overfitting, we use early stopping during training, halting after only five epochs. The final fitted equation solved by the emulator is

$$\mathbf{x}(t+1) = \mathbf{x}(t) + \mathbf{f}(\mathbf{x}(t)) + \mathbf{\Sigma}(\mathbf{x}(t))\xi(t), \quad (\text{C4})$$

where again $\xi(t) \in \mathbb{R}^N$ is a Gaussian white noise with zero mean, unit variance.

Furthermore, even in the real-world case of sea surface temperature fields and radiative fluxes, reducing the dimensionality to a few, $N = 21$, components led to a simplified computational problem where all neural networks could be trained on CPUs and on a personal computer.

2. Scalar data-driven stochastic models

Consider a trajectory of length T of a 1-dimensional dynamical system. As in the Appendix above, we assume that this scalar time series can be modeled as a discrete stochastic process with time stepping $\Delta t = 1$, therefore rescaling our time stepping by the sampling interval. The scalar equation we aim to fit is therefore

$$\Delta x = x(t+1) - x(t) = f(x(t)) + \sigma(x)\xi(t). \quad (\text{C5})$$

The training of the deterministic (i.e., drift) dynamics follows exactly the same procedure of the multivariate case. Even the network considered are the same, with one hidden layer and 30 neurons. So we refer the reader to the Appendix above regarding the fitting of the deterministic dynamics and proceed in focusing on the stochastic case.

Fitting the stochastic term

(1) Additive noise: We compute the residual time series $r(t) = \Delta x - f(x(t))$ from the drift term. We then proceed in computing the standard deviation σ of such residual time series. The full fitted equation solved by the emulator can then be written as

$$x(t+1) = x(t) + f(x(t)) + \sigma\xi(t), \quad (\text{C6})$$

and we remind the reader that $\xi(t)$ is a scalar Gaussian white noise with zero mean, unit variance.

(2) Multiplicative noise: We again consider the residual time series $r(t)$ from the drift term. We then compute at each time step t the square of the residual at that time step, i.e., $r(t)^2$. We then train a neural network to have $g(x(t)) \approx r(t)^2$ by minimizing the mean squared error [152]. The neural network is again an MLP with one hidden layer with 30 neurons. Differently, from the neural networks considered for the drift, here we added a softplus activation function after the last layer of the network to ensure that the variance estimates are nonnegative. The final fitted equation solved by the emulator is

$$x(t+1) = x(t) + f(x(t)) + \sqrt{g(x(t))}\xi(t). \quad (\text{C7})$$




FIG. 10. (a)–(t) Stationary distributions obtained the first 20 principal components of the SST field as in the original GFDL-CM4 climate model and emulated by the neural emulator with multiplicative [Eq. (11)] and additive [Eq. (D1)] noise. Units: nondimensional. (u) Same as panels (a)–(t) but for the global-mean radiative flux at the top of the atmosphere. Units: W m^{-2} . The stationary distributions have been obtained after simulating the neural emulators for 600 yr. Note that the system was trained on only 7200 time steps, which limits the accuracy of the learned multiplicative noise term.

APPENDIX D: EMULATOR OF THE SST FIELD AND NET RADIATIVE FLUX: ADDITIVE VERSUS MULTIPLICATIVE NOISE

In the main text, we considered the model defined in Eq. (11) to emulate the joint variability of SST and net radiative flux at the TOA. Here, we examine a simpler variant with additive noise:

$$dx = (Lx + n(x))dt + \Sigma dW. \quad (\text{D1})$$

The drift term $Lx + n(x)$ is identical to the one used in the multiplicative noise model in the main text. We performed a 600-yr-long emulation following the same protocol as in the main text.

The resulting autocorrelation functions, the reconstruction of net radiative flux, and the ensemble-mean response are all nearly indistinguishable from those obtained with multiplicative noise [Eq. (11)], as shown in Figs. 6(c)–6(f) and 7.

This demonstrates that the deterministic component (shared by both emulators) is responsible for setting these features.

The main difference arises in the reconstruction of probability distributions. Figure 10 shows the PDFs for all 20 SST modes, as well as for the global-mean net radiative flux at the TOA. The PDF of the leading mode, associated with ENSO variability, is accurately reproduced by the emulator with multiplicative noise. Table I reports the first four moments of the

TABLE I. Mean, standard deviation, skewness, and kurtosis of SST Mode 1 (ENSO variability) for 600-yr-long runs.

Mode 1 (ENSO)	Mean	St. dev.	Skewness	Kurtosis
CM4 climate model	0.	28.27	−0.10	3.16
Eq. (11) (multiplicative noise)	2.83	29.07	−0.08	2.88
Eq. (D1) (additive noise)	−0.65	42.31	−0.13	3.33

distributions of this first mode for the climate model and for the two emulators. The most pronounced discrepancy appears in the standard deviation, where the multiplicative noise emulator shows substantially closer agreement. The multiplicative formulation also succeeds in capturing the PDF of the second mode.

Interestingly, for higher modes, each explaining a much smaller fraction of the variance, the emulator with additive

noise performs slightly better. These differences, while visible, are modest if comparing the scale with the values of the first mode. We suspect that the very limited dataset (only 7200 samples in total for training and validation) is the main factor preventing the network from learning the multiplicative noise structure with higher accuracy. With a larger dataset, we expect the performance of the multiplicative noise emulator to achieve more accurate reconstructions.

- [1] M. Ghil and V. Lucarini, The physics of climate variability and climate change, *Rev. Mod. Phys.* **92**, 035002 (2020).
- [2] V. Lucarini and M. Chekroun, Theoretical tools for understanding the climate crisis from Hasselmann's programme and beyond, *Nat. Rev. Phys.* **5**, 744 (2023).
- [3] C. Penland, Random forcing and forecasting using principal oscillation pattern analysis, *Mon. Weather Rev.* **117**, 2165 (1989).
- [4] C. Penland and P. Sardeshmukh, The optimal growth of tropical sea surface temperature anomalies, *J. Clim.* **8**, 1999 (1995).
- [5] C. Penland, A stochastic model of IndoPacific sea surface temperature anomalies, *Phys. D* **98**, 534 (1996).
- [6] S. Kravtsov, D. Kondrashov, and M. Ghil, Multilevel regression modeling of nonlinear processes: Derivation and applications to climatic variability, *J. Clim.* **18**, 4404 (2005).
- [7] D. Kondrashov, M. Chekroun, and M. Ghil, Data-driven non-Markovian closure models, *Phys. D* **297**, 33 (2015).
- [8] K. Strounine, S. Kravtsov, D. Kondrashov, and M. Ghil, Reduced models of atmospheric low-frequency variability: Parameter estimation and comparative performance, *Phys. D* **239**, 145 (2010).
- [9] N. Agarwal, D. Kondrashov, P. Dueben, E. Ryzhov, and P. Berloff, A comparison of data-driven approaches to build low-dimensional ocean models, *J. Adv. Model. Earth Syst.* **13**, e2021MS002537 (2021).
- [10] A. J. Majda and Y. Yuan, Fundamental limitations of ad hoc linear and quadratic multi-level regression models for physical systems, *Discrete Contin. Dyn. Syst.—Ser. B* **17**, 1333 (2012).
- [11] A. Majda and J. Harlim, Physics constrained nonlinear regression models for time series, *Nonlinearity* **26**, 201 (2013).
- [12] A. Souza, Representing turbulent statistics with partitions of state space. Part 1. Theory and methodology, *J. Fluid Mech.* **997**, A1 (2024).
- [13] A. Souza, Representing turbulent statistics with partitions of state space. Part 2. The compressible Euler equations, *J. Fluid Mech.* **997**, A2 (2024).
- [14] L. T. Giorgini, A. N. Souza, and P. J. Schmid, Reduced Markovian models of dynamical systems, *Phys. D* **470**, 134393 (2024).
- [15] L. Giorgini, Data-driven decomposition of conservative and non-conservative dynamics in multiscale systems, *arXiv:2505.01895*.
- [16] L. Giorgini, T. Bischoff, and A. Souza, Reduced-order modeling of cyclo-stationary time series using score-based generative methods, *arXiv:2508.19448*.
- [17] H. Fan, Y. Xiao, Y. Qu, F. Ling, B. Fei, L. Bai, and P. Gentile, Incorporating multivariate consistency in ML-based weather forecasting with latent-space constraints, *arXiv:2510.04006*.
- [18] A. Majda and D. Qi, Strategies for reduced-order models for predicting the statistical responses and uncertainty quantification in complex turbulent dynamical systems, *SIAM Rev.* **60**, 491 (2018).
- [19] D. Qi and A. Majda, Linear and nonlinear statistical response theories with prototype applications to sensitivity analysis and statistical control of complex turbulent dynamical systems, *Chaos* **29**, 103131 (2019).
- [20] M. Baldovin, F. Cecconi, and A. Vulpiani, Understanding causation via correlations and linear response theory, *Phys. Rev. Res.* **2**, 043436 (2020).
- [21] F. Falasca, P. Perezhgin, and L. Zanna, Data-driven dimensionality reduction and causal inference for spatiotemporal climate fields, *Phys. Rev. E* **109**, 044202 (2024).
- [22] J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath, and A. Anandkumar, FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators, *arXiv:2202.11214*.
- [23] I. Price, A. Sanchez-Gonzalez, F. Alet, T. R. Andersson, A. El-Kadi, D. Masters, T. Ewalds, J. Stott, S. Mohamed, P. Battaglia, R. Lam, and M. Willson, Probabilistic weather forecasting with machine learning, *Nature (London)* **637**, 84 (2025).
- [24] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Meroze, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, and P. Battaglia, Learning skillful medium-range global weather forecasting, *Science* **382**, 1416 (2023).
- [25] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian, Accurate medium-range global weather forecasting with 3D neural networks, *Nature (London)* **619**, 533 (2023).
- [26] B. Bonev *et al.*, FourCastNet 3: A geometric approach to probabilistic machine-learning weather forecasting at scale, *arXiv:2507.12144*.
- [27] D. Kochkov, J. Yuval, I. Langmore, P. Norgaard, J. Smith, G. Mooers, M. Klöwer, J. Lottes, S. Rasp, P. Düben, S. Hatfield, P. Battaglia, A. Sanchez-Gonzalez, M. Willson, M. P. Brenner, and S. Hoyer, Neural general circulation models for weather and climate, *Nature (London)* **632**, 1060 (2024).
- [28] O. Watt-Meyer, G. Dresdner, J. McGibbon, S. K. Clark, B. Henn, J. Duncan, N. D. Brenowitz, K. Kashinath, M. S. Pritchard, B. Bonev, M. E. Peters, and C. S. Bretherton, ACE: A fast, skillful learned global atmospheric model for climate prediction, *arXiv:2310.02074*.

- [29] O. Watt-Meyer, B. Henn, J. McGibbon, S. K. Clark, A. Kwa, W. A. Perkins, and C. S. Bretherton, ACE2: Accurately learning subseasonal to decadal atmospheric variability and forced responses, [arXiv:2411.11268](#).
- [30] H. Guan, T. Arcomano, A. Chattopadhyay, and R. Maulik, LUCIE: A lightweight uncoupled climate emulator with long-term stability and physical consistency for O(1000)-member ensembles, [arXiv:2405.16297](#).
- [31] A. Chattopadhyay, M. Gray, T. Wu, A. B. Lowe, and R. He, OceanNet: A principled neural operator-based digital twin for regional oceans, *Sci. Rep.* **14**, 21181 (2024).
- [32] S. Bire, B. Lütjens, K. Azizzadenesheli, A. Anandkumar, and C. Hill, Ocean emulation with Fourier neural operators: Double gyre, *J. Adv. Model. Earth Syst.* **17**, e2023MS004137 (2025).
- [33] S. Dheeshjith, A. Subel, A. Adcroft, J. Busecke, C. Fernandez-Granda, S. Gupta, and L. Zanna, Samudra: An AI global ocean emulator for climate, *Geophys. Res. Lett.* **52**, e2024GL114318 (2025).
- [34] W. Chapman, J. Schreck, Y. Sha, D. Gagne II, D. Kimpara, L. Zanna, K. Mayer, and J. Berner, CAMulator: Fast emulation of the community atmosphere model, [arXiv:2504.06007](#).
- [35] H. Guan, T. Arcomano, A. Chattopadhyay, and R. Maulik, LUCIE-3D: A three-dimensional climate emulator for forced responses, [arXiv:2509.02061](#).
- [36] J. Duncan *et al.*, SamudrACE: Fast and accurate coupled climate modeling with 3D ocean and atmosphere emulators, [arXiv:2509.12490](#).
- [37] L. Zanna *et al.*, A framework for hybrid physics-AI coupled ocean models, [arXiv:2510.22676](#).
- [38] A. Lancelin *et al.*, AI-boosted rare event sampling to characterize extreme weather, [arXiv:2510.27066](#).
- [39] C. Pedersen, L. Zanna, and J. Bruna, Thermalizer: Stable autoregressive neural emulation of spatiotemporal chaos, [arXiv:2503.18731](#).
- [40] R. Jiang, X. Zhang, K. Jakhar, P. Lu, P. Hassanzadeh, M. Maire, and R. Willett, Hierarchical implicit neural emulators, [arXiv:2506.04528](#).
- [41] Y. Sun, P. Hassanzadeh, M. Zand, and D. Abbot, Can AI weather models predict out-of-distribution gray swan tropical cyclones? *Proc. Natl. Acad. Sci. USA* **122**, e2420914122 (2025).
- [42] I. Baxter, H. Pahlavan, P. Hassanzadeh, K. Rucker, and T. Shaw, Benchmarking atmospheric circulation variability in an AI emulator, ACE2, and a hybrid model, neuralgcm, [arXiv:2510.04466](#).
- [43] S. Van Loon, M. Rugenstein, and E. A. Barnes, Reanalysis-based global radiative response to sea surface temperature patterns: Evaluating the ai2 Climate emulator, *Geophys. Res. Lett.* **52**, e2025GL115432 (2025).
- [44] J. Bloch-Johnson, M. A. A. Rugenstein, M. J. Alessi, C. Proistosescu, M. Zhao, B. Zhang, A. I. L. Williams, J. M. Gregory, J. Cole, Y. Dong, M. L. Duffy, S. M. Kang, and C. Zhou, The Green's function model intercomparison project (gf mip) protocol, *J. Adv. Model. Earth Syst.* **16**, e2023MS003700 (2024).
- [45] Y. Dong, C. Proistosescu, K. C. Armour, and D. S. Battisti, Attributing historical and future evolution of radiative feedbacks to regional warming patterns using a Green's function approach: The preeminence of the western Pacific, *J. Clim.* **32**, 5471 (2019).
- [46] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara *et al.*, The ERA5 global reanalysis, *Q. J. R. Meteorol. Soc.* **146**, 1999 (2020).
- [47] B. Zhang and T. Merlis, The equilibrium response of atmospheric machine-learning models to uniform sea surface temperature warming, [arXiv:2510.02415](#).
- [48] E. Wu, F. Rebassoo, P. Paul, C. Proistosescu, J. Nugent, D. McCoy *et al.*, Applying the ACE2 emulator to SST Green's functions for the E3SMv3 global atmosphere model, *J. Geophys. Res.: Mach. Learn. Comput.* **2**, e2025JH000774 (2025).
- [49] F. Cecconi, M. Cencini, M. Falcioni, and A. Vulpiani, Predicting the future from the past: An old problem from a modern perspective, *Am. J. Phys.* **80**, 1001 (2012).
- [50] H. Hosni and A. Vulpiani, Forecasting in light of big data, *Philos. Technol.* **31**, 557 (2018).
- [51] M. Baldovin, F. Cecconi, M. Cencini, A. Puglisi, and A. Vulpiani, The role of data in model building and prediction: A survey through examples, *Phys. Rev. Res.* **20**, 807 (2018).
- [52] N. Chen and A. J. Majda, Conditional Gaussian systems for multiscale nonlinear stochastic systems: Prediction, state estimation and uncertainty quantification, *Entropy* **20**, 509 (2018).
- [53] A. Majda, C. Franzke, and D. Croomlin, Normal forms for reduced stochastic climate models, *Proc. Natl. Acad. Sci. USA* **106**, 3649 (2010).
- [54] A. Majda, B. Gershgorin, and Y. Yuan, Low-frequency climate response and fluctuation-dissipation theorems: Theory and practice, *J. Atmos. Sci.* **67**, 1186 (2010).
- [55] A. J. Majda, R. V. Abramov, and M. J. Grote, *Information Theory and Stochastics for Multiscale Nonlinear Systems*, CRM Monograph Series (American Mathematical Society, Providence, RI, 2005).
- [56] U. Marconi, A. Puglisi, L. Rondoni, and A. Vulpiani, Fluctuation-dissipation: Response theory in statistical physics, *Phys. Rep.* **461**, 111 (2008).
- [57] G. L. and A. Vulpiani, Prediction and inference: From models and data to artificial intelligence, *Found. Phys.* **54**, 67 (2024).
- [58] A. Dalmedico, History and epistemology of models: Meteorology (1946–1963) as a case study, *Arch. Hist. Exact Sci.* **55**, 395 (2001).
- [59] L. A. Smith, What might we learn from climate forecasts? *Proc. Natl. Acad. Sci. USA* **99**, 2487 (2002).
- [60] J. G. Charney, On a physical basis for numerical prediction of large-scale motions in the atmosphere, *J. Meteorol.* **6**, 372 (1949).
- [61] J. G. Charney, R. Fjörtoft, and J. Von Neumann, Numerical integration of the barotropic vorticity equation, *Tellus* **2**, 237 (1950).
- [62] A. Timmermann *et al.*, El Niño–Southern Oscillation complexity, *Nature (London)* **559**, 535 (2018).
- [63] B. Dubrulle, F. Daviaud, D. Faranda, L. Marié, and B. Saint-Michel, How many modes are needed to predict climate bifurcations? Lessons from an experiment, *Nonlin. Processes Geophys.* **29**, 17 (2022).

- [64] J. Wouters and V. Lucarini, Cluster-based reduced-order modelling of a mixing layer, *J. Stat. Mech.* (2012) P03003.
- [65] G. Vissio and V. Lucarini, A proof of concept for scale-adaptive parametrizations: The case of the Lorenz '96 model, *Q. J. R. Meteorol. Soc.* **144**, 63 (2018).
- [66] I. Held, The gap between simulation and understanding in climate modeling, *Bull. Am. Meteorol. Soc.* **86**, 1609 (2005).
- [67] K. Hasselmann, Stochastic climate models part I. Theory, *Tellus* **28**, 473 (1976).
- [68] N. Chen and H. Liu, Minimum reduced-order models via causal inference, *Nonlinear Dyn.* **113**, 11327 (2025).
- [69] A. J. Majda, I. Timofeyev, and E. Vanden-Eijnden, Models for stochastic climate prediction, *Proc. Natl. Acad. Sci. USA* **96**, 14687 (1999).
- [70] A. J. Majda, I. Timofeyev, and Vanden-Eijnden, A mathematical framework for stochastic climate models, *Communications on Pure and Applied Mathematics* **54**, 891 (2001).
- [71] A. J. Majda, I. Timofeyev, and Vanden-Eijnden, A priori test of a stochastic mode reduction strategy, *Phys. D* **170**, 206 (2002).
- [72] A. J. Majda, I. Timofeyev, and Vanden-Eijnden, Systematic strategies for stochastic mode reduction in climate, *J. Atmos. Sci.* **60**, 1705 (2003).
- [73] T. J. O'Kane, D. Harries, and M. A. Collier, Bayesian structure learning for climate model evaluation, *J. Adv. Model. Earth Syst.* **16**, e2023MS004034 (2024).
- [74] J. Donges, Y. Zou, N. Marwan, and J. Kurths, Complex networks in climate dynamics, *Eur. Phys. J. Spec. Top.* **174**, 157 (2009).
- [75] I. Fountalis, C. Dovrolis, A. Bracco, B. Dilkina, and S. Keilholz, δ -MAPS from spatio-temporal data to a weighted and lagged network between functional domain, *Appl. Netw. Sci.* **3**, 21 (2018).
- [76] A. Majda, C. Franzke, and K. Boualem, An applied mathematics perspective on stochastic modelling for climate, *Phil. Trans. R. Soc. A* **366**, 2427 (2008).
- [77] D. Crommelin and A. Majda, Strategies for model reduction: Comparing different optimal bases, *J. Atmos. Sci.* **61**, 2206 (2004).
- [78] H. Fan, B. Fei, P. Gentile, Y. Xiao, K. Chen, Y. Liu, Y. Qu, F. Ling, and L. Bai, Physically consistent global atmospheric data assimilation with machine learning in a latent space, [arXiv:2502.02884](https://arxiv.org/abs/2502.02884).
- [79] N. Lutsko, I. Held, and P. Zurita-Gotor, Applying the fluctuation–dissipation theorem to a two-layer model of quasi-geostrophic turbulence, *J. Atmos. Sci.* **72**, 3161 (2015).
- [80] P. Hassanzadeh and Z. Kuang, The linear response function of an idealized atmosphere. Part II: Implications for the practical use of the fluctuation–dissipation theorem and the role of operator's nonnormality, *J. Atmos. Sci.* **73**, 3441 (2016).
- [81] C. Penland, M. D. Fowler, D. L. Jackson, and R. Cifelli, Forecasts of opportunity for northern California soil moisture, *Land* **10**, 713 (2021).
- [82] Y. Zhao and A. Capotondi, The role of the tropical atlantic in tropical Pacific climate variability, *npj Clim. Atmos. Sci.* **7**, 140 (2024).
- [83] J. Lien, Y.-N. Kuo, H. Ando, and S. Kido, Colored linear inverse model: A data-driven method for studying dynamical systems with temporally correlated stochasticity, *Phys. Rev. Res.* **7**, 023042 (2025).
- [84] J. Lien and H. Ando, Nonlinear extensions of linear inverse models under memoryless or persistent random forcing, *Phys. Rev. Res.* **7**, L032038 (2025).
- [85] D. Kondrashov, S. Kratsov, A. Robertson, and M. Ghil, A hierarchy of data-based ENSO models, *J. Clim.* **18**, 4425 (2005).
- [86] H. Risken, *The Fokker-Planck Equation—Methods of Solution and Applications* (Springer, New York, 1996).
- [87] J. Ismael, Reflections on the asymmetry of causation, *Interface Focus* **13**, 20220081 (2023).
- [88] E. Aurell and G. Del Ferraro, Causal analysis, correlation-response, and dynamic cavity, *J. Phys.: Conf. Ser.* **699**, 012002 (2016).
- [89] V. Lucarini, Revising and extending the linear response theory for statistical mechanical systems: Evaluating observables as predictors and predictands, *J. Stat. Phys.* **173**, 1698 (2018).
- [90] V. Lucarini and M.D. Chekroun, Detecting and attributing change in climate and complex systems: Foundations, Green's functions, and nonlinear fingerprints, *Phys. Rev. Lett.* **133**, 244201 (2024).
- [91] A. Majda, R. Abramov, and B. Gershgorin, High skill in low-frequency climate response through fluctuation dissipation theorems despite structural instability, *Proc. Natl. Acad. Sci. USA* **107**, 581 (2010).
- [92] C. Franzke and A. Majda, Low-order stochastic mode reduction for a prototype atmospheric GCM, *J. Atmos. Sci.* **63**, 457 (2006).
- [93] J. Berner and G. Branstator, Linear and nonlinear signatures in the planetary wave dynamics of an AGCM: Probability density functions, *J. Atmos. Sci.* **64**, 117 (2007).
- [94] P. Sura, M. Newman, C. Penland, and P. Sardeshmukh, Multiplicative noise and non-Gaussianity: A paradigm for atmospheric regimes? *J. Clim.* **62**, 1391 (2005).
- [95] G. Lacorata and A. Vulpiani, Fluctuation-response relation and modeling in systems with fast and slow dynamics, *Nonlin. Processes Geophys.* **14**, 681 (2007).
- [96] M. Baldovin, A. Puglisi, and A. Vulpiani, Langevin equations from experimental data: The case of rotational diffusion in granular media, *PLoS ONE* **14**, e0212135 (2019).
- [97] H. Arnold, I. Moroz, and T. Palmer, Stochastic parametrizations and model uncertainty in the Lorenz '96 system, *Phil. Trans. R. Soc. A* **371**, 20110479 (2013).
- [98] A. Dawson and T. Palmer, Simulating weather regimes: Impact of model resolution and stochastic parameterization, *Clim Dyn* **44**, 2177 (2015).
- [99] J. Berner *et al.*, Stochastic parameterization: Toward a new view of weather and climate models, *Bull. Am. Meteorol. Soc.* **98**, 565 (2017).
- [100] F. Cecconi, G. Costantini, C. Guardiani, M. Baldovin, and A. Vulpiani, Correlation, response and entropy approaches to allosteric behaviors: A critical comparison on the ubiquitous case, *Phys. Biol.* **20**, 056002 (2020).
- [101] B. Meyssignac, R. Guillaume-Castel, and R. Roca, Revisiting the global energy budget dynamics with a multivariate earth energy balance model to account for the warming pattern effect, *J. Clim.* **36**, 8113 (2023).
- [102] J. J. Barsugli and P. D. Sardeshmukh, Global atmospheric sensitivity to tropical SST anomalies throughout the Indo-Pacific basin, *J. Clim.* **15**, 3427 (2002).
- [103] C. Zhou, M. D. Zelinka, and S. A. Klein, Analyzing the dependence of global cloud feedback on the spatial pattern of sea

- surface temperature change with a Green's function approach, *J. Adv. Model. Earth Syst.* **9**, 2174 (2017).
- [104] B. Zhang, M. Zhao, and Z. Tan, Using a Green's function approach to diagnose the pattern effect in GFDL AM4 and CM4, *J. Clim.* **36**, 1105 (2023).
- [105] Y. Dong, K. C. Armour, M. D. Zelinka, C. Proistosescu, D. S. Battisti, C. Zhou, and T. Andrews, Intermodel spread in the pattern effect and its contribution to climate sensitivity in CMIP5 and CMIP6 Models, *J. Clim.* **33**, 7755 (2020).
- [106] M. J. Alessi and M. A. Rugenstein, Surface temperature pattern scenarios suggest higher warming rates than current projections, *Geophys. Res. Lett.* **50**, e2023GL105795 (2023).
- [107] A. I. Williams, N. Jeevanjee, and J. Bloch-Johnson, Circus tents, convective thresholds, and the non-linear climate response to tropical ssts, *Geophys. Res. Lett.* **50**, e2022GL101499 (2023).
- [108] F. Falasca, A. Basinski, L. Zanna, and M. Zhao, A fluctuation-dissipation theorem perspective on radiative responses to temperature perturbations, *J. Clim.* **38**, 3239 (2025).
- [109] I. M. Held, H. Guo, A. Adcroft, J. P. Dunne, L. W. Horowitz, J. Krasting, E. Shevliakova, M. Winton, M. Zhao, M. Bushuk, A. T. Wittenberg, B. Wyman, B. Xiang, R. Zhang, W. Anderson, V. Balaji, L. Donner, K. Dunne, J. Durachta, P. P. G. Gauthier *et al.*, Structure and performance of GFDL's CM4.0 climate model, *J. Adv. Model. Earth Syst.* **11**, 3691 (2019).
- [110] A. Adcroft, W. Anderson, V. Balaji, C. Blanton, M. Bushuk, C. O. Dufour *et al.*, The GFDL global ocean and sea ice model OM4.0: Model description and simulation features, *J. Adv. Model. Earth Syst.* **11**, 3167 (2019).
- [111] M. Zhao *et al.*, The GFDL global atmosphere and land model AM4.0/LM4.0: 1. Simulation characteristics with prescribed ssts, *J. Adv. Model. Earth Syst.* **10**, 691 (2018).
- [112] M. Zhao *et al.*, The GFDL global atmosphere and land model am4.0/LM4.0: 2. Model description, sensitivity studies, and tuning strategies, *J. Adv. Model. Earth Syst.* **10**, 735 (2018).
- [113] R. T. Pierrehumbert, *Principles of Planetary Climate* (Cambridge University Press, Cambridge, 2010).
- [114] D. M. Romps, J. T. Seeley, and J. P. Edman, Why the forcing from carbon dioxide scales as the logarithm of its concentration, *J. Clim.* **35**, 4027 (2022).
- [115] M. Zhao, An investigation of the effective climate sensitivity in gfdl's new climate models CM4.0 and SPEAR, *J. Clim.* **35**, 5637 (2022).
- [116] S. M. Kang, P. Ceppi, Y. Yu, and I.-S. Kang, Recent global climate feedback controlled by southern ocean cooling, *Nat. Geosci.* **16**, 775 (2023).
- [117] S. Van Loon, M. Rugenstein, and E. A. Barnes, Observation-based estimate of Earth's effective radiative forcing, *Proc. Natl. Acad. Sci. USA* **122**, e2425445122 (2025).
- [118] L. Fredericks, M. Rugenstein, D. W. J. Thompson, S. Van Loon, F. Falasca, R. Basinski-Ferris, P. Ceppi, Q. Wu, J. Bloch-Johnson, M. Alessi, S. M. Kang, Quantifying the radiative response to surface temperature variability: A critical comparison of current methods, [arXiv:2511.00731](https://arxiv.org/abs/2511.00731).
- [119] E. Guilyardi, P. Braconnot, F.-F. Jin, S. T. Kim, M. Kolasinski, T. Lee, and I. Musat, Atmosphere feedbacks during ENSO in a coupled GCM with a modified atmospheric convection scheme, *J. Clim.* **22**, 5698 (2009).
- [120] L. A. Smith, Disentangling uncertainty and error: On the predictability of nonlinear systems, in *Nonlinear Dynamics and Statistics*, edited by A. I. Mees (Birkhäuser Boston, Boston, MA, 2001), pp. 31–64.
- [121] P. Sardeshmukh, C. Penland, and G. Compo, Learning ENSO dynamics from data, *J. Clim.* **38**, 7437 (2025).
- [122] S. Lang *et al.*, AIFS–ecmwf's data-driven forecasting system, [arXiv:2406.01465](https://arxiv.org/abs/2406.01465).
- [123] B. Lütjens, R. Ferrari, D. Watson-Parris, and N. E. Selin, The impact of internal variability on benchmarking deep learning climate emulators, *J. Adv. Model. Earth Syst.* **17**, e2024MS004619 (2025).
- [124] G. Geogdzhayev, A. N. Souza, G. R. Flierl, and R. Ferrari, An eof-based emulator of means and covariances of monthly climate fields, EGUsphere [preprint], <https://egusphere.copernicus.org/preprints/2025/egusphere-2025-3768/>.
- [125] S. Bouabid, A. Souza, and R. Ferrari, Score-based generative emulation of impact-relevant earth system model outputs, [arXiv:2510.04358](https://arxiv.org/abs/2510.04358).
- [126] T. Palmer, A nonlinear dynamical perspective on model error: A proposal for non-local stochastic-dynamic parametrization in weather and climate prediction models, *Proc. Natl. Acad. Sci. USA* **127**, 279 (2001).
- [127] F. Alet *et al.*, Skillful joint probabilistic weather forecasting from marginals, [arXiv:2506.10772](https://arxiv.org/abs/2506.10772).
- [128] Y. Chen, M. Goldstein, M. Hua, M. Albergo, N. Boffi, and E. Vanden-Eijnden, Probabilistic forecasting with stochastic interpolants and föllmer processes, [arXiv:2403.13724](https://arxiv.org/abs/2403.13724).
- [129] A. Zhou, A. Wikner, A. Lancelin, P. Hassanzadeh, and A. B. Farimani, Reframing generative models for physical systems using stochastic interpolants, [arXiv:2509.26282](https://arxiv.org/abs/2509.26282).
- [130] N. Mücke and B. Sanderse, Physics-aware generative models for turbulent fluid flows through energy-consistent stochastic interpolants, [arXiv:2504.05852](https://arxiv.org/abs/2504.05852).
- [131] E. Lorenz, Predictability: A problem partly solved, in *Seminar on Predictability, 4–8 September 1995* (ECMWF, 1995).
- [132] A. Gritsun and G. Branstator, Climate response using a three-dimensional operator based on the fluctuation–dissipation theorem, *J. Atmos. Sci.* **64**, 2558 (2007).
- [133] L. Giorgini, F. Falasca, and A. Souza, Predicting forced responses of probability distributions via the fluctuation-dissipation theorem and generative modeling, *Proc. Natl. Acad. Sci. USA* **122**, e2509578122 (2025).
- [134] L. Onsager and S. Machlup, Fluctuations and irreversible processes, *Phys. Rev.* **91**, 1505 (1953).
- [135] D. Lucente, A. Baldassarri, A. Puglisi, A. Vulpiani, and M. Viale, Inference of time irreversibility from incomplete information: Linear systems and its pitfalls, *Phys. Rev. Res.* **4**, 043103 (2022).
- [136] A general framework for linking free and forced fluctuations via Koopmanism, *Chaos, Solitons Fractals* **202**, 117540 (2026).
- [137] M. J. Kusner, J. Loftus, C. Russell, and R. Silva, Counterfactual fairness, in *Advances in Neural Information Processing Systems*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., Long Beach, CA, USA, 2017), Vol. 30.
- [138] J. Runge, V. Petoukhov, J. F. Donges, J. Hlinka, N. Jajcay, M. Vejmelka, D. Hartman, N. Marwan, M. Paluš, and J. Kurths, Identifying causal gateways and mediators in complex spatio-temporal systems, *Nat. Commun.* **6**, 8502 (2015).

- [139] N. Ay and D. Polani, Information flows in causal networks, *Adv. Complex Syst.* **11**, 17 (2008).
 - [140] M. Falcioni and A. Vulpiani, The relevance of chaos for the linear response theory, *Phys. A* **215**, 481 (1995).
 - [141] M. Falcioni, S. Isola, and A. Vulpiani, Correlation functions and relaxation properties in chaotic dynamics and statistical mechanics, *Phys. Lett. A* **144**, 341 (1990).
 - [142] C. E. Leith, Climate response and fluctuation dissipation, *J. Atmos. Sci.* **32**, 2222 (1975).
 - [143] V. Lucarini, Interpretable and equation-free response theory for complex systems, [arXiv:2502.07908](https://arxiv.org/abs/2502.07908).
 - [144] G. Di Antonio and G. V. Vinci, Nonlinear correlations underlie linear response and causality, *Phys. Rev. Res.* **7**, L042029 (2025).
 - [145] E. Lorenz, Deterministic nonperiodic flow, *J. Atmos. Sci.* **20**, 130 (1963).
 - [146] J. Charney and J. DeVore, Multiple flow equilibria in the atmosphere and blocking, *J. Atmos. Sci.* **36**, 1205 (1979).
 - [147] A. J. Roberts, Resolving the multitude of microscale interactions accurately models stochastic partial differential equations, *LMS J. Comput. Math.* **9**, 193 (2006).
 - [148] A. Roberts, Normal form transforms separate slow and fast modes in stochastic dynamical systems, *Phys. A* **387**, 12 (2008).
 - [149] S. Kravtsov, M. Ghil, and D. Kondrashov, Empirical model reduction and the modeling hierarchy in climate dynamics and the geosciences, in *Stochastic Physics and Climate Modeling*, edited by T. Palmer and P. Williams (Cambridge University Press, Cambridge, 2009), pp. 35–72.
 - [150] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
 - [151] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
 - [152] J. Adler and O. Öktem, Deep Bayesian inversion, *Data-driven Models in Inverse Problems*, edited by T. A. Bubba (De Gruyter, Berlin, Boston, 2025), pp. 359–412.
 - [153] F. Takens, Detecting strange attractors in turbulence, in *Dynamical Systems and Turbulence, Lecture Notes in Mathematics*, edited by D. Rand and L. Young (Springer, Berlin, 1981), Vol. 898, p. 21.
- Correction:* The first sentence of the second paragraph of Sec. II B entitled “Impulse response operator” contained an error in wording and has been fixed.
- Second Correction:* The first sentence of the second paragraph of Sec. II B has been further modified. The first sentence of the paragraph numbered (2) in Appendix A 1 has been modified.