

Choosing observables that capture critical slowing down before tipping points: A Fokker-Planck operator approach

Johannes Lohmann^{1,*} and Georg A. Gottwald²

¹*Physics of Ice, Climate and Earth, Niels Bohr Institute, University of Copenhagen, Denmark*

²*School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia*



(Received 13 June 2025; accepted 12 November 2025; published 1 December 2025)

Tipping points (TP) are abrupt transitions between metastable states in complex systems, most often described by a bifurcation or crisis of a multistable system induced by a slowly changing control parameter. An avenue for predicting TPs in real-world systems is critical slowing down (CSD), which is a decrease in the relaxation rate after perturbations prior to a TP that can be measured by statistical early warning signals (EWS) in the autocovariance of observational time series. In high-dimensional systems, we cannot expect *a priori* chosen scalar observables to show significant EWS, and some may even show an opposite signal. Thus, to avoid false negative or positive early warnings, it is desirable to monitor fluctuations only in observables that are designed to capture CSD. Here we propose that a natural observable for this purpose can be obtained by a data-driven approximation of the first nontrivial eigenfunction of the backward Fokker-Planck (or Kolmogorov) operator, using the diffusion map algorithm.

DOI: [10.1103/PhysRevE.112.064204](https://doi.org/10.1103/PhysRevE.112.064204)

I. INTRODUCTION

We consider how critical transitions in stochastically forced complex systems may be anticipated by measuring increases in amplitude and temporal correlation of fluctuations in certain observables as early warning signals (EWS). The complex systems of interest are a general class of heterogeneous systems with many interacting agents or scales, arising for instance in ecology, biology, social science, and the Earth system [1]. They are often modeled by a first-order stochastic differential equation with a nonlinear, deterministic *drift* giving rise to possibly chaotic dynamics, and a noise process that represents unresolved scales and random disturbances by the environment, as well as control parameters that modify the drift and represent slow changes in external boundary conditions. A critical transition occurs when upon a parameter change a base state, i.e., a stable invariant set of the drift, loses stability and the system undergoes an abrupt transition to an alternative state. This is usually due to a collision of the base state with an *edge* state, which is an unstable invariant set of the drift. The stable manifold of the edge state is the basin boundary separating the base state from the alternative state. The simplest case of such a transition is a noisy saddle-node bifurcation (SNB) [2], which is often considered the archetype of a tipping point (TP).

EWS arise due to so-called critical slowing down (CSD) [3–9], which is most easily understood for the SNB. Here, as a control parameter μ crosses a critical value (e.g., $\mu_c = 0$), the leading eigenvalue of the Jacobian describing the linearized dynamics around the base fixed point crosses the imaginary axis. Leading up to this, for $0 < \mu \ll 1$, the dynamics along one degree of freedom (d.o.f.) becomes much slower compared to all others, and after a short relaxation time the system is confined to an (extended) center manifold, or a neighborhood thereof due to the noise forcing [10]. The drift on the manifold is one-dimensional and given by the SNB normal form after a suitable coordinate transformation. As $\mu \rightarrow 0$, CSD refers to the slowing down of the relaxation dynamics toward the equilibrium along the center manifold after an arbitrary perturbation.

In more general cases, where the base state is a limit cycle or chaotic attractor, we also expect a decrease of the relaxation rate back to steady state after a perturbation of the system. This is plausible since upon control parameter change the underlying deterministic dynamics experience a continuous change from being stable to neutrally stable in one d.o.f., before finally becoming unstable. This generic feature makes the detection of loss of resilience to perturbations the primary avenue for predicting TPs [11].

For large-scale systems in the real world, controlled perturbations are difficult to find. But there is a permanent influence of random disturbances from the environment. Such noise-driven, natural fluctuations of the unperturbed system allow one to infer the system's response to perturbations if linear response theory guarantees a fluctuation-dissipation theorem [12,13]. The size and correlation of the fluctuations are expected to grow in tandem with the system's slowing response as the critical transition is approached, thereby forming statistical EWS. This is the other side of the coin of CSD. Growing

*Contact author: johannes.lohmann@nbi.ku.dk

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

fluctuations toward the basin boundary imply a flattening of the quasipotential [14]. This happens in the direction of a particular d.o.f. that is related to the location of the edge state, since the latter usually lies on the most probable path of a noise-induced escape [15].

Real-world observations have been analyzed for CSD by measuring statistical EWS of presumed critical transitions, including financial crises [16], depression [17,18], neuron spiking [19], and climate tipping points, such as the Greenland ice sheet [20], Amazon forest [21], and Atlantic meridional overturning circulation (AMOC) [22–24]. But statistical false positives and false negatives can occur. The destabilization of the system in a single (critical) d.o.f. implies that the increase in noise-driven variability occurs also predominantly in a single d.o.f., which gives rise to a scalar observable where the variance is expected to diverge and the autocorrelation to approach 1 at the bifurcation. Thus, EWS can be masked if measurements have been taken from a dynamical observable that does not sufficiently project on the critical d.o.f. [2,15,25–27]. In the simple case of the SNB, this means that the observable does not follow the SNB normal form to any good approximation and hence is not subjected to significant CSD.

Consequently, the central question that will be addressed here is what observables should be used to detect CSD. This depends on how the system under the influence of noise responds to perturbations away from its steady state, and how this response changes as a control parameter approaches its critical value. This can be understood in terms of the Fokker-Planck (FP) equation, which governs the temporal evolution of the probability density of the state in phase space. The density can be written as an expansion in the eigenfunctions ψ_n of the FP operator $\mathcal{L}$. For a fixed control parameter any initial density will converge to the unique stationary density $\pi(\mathbf{x})$, which is the first eigenfunction ψ_0 with eigenvalue $\lambda_0 = 0$. The system is then in statistical equilibrium, where the contributions of all other ψ_n with $|\lambda_n| > 0$ have decayed. The first few ψ_n (with $|\lambda_n|$ closest to 0, and assuming a discrete spectrum) signify locations in phase space where fluctuations tend to linger on finite, but long-term time horizons.

We consider systems with a possible TP, i.e., the deterministic drift term in the governing equations gives rise to (at least) bistability, with a base and an alternative attractor. A low noise strength is assumed, meaning it is small compared to the (quasi)potential barrier height between competing attractors of the drift. It is assumed to hold for most values of μ , except when very close to the bifurcation where the barrier height goes to zero. This is required for the paradigm of bifurcation-induced tipping [28] to be meaningful and for the concept of EWS to be useful. Hence, when far from the bifurcation, the system spends very long periods of time in distinct regions around the attractors, referred to as metastable states. Transitions between the states can be ignored, as they occur very rarely on timescales of $\mathcal{O}(1/|\lambda_1|)$ with $|\lambda_1| \ll 1$. While part of the invariant density $\pi(\mathbf{x})$ occupies the alternative metastable state, we hence assume that on a finite time horizon the system is in a quasistationary distribution concentrated entirely around the base state, where the contribution of ψ_1 , which signifies a very slow transfer of density between the metastable states, has not decayed. Due to CSD, the relaxation

toward the base state within the quasipotential well of the base state becomes slower along a particular mode. When close enough to the TP, this mode becomes the slowest in the system and will be expressed by the next eigenfunction ψ_2 .

An observable that naturally expresses increases in fluctuations related to this critical mode is proposed here to be given by the corresponding eigenfunction ϕ_2 of the backward (adjoint) operator $\mathcal{L}^*$, also known as the generator. $\mathcal{L}^*$ governs the temporal evolution of expectation values of observables as a function of initial states (e.g., states after a perturbation), and the first few ϕ_n can be interpreted as patterns of initial conditions with slowest decay toward $\pi(\mathbf{x})$. This proposition is in agreement with the framework of optimal fingerprints presented in Ref. [29].

To obtain $\mathcal{L}^*$ from data we propose to use the diffusion map (DM) algorithm [30–32]. DM has been successfully used to define generalized collective coordinates that capture the effective dynamics of complex systems [32–34]. It gives an approximation (discretized on the set of data points) of $\mathcal{L}^*$ induced by a stochastic differential equation with drift $\nabla \ln[\pi(\mathbf{x})]$, i.e., a gradient system related to the quasipotential of the underlying stochastic dynamic system [14,35]. We show here that it yields physical observables carrying excellent EWS, also for nongradient systems including a high-dimensional global ocean model exhibiting a TP of the AMOC. It preserves the flattening of the quasipotential (in a particular critical d.o.f.) as a key property of CSD, which is not affected by nongradient terms of the drift. Nevertheless, methods to approximate $\mathcal{L}^*$ for nongradient systems can be applied, such as the extension of DM for general diffusions with multiplicative noise [36], which is also implemented here and shown to give very similar results (Appendix A).

The paper is structured as follows. Section II reviews some fundamentals of FP operators and introduces notation, as well the DM algorithm to estimate the eigenfunctions of $\mathcal{L}^*$ from data. In Secs. III A and III B we motivate the usage of the backward FP eigenfunctions for the purposes of EWS with simple double well potential systems in one and two dimensions. In Sec. III C we show with conceptual models that the reconstruction of the eigenfunctions with DM indeed yields observables that carry strong EWS. Further, in Sec. III D we show that such observables are strictly necessary if one wants extrapolate from increasing fluctuations to forecast the timing of a TP. In Sec. III E we apply our method successfully to a high-dimensional model of the global ocean circulation, and conclude with a discussion in Sec. IV.

II. FOKKER-PLANCK EIGENFUNCTIONS AND DIFFUSION MAPS

Consider the d -dimensional state $\mathbf{X} = (X^1, \dots, X^d)$, governed by an Ito diffusion equation with time-independent coefficients written component-wise as

$$dX^\nu = b^\nu(\mathbf{X})dt + \sqrt{\epsilon}\sigma_\nu^\gamma(\mathbf{X})dW^\gamma, \quad (1)$$

with drift $b^\nu(\mathbf{X})$ and diffusion $\sigma_\nu^\gamma(\mathbf{X})$. The transition probability density $P(\mathbf{X}(t) = \mathbf{x} | \mathbf{X}(0) = \mathbf{x}_0) \equiv P(\mathbf{x}, t | \mathbf{x}_0)$ is governed by the FP equation

$$\partial_t P(\mathbf{x}, t | \mathbf{x}_0) = \mathcal{L}P(\mathbf{x}, t | \mathbf{x}_0), \quad (2)$$

with FP operator

$$\mathcal{L}(x) = -\frac{\partial}{\partial x^\gamma} b^\gamma(\mathbf{x}) + \frac{\epsilon}{2} \frac{\partial^2}{\partial x^\gamma \partial x^\nu} A^{\gamma\nu}(\mathbf{x}) \quad (3)$$

and diffusion tensor $A^{\gamma\nu}(\mathbf{x}) = \sigma_\lambda^\gamma(\mathbf{x}) \sigma_\sigma^\nu(\mathbf{x}) \delta^{\lambda\sigma}$. The stationary distribution $\pi(\mathbf{x}) \equiv \lim_{t \rightarrow \infty} P(\mathbf{x}, t | \mathbf{x}_0)$ is an eigenfunction of $\mathcal{L}$ with eigenvalue $\lambda_0 = 0$, satisfying $\mathcal{L}\pi = 0$. The eigenfunctions of the FP operator

$$\mathcal{L}\psi_n(\mathbf{x}) = \lambda_n \psi_n(\mathbf{x}) \quad (4)$$

have eigenvalues with $|\lambda_0| = 0 < |\lambda_1| \leq |\lambda_2| \leq \dots < \infty$ and negative real parts. We restrict our attention to a discrete FP operator spectrum. For certain sufficiently chaotic, deterministic dynamics $\mathcal{L}$ is quasicompact in appropriate function spaces, which implies that the dominant eigenvalues are discrete and well separated from the essential spectrum [37]. Further, sufficiently smooth stochastic dynamical systems have compact operators provided the domain is compact [38], or in unbounded domains given certain conditions on the drift [39]. We only consider stochastic dynamical systems, since otherwise the concept of EWS is not meaningful. This can also be motivated physically in terms of the Hasselmann program [40], where stochastic dynamical systems are viewed as the limiting equations for macroscopic slow processes in high-dimensional deterministic systems. In particular, the interaction between chaotic multiscale microscopic constituents in deterministic high-dimensional systems leads to an emergent stochastic behavior at the slow macroscopic level (see, for example, Ref. [41]). A bounded domain may be physically justified too, since for practical purposes in reasonable time horizons the dynamics should be sufficiently contained to access only a bounded region.

Time-varying solutions of the FP equation can be written in the eigenfunction basis as

$$P(\mathbf{x}, t) = \sum_{n=0}^{\infty} c_n \psi_n(\mathbf{x}) e^{\lambda_n t}, \quad (5)$$

with $c_n = \int \psi_n(\mathbf{x}) \pi^{-1}(\mathbf{x}) \pi_0(\mathbf{x}) d\mathbf{x}$, where $\pi_0(\mathbf{x}) = P(\mathbf{x}, t = 0)$. Since $|\lambda_0| = 0 < |\lambda_1| \leq |\lambda_2| \leq \dots$ with nonpositive real parts, eigenfunctions with small indices decay slowest. In view of studying TPs, we consider here dynamical systems that spend a very long time in one part of phase space (a metastable set) before exhibiting a rare transition to another, and so on. In the simplest case there are two metastable sets, corresponding to neighborhoods of the attractors of the underlying deterministic dynamics, which is reflected in the spectrum of $\mathcal{L}$ as $|\lambda_1| \ll 1$ and a spectral gap with $|\lambda_2 - \lambda_1| \gg |\lambda_1|$. To study the dynamics before the TP, it is sufficient to consider this bistable case, since a critical transition generically consists of the collision of only a single base state and a single boundary.

The adjoint of the FP operator is the generator

$$\mathcal{L}^*(x) = a^\gamma(\mathbf{x}) \frac{\partial}{\partial x^\gamma} + \frac{\epsilon}{2} b^{\gamma\nu}(\mathbf{x}) \frac{\partial^2}{\partial x^\gamma \partial x^\nu}. \quad (6)$$

It governs the backward Kolmogorov equation

$$-\partial_s u(\mathbf{x}, s) = \mathcal{L}^* u(\mathbf{x}, s), \quad (7)$$

which is defined on the time interval $s \in [0, T]$ for functions $u(\mathbf{x}, s) = \mathbb{E}_{\mathbf{x}, s}[f(X_T)] \equiv \mathbb{E}[f(X_T) | X_s = \mathbf{x}]$, i.e., conditional expectation values of observables $f(\mathbf{x}, s)$ (initialized at $\mathbf{x}$), and with the final condition $u(\mathbf{x}, T) = f(\mathbf{x})$. Employing the transformation $t = T - s$ the equation can be formulated as an initial value problem

$$\partial_t u(\mathbf{x}, t) = \mathcal{L}^* u(\mathbf{x}, t), \quad (8)$$

with initial condition $u(\mathbf{x}, 0) = f(\mathbf{x})$. Here, u is the conditional expectation $u(\mathbf{x}, t) = \mathbb{E}[f(X_t) | X_0 = \mathbf{x}]$ with initial position $\mathbf{x}$. The solution of Eq. (8) can be expressed as an eigenfunction expansion

$$u(\mathbf{x}, t) = \sum_{n=0}^{\infty} d_n \phi_n(\mathbf{x}) e^{\lambda_n t}, \quad (9)$$

with $d_n = \int \phi_n(\mathbf{x}) f(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}$, and eigenfunctions ϕ_n satisfying $\mathcal{L}^* \phi_n = \lambda_n \phi_n$. The leading eigenfunction $\phi_0(\mathbf{x})$, corresponding to the eigenvalue $\lambda_0 = 0$, is the unique solution of $\mathcal{L}^* \phi_0 = 0$ and is given by $\phi_0 = \text{const}$. This reflects the ergodicity of the underlying system, which implies that (long-term) expectation values do not depend on the initial conditions.

Any observable $g(x)$ can be expressed by an expansion in the backward eigenfunction basis with

$$g(\mathbf{x}) = \sum_{n=0}^{\infty} g_n \phi_n(\mathbf{x}), \quad (10)$$

and possibly approximated by a truncation thereof. Here, $g_n = \int g(x) \phi_n(x) \pi(x) dx$. Thus, $\phi_n(\mathbf{x})$ can themselves be considered as observables, and in particular the first few $\phi_n(\mathbf{x})$ are special observables with expectation values that converge only slowly because they are nonconstant functions in regions of phase space where there is a slow relaxation to equilibrium. From a different point of view, the subset of leading $\phi_n(x)$ are transformations of the system from the original coordinates to reduction coordinates. The reduction is meaningful in case of timescale separation, which is expected to emerge when the deterministic drift of the system approaches a bifurcation. In particular, it can be shown that the evolution of the first k eigenfunctions is approximately Markovian [42]. In this case, the long-term evolution of the system is governed by the first k backward eigenfunctions.

For systems obeying detailed balance, i.e., when the drift $b(\mathbf{X})$ is the gradient of a potential and the diffusion $\sigma(\mathbf{X})$ is independent of the position $\mathbf{X}$ (additive noise), a discrete approximation to $\phi_n(x)$ can be obtained by the DM algorithm given sufficient data. For the majority of this work, we will use this method and show that it gives good results in the context of EWS, even for nongradient systems, where we thus effectively reconstruct a gradient system based on the quasipotential $V_q(\mathbf{x}) \propto \ln \pi(\mathbf{x})$ of the full system. The presence of nongradient terms does not change $\pi(\mathbf{x})$, but it gives rise to oscillatory modes that are filtered out here.

Importantly, a generalized DM method has been introduced that allows to obtain $\phi_n(x)$ also for nongradient drift and position-dependent (multiplicative) noise [36], if one can estimate drift and diffusion at the data points. We implemented this generalized method for the high-dimensional ocean model Veros [50], showing that it gives very sim-

ilar results compared to the simpler DM method (see Appendix A). While we do not consider examples with multiplicative noise, being able to properly account for it may be crucial in practice, as variance-based methods for EWS can fail [43,44]. The advantage of using the original DM method is that it requires tuning fewer parameters, and it will be more widely applicable for EWS, since drift and diffusion matrix do not need to be estimated, which can be challenging for smaller data sets in high dimensions.

The DM algorithm defines a weighted graph on the data points, and subsequently computes the eigenvalues and eigenvectors of a random walk on this graph. To this end we define a kernel with bandwidth $\epsilon > 0$ measuring the distance of two data points $\mathbf{x}$ and $\mathbf{y}$

$$K(\mathbf{x}, \mathbf{y}) = \exp(-(4\epsilon)^{-1} \|\mathbf{x} - \mathbf{y}\|^2). \quad (11)$$

With this, given N data points $\{\mathbf{x}_i\}_{i=1}^N$, construct the $N \times N$ matrix for all pairs of data points

$$\tilde{K}_{ij} = \frac{K(\mathbf{x}_i, \mathbf{x}_j)}{\sqrt{p_\epsilon(\mathbf{x}_i)p_\epsilon(\mathbf{x}_j)}}, \quad (12)$$

with $p_\epsilon(\mathbf{x}) = \sum_{j=1}^N K(\mathbf{x}, \mathbf{x}_j)$. Finally, construct the row-stochastic Markov matrix

$$M_{ij} = \frac{\tilde{K}_{ij}}{D_i}, \quad (13)$$

with $D_i = \sum_{j=1}^N \tilde{K}_{ij}$. The first few eigenvectors v_n of M , corresponding to eigenvalues $\lambda_n^{(M)}$, define the so-called diffusion coordinates $\xi_n = \lambda_n^{(M)} v_n$. In the limit $N \rightarrow \infty$ and $\epsilon \rightarrow 0$, the operator $(M - I)/\epsilon$ converges to $\mathcal{L}^*$ (i.e., the backward FP operator) [31,45,46], with $\lambda_n^{(M)} = e^{\lambda_n}$. The Euclidean distance between data points in the DM coordinates is called the diffusion distance and measures how closely two points are connected via diffusion of the Markov chain M . Two points $\mathbf{x}$ and $\mathbf{y}$ may have small Euclidean distance, but large diffusion distance, which can reflect that the dynamics evolve on a lower dimensional manifold.

To evaluate the eigenfunctions approximately at points $\mathbf{y}$ not in the given data set $\{\mathbf{x}_i\}_{i=1}^N$, the so-called Nyström eigenspace interpolation can be employed,

$$\xi_n(\mathbf{y}) = \lambda_n^{-1} \sum_{i=1}^N \frac{\tilde{K}(\mathbf{y}, \mathbf{x}_i)}{D(\mathbf{y})} \xi_n(\mathbf{x}_i), \quad (14)$$

with

$$\tilde{K}(\mathbf{y}, \mathbf{x}_i) = \frac{K(\mathbf{y}, \mathbf{x}_i)}{\sqrt{p_\epsilon(\mathbf{y})p_\epsilon(\mathbf{x}_i)}} \quad (15)$$

and $D(\mathbf{y}) = \sum_{i=1}^N \tilde{K}(\mathbf{y}, \mathbf{x}_i)$. In Eq. (14), $\xi_n(\mathbf{x}_i)$ denotes the entry of the n th eigenvector corresponding to the i -th data point.

In our DM implementation we normalize each data variable to have unit variance. Furthermore, we remove a small number of outliers so that we can use a smaller ϵ to obtain a better approximation of the backward eigenfunctions. Individual outliers are those data points where the minimum of the distance to all other points is largest. We find $n = 15$ points with the largest minimum distance and remove the respective columns and rows in the distance matrix $K(\mathbf{x}_i, \mathbf{x}_j)$ in Eq. (12).

Thereafter we remove $n = 10$ double outliers, which are those pairs of points where the second smallest distance to all other points is largest. The number of removed outliers has been chosen by trial and error to give the best performance across all data sets with typical sample size of about 10 000. If this step is skipped, often a quite large ϵ is needed to prevent the first diffusion coordinates from merely acting to cluster individual outliers against the rest of the data.

III. RESULTS

A. Interpretation of Fokker-Planck eigenfunctions in one dimension

We first study the FP eigenfunctions in a one-dimensional system with additive noise

$$dX_t = -\left(\frac{d}{dx}V(X_t)\right)dt + \sigma dW_t, \quad (16)$$

and double-well potential

$$V(x) = x^4 - x^2 + \beta x, \quad (17)$$

where β is a control parameter. The potential is shown in Fig. 1(e) for different values of β . The deterministic drift of the system undergoes a saddle-node bifurcation at $\beta_c = \sqrt{8/27} \approx 0.544331$, where one of the potential wells disappears. From here on we will simply refer to this as bifurcation or TP, also when referring to the stochastic system. We determine the forward eigenfunctions ψ_n numerically by a discrete approximation of $\mathcal{L}$ using the finite difference method by Chang and Cooper [47], and then an eigendecomposition of the obtained matrix using an implicitly restarted Arnoldi method (scipy.sparse.linalg.eigs package implementation of ARPACK). Since this is a gradient system the forward and backward eigenfunctions are related by $\phi_n(\mathbf{x}) = \pi(\mathbf{x})^{-1} \psi_n(\mathbf{x})$, with the same, real λ_n . Hence, we can determine both ψ_n and ϕ_n , which is shown in Fig. 1.

As β is increased from zero toward the bifurcation, for low noise $\pi(\mathbf{x})$ quickly becomes heavily asymmetric with a dominant peak at the deeper potential well [Fig. 1(a)], and a peak around the shallower well that is orders of magnitude smaller. The ψ_n for small $n > 0$ correspond to distinct patterns that modify the density such that it takes longest until statistical equilibrium is reached, given that the initial density $\pi_0(\mathbf{x})$ projects significantly on those pattern.

The slowest decaying pattern ψ_1 describes the situation where the probability mass in one of the wells is initially larger as it should be according to $\pi(\mathbf{x})$ [Fig. 1(b)]. For the equilibration of such a configuration, part of the probability mass needs to diffuse uphill and overcome the potential barrier. In systems with low noise this is associated with a long timescale, and $|\lambda_1|$ is approximately equal to the escape rate out of the shallow well.

The pattern ψ_2 gives a large contribution when the initial density is concentrated more prominently in the vicinity of the saddle [as compared to $\pi(\mathbf{x})$]. The function shows minima slightly outwards [larger $|x|$] of the two stable fixed points, and a broad double maximum around the saddle point [Fig. 1(c)]. This pattern can be interpreted as additional mass that survives outside the vicinity of the two minima for some time $|\lambda_2|^{-1}$ due to the asymmetry of each well, i.e., the smaller




FIG. 1. Eigenfunctions of the Fokker-Planck operator for the one-dimensional double-well model (16) for a range of parameter values from $\beta = 0.05$ to $\beta = 0.53$ drawn with increasing contrast. The bifurcation occurs at $\beta \approx 0.544$. The eigenfunctions are estimated by eigendecomposition of the discrete approximation of the FP operator via the scheme of Chang and Cooper [47]. (a)–(d) The first four eigenfunctions of the forward FP operator for $\sigma = 0.25$. (e) Associated potential $V(x)$. (f)–(h) The first three nontrivial eigenfunctions of the backward operator $\phi_{1,2,3}$ (rescaled for each parameter value to have a maximum value of 1). The green dots in panel (g) indicate the locations of the saddle point for the respective parameter values. The vertical dashed line is the inflection point of the potential in the shallower well, which is independent of β .

curvature of the potential toward the saddle. In other words, the relaxation toward equilibrium is slower in the vicinity of the saddle and on the sides of the wells that are facing the saddle. As the bifurcation is approached, the segment of the potential within the shallow well that faces the saddle becomes more flat, and thus ψ_2 is the relevant mode carrying CSD. In the one-dimensional case, higher eigenfunctions are less important for our analysis, representing higher-order corrections [see Fig. 1(d) for ψ_3].

With these considerations on ψ_n one may interpret the backward eigenfunctions ϕ_n . As mentioned above, ϕ_0 is constant due to the ergodicity of the system. ϕ_1 shows a sigmoidal shape, with plateaus around the two fixed points. On timescales smaller than the mean escape time from the shallow well, observables thus have different, quasiconstant expectation values that depend on which is the starting basin. The transition zone of the sigmoid function with its midpoint at the saddle becomes narrower for decreasing noise levels.

While the constant ϕ_0 can be called the “trivial” eigenfunction, and ϕ_1 the dominant eigenfunction since $|\lambda_2 - \lambda_1| \gg |\lambda_1|$, we refer to ϕ_2 as the first subdominant eigenfunction. ϕ_2 peaks close to the saddle point, and converges to 0 at the deep well while reaching lower values in the shallow well. On timescales of order $|\lambda_2|^{-1}$, expectation values are thus altered when starting close to the saddle. ϕ_3 (and similarly higher eigenfunctions) is nonmonotonic within the shallow well. It first increases toward the inflection point in the shallow well [Fig. 1(h)], and then decreases again toward the saddle. For

the purpose of EWS, we are, however, interested in observables that are monotonic from the base attractor toward the edge state, because otherwise fluctuations of increasing length along the critical d.o.f. toward the edge state due to CSD are suppressed in the measured observable.

Depending on the noise level, the CSD as β is changed toward the bifurcation may be reflected in the eigenvalues. Since ψ_2 captures the slowing of relaxation toward equilibrium as the curvature in the shallow well decreases, the relaxation rate $|\lambda_2|$ should go toward zero as a manifestation of CSD. This can indeed be seen for low noise levels in Figs. 2(b) and 2(c). But for finite noise levels $|\lambda_2|$ remains finite [Fig. 2(a)], and even slightly increases close to the bifurcation. It is bounded by the noise-induced escape rate $|\lambda_1|$, which increases drastically and becomes $O(1)$ at the bifurcation, at which point the potential is so flat that the relevant timescale for the decay of ψ_2 is not set by deterministic relaxation, but pure diffusion dynamics. Additionally, the decreasing distance of saddle and fixed point in the shallow well may play a role.

B. Eigenfunctions in two dimensions

In the previous one-dimensional example, there were no other slow d.o.f. that compete with the critical d.o.f. to be in the position of mode ϕ_2 . But generally the situation is different, especially in high-dimensional systems with multiscale behavior, where the correct physical mode first needs to slow down enough so that it emerges as ϕ_2 . We illustrate this with




FIG. 2. (a)–(c) Eigenvalues $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$ of the FP operator for the one-dimensional double-well potential as function of the control parameter β , for different noise levels σ . The critical value corresponding to the bifurcation is marked by the vertical dashed line. (d) Eigenvalues $\{\lambda_1, \lambda_2, \dots, \lambda_7\}$ of the FP operator of the two-dimensional double-well (19) as a function of the control parameter e , using the noise level $\sigma = 0.3$. The bifurcation point is marked with the vertical dashed line.

a system of two variables x and y in the double-well potential

$$V(x, y) = x^2(x^2 + y^2 - a) + y \frac{cy + d}{x^2 + b} + ex. \quad (18)$$

Adding Gaussian white noise independently to both variables yields the system

$$\begin{pmatrix} dx_t \\ dy_t \end{pmatrix} = \begin{pmatrix} -\frac{\partial V}{\partial x} \\ -\frac{\partial V}{\partial y} \end{pmatrix} dt + \begin{pmatrix} \sigma_x dW_{x,t} \\ \sigma_y dW_{y,t} \end{pmatrix}, \quad (19)$$

where $W_{x,t}$ and $W_{y,t}$ are independent, standard Wiener processes. $a = 2.5$, $b = 0.5$, $c = 0.2$, $d = 0.5$ are chosen, and e is the control parameter. For small e , there are two stable fixed points and one saddle point in the deterministic system. There is a saddle-node bifurcation of the deterministic drift at $e_c \approx 1.73$, where the potential well with $x > 0$ disappears. Figure 3(f) shows isolines of the potential, as well as the fixed points and basin boundary at $e = 0.5$.

The first nontrivial eigenfunction ψ_1 is again related to the slow transport of density from one well to the other [Fig. 3(b)], and accordingly the backward function ϕ_1 is approximately constant in the two basins [Fig. 3(g)]. Next, compared to the one-dimensional (1D) double well there is an additional eigenmode because of the slow timescale from the generally slower deterministic dynamics in the y direction. Far from the bifurcation, ψ_2 represents probability mass that is more slowly contracted in the y direction and for some time ($|\lambda_2|^{-1}$) has the tendency to linger at strongly negative y values [Fig. 3(c)] instead of converging to either fixed point. Hence, ϕ_2 identifies initial conditions that take longest to converge to either of the two wells along the y direction [Fig. 3(h)]. The next mode ψ_3 corresponds (at this value of the control parameter e) to ψ_2 of the one-dimensional case, i.e., a result of slow convergence to the fixed points in the more flat parts of the asymmetric potential wells toward the saddle point [Fig. 3(d)]. Correspondingly, ϕ_3 peaks near the saddle, and it shows that in particular initial conditions starting near the stable manifold of the saddle will lead to a slow relaxation of conditional expectation values. ψ_4 is a higher-order, antisymmetric pattern, analogous to ψ_3 of the 1D double well.

As e approaches the bifurcation, the eigenvalues of the abovementioned patterns cross [Fig. 2(d)], leading to a differ-

ent ordering of the modes [Figs. 3(k)–3(t)]. The eigenvalue of the pattern associated with the d.o.f. in the y direction decreases [gray band in Fig. 2(d)], and the pattern drops to higher n . Instead, the pattern related to the low potential curvature toward the saddle [yellow band in Fig. 2(d)] becomes ϕ_2 [Fig. 3(r)], i.e., the pattern subdominant only to the pattern ϕ_1 that reflects noise-induced escape. For the critical d.o.f. to reach the position of ϕ_2 before a transition, it is crucial that the noise is low enough. If this is not the case and there is a very slow (and unchanging) d.o.f., the moment when the critical d.o.f. gets slow enough, i.e., when $|\lambda_3|$ approaches $|\lambda_2|$, may only occur at a parameter value where noise-induced escape is already common, i.e., when $|\lambda_1|$ approaches $|\lambda_2|$ and $|\lambda_3|$ from below and reaches the same order of magnitude.

C. Eigenfunction reconstruction from diffusion maps and observables for early warning signals

We now reconstruct $\phi_n(x)$ via the DM approach from data of the two-dimensional (2D) double-well (19) obtained by simulation with an Euler-Maruyama scheme with time step $dt = 0.005$. We simulate an ensemble of 100 uniformly distributed initial conditions covering both wells for a fixed simulation time $T = 100$, allowing the ensemble to converge to $\pi(\mathbf{x})$. We only use data after $t = 75$, i.e., we discard any transient dynamics. Finally, the simulated data is subsampled and all ensemble members are combined to yield a set of 15 000 data points.

The scaled eigenvectors ξ_n (i.e., the diffusion coordinates) obtained from a spectral decomposition of the Markov matrix (13) define functions that can be evaluated approximately at any point in the original phase space via Eqs. (14) and (15). Evaluation on an evenly spaced grid for the 2D double-well shows that the first few nontrivial diffusion coordinates $\xi_{1,2,3,4}$ are indeed in good qualitative agreement with the (scaled) eigenfunctions $\phi_{1,2,3,4}$ obtained from the discretized FP operator [compare Figs. 4(a)–4(d) and Figs. 3(g)–3(j)].

We now restrict our attention to the scenario of bifurcation-induced tipping, where the system resides in one of the metastable sets, and where we consider timescales much shorter than $|\lambda_1|^{-1}$ associated with noise-induced escape. We consider the well that contains the base state with $x > 0$




FIG. 3. (a)–(j) Forward (a)–(e) and backward (g)–(j) FP eigenfunctions of the two-dimensional double well model (19) with $\sigma = 0.6$ and control parameter $e = 0.5$, computed using the method by Chang and Cooper [47]. The black contour depicts the level where $\psi_n = \phi_n = 0$. The potential $V(x, y)$ of the system is shown as level sets in panel (f). The instanton (computed by the method in Ref. [48]) is drawn in purple, and the basin boundary in green. (k)–(t) Same but for the model with $\sigma = 0.3$ and control parameter $e = 1.0$, which is closer to the bifurcation at $e \approx 1.73$ compared to the case in panels (a)–(j). Note there is numerical noise due to the very low probabilities that occur at the steepest parts of the potential around the boundaries of the domain. This produces numerical artefacts in the zero contour-line of the eigenfunctions, where erroneously the values in the computed eigenfunctions rapidly alternate in sign.

(cf. Fig. 4 or Fig. 3). Assuming a slowly varying control parameter, the system relaxes toward a quasistationary distribution $p_{qs}(x) \approx \psi_0(x) + c_1\psi_1(x)$ at any given instantaneous control parameter value. Here, $c_1\psi_1$ compensates ψ_0 such that all mass is concentrated in the shallow well. The dominant eigenfunction ϕ_1 is approximately constant in the shallow well, and thus variations in expectation values are determined by ϕ_2 and onward. When sufficiently close to the bifurcation, the first backward eigenfunctions are almost constant $\phi_2 \approx$

$\phi_3 \approx \phi_4 \approx 0$ in the basin of the alternative state [Figs. 3(q)–3(t)]. Hence, it should be possible to approximate them from data restricted to the basin of the base state.

We compute DMs from simulation data with initial conditions restricted to a square domain around the base state that lies entirely within its basin. The initial conditions quickly converge to the quasistationary distribution $p_{qs}(x)$ in the shallow well and the transient during equilibration is discarded. A small noise strength $\sigma = 0.3$ is used, ensuring




FIG. 4. Reconstruction of eigenfunctions of the 2D double-well (19) using the DM algorithm. (a)–(d) First four diffusion coordinates $\xi_{1,2,3,4}$ that approximate $\phi_{1,2,3,4}$, obtained at control parameter $e = 0.5$ and noise strength $\sigma = 0.6$, and estimated by DM on simulated data sampling the whole phase space. The diffusion coordinates are evaluated at evenly spaced grid points using Eqs. (14) and (15). (e), (f) First two nontrivial diffusion coordinates of Eq. (19) from simulated data restricted to dynamics that remains in the shallow well, with control parameter $e = 0.5$ and noise strength $\sigma = 0.3$.

that noise-induced transitions to the other well during the simulation time are extremely rare. Realization leading to transitions are discarded. The features of the first two nontrivial eigenfunctions [Figs. 4(e) and 4(f)] are consistent with the corresponding (higher) eigenfunctions obtained from the full

state space [Figs. 3(r), 3(s) or Figs. 3(i), 3(j)]. Specifically, the level sets show that ξ_1 is increasing monotonically and nonlinearly toward the saddle. Level sets of ξ_2 have similar shape, but do not exhibit a monotonic increase toward the saddle. In fact, there is a quadratic relationship between ξ_2 and ξ_1 .




FIG. 5. (a), (d) Simulated data points of the 2D double-well (19) in the space spanned by the first three diffusion coordinates, for parameter value $e = 0.5$ far from the bifurcation (left) and for $e = 1.5$ (right), which is closer to the bifurcation at $e \approx 1.73$. The lower panels (b), (c), (e), (f) show the same data in two-dimensional projections onto the diffusion coordinates (ξ_1, ξ_2) and (ξ_1, ξ_3) . We use a lower noise level of $\sigma = 0.09$ to obtain simulation data restricted to the shallow well when very close to the bifurcation, and thus the relation of ξ_1 and ξ_2 is different compared to Figs. 4(e) and 4(f), where $\sigma = 0.3$.




FIG. 6. Observables for EWS in the 2D double well model (19). (a) Distribution of the ξ_1 values evaluated at simulated data points using $e = 1.5$ and $\sigma = 0.09$ (red line). The black dashed line is a distribution of the same ξ_1 (i.e., obtained from the DM at $e = 1.5$), but evaluated on data points simulated at $e = 0.5$ via the Nyström interpolation (14) and (15). (b)–(d) The diffusion coordinate ξ_1 , estimated for simulation data close to the tipping point ($e = 1.65$), is fit to a polynomial of the state variables (x, y), and used as observable to detect CSD by evaluating it on residual data of simulations at lower values of the control parameter e . All simulations are initialized in the shallow well. Shown are distributions of the values of a linear (d) and cubic (c) polynomial, normalized to the fluctuations at $e = 0.2$. Panel (b) shows the variance increase of the linear and nonlinear observable, as well as the observable $x - y$, normalized to the variance of the fluctuations at $e = 0.25$. (e)–(g) Polynomial fits of ξ_1 estimated from simulations (dots) with $\sigma = 0.09$. Panel (e) is a cubic fit to data at the parameter value $e = 0.5$ far from the bifurcation, and (f), (g) is a cubic and linear fit for $e = 1.5$, which is closer to the bifurcation.

The leading ξ_n are not necessarily all independent d.o.f.'s. In the system restricted to the shallow well, the critical d.o.f. becomes the slowest upon approaching the bifurcation and the timescale separation with respect to all other d.o.f.'s becomes larger, at which point the first few backward eigenfunctions (and hence the first few diffusion coordinates) are all expected to parametrize the slowest d.o.f. This is particular for single-well systems. For instance, in a multidimensional parabolic potential with a slow variable x and a spectral gap, ξ_1 is a function of x and the next k eigenfunctions ξ_k (with k dependent on the magnitude of the spectral gap) are polynomially related to ξ_1 [32,42]. In this case, ξ_1 is sufficient as a reduction coordinate and is the only diffusion coordinate that indicates monotonically how far fluctuations evolve toward the saddle. Accordingly, in the 2D double-well restricted to the shallow well, the dynamics in the space of the diffusion coordinates $\xi_{1,2,3}$ evolves from a two-dimensional to a one-dimensional manifold as the bifurcation is approached (Fig. 5).

The observations above suggest three main ways to leverage information contained in the diffusion coordinates ξ_n for early warning of TPs. First, one may observe the qualitative change of the functional dependencies of the first ξ_n as a result of the emerging timescale separation, as just discussed.

Second, one can directly evaluate and compare ξ_1 for data sets obtained at different observational time slices (i.e., for different values of the control parameter) via the Nyström interpolation (14) and (15). In particular, we can estimate ξ_1 from a data set believed to be closest to a TP, for example from climate observations closest to present-day, and then evaluate the observable ξ_1 on data sets sampled further away from the TP, for example using climate observations of the past. If the variability and correlation of the values of ξ_1 is increased significantly in the former data set, this is an indication of decrease in resilience in the critical d.o.f. and an impending TP. This is illustrated in Fig. 6(a), where a clear change in variability is seen when evaluating ξ_1 estimated at $e = 1.5$ on data simulated farther away from the bifurcation at $e = 0.5$. Note that the same normalization of the variables that is applied before the DM algorithm to the data set where ξ_1 is estimated has to be applied to data at parameter values further away from the bifurcation. Here and in most of the following we focus our presentation on the variance as EWS, but similar plots could be shown for the autocorrelation.

Third, an explicit expression of a physical observable can be constructed from ξ_1 as function of (possibly a subset of) the state variables. ξ_1 is estimated for an observational time slice




FIG. 7. (a)–(d) Scatterplots of simulated data from the five-box AMOC model in projections onto the first three diffusion coordinates, for two values of the control parameter $H = 0.25$ (a), (b) and $H = 0.315$ (c), (d). (e) Physical observable associated with the leading diffusion coordinate ξ_1 , using as cubic polynomial of the variables S_N and S_T , at $H = 0.315$. The position of the edge state is marked by the green triangle, and the base attractor is the red dot. The blue point cloud is the data at $H = 0.315$ used for the construction of the DM. The green point cloud is corresponding simulation data further from the bifurcation ($H = 0.25$), shifted such that the mean is at the base attractor at $H = 0.315$.

closest to the TP and the values of ξ_1 at the data points (i.e., the entries of the eigenvectors of the Markov matrix M) are fit to a suitable function, e.g., a polynomial. The fitted function can then be evaluated for any data sets further away from the TP. In Figs. 6(e)–6(g) we show polynomial fits to ξ_1 estimated from simulation data of the 2D double-well. The directionality of the level sets is consistent with the direction of the edge state, as expected. Far from the bifurcation (assuming low noise) the dynamics samples only the relatively flat part of ξ_1 far from the saddle point, provided the noise is sufficiently small, as shown in Fig. 6(e). Still, the fitted function shows a more rapid decrease toward the saddle point, thus indicating that it already carries the crucial information for detecting CSD. When close to the TP, nonlinear functions tend to be required for an adequate fit of the ξ_1 data [Fig. 6(f)]. However, linear fits preserve the directionality of the edge state well [Fig. 6(g)]. Using the fits estimated from data sampled close to the TP [Figs. 6(f) and 6(g)], we again see that the variability of the values of ξ_1 decreases when evaluated for data sample further away from the TP [Figs. 6(c) and 6(d)]. The nonlinear observable shows a significantly stronger change in variance compared to the linear one [Fig. 6(b)].

Similar results are obtained when applying the method to a slightly more complex nongradient system, a four-dimensional conceptual model of the AMOC [49], which is summarized in the following. The variables ($\{S_N, S_T, S_S, S_I\}$) are the average salinities in four boxes of the global ocean (see Sec. B for the equations). Large differences in the box volumes gives a timescale separation, with S_S and S_I being the slowest variables, and S_N being the fastest. The model is

bistable for a range of the control parameter H from $H \approx 0.04$ until the bifurcation at $H \approx 0.3214$, where the stable fixed point corresponding to a present-day AMOC disappears.

We use simulation data restricted to the present-day AMOC state. Far from the bifurcation, ξ_1 represents a correlated relaxation mode in the slow variables S_S and S_I . ξ_2 is a quadratic function of ξ_1 , and ξ_3 is independent and strongly correlated with S_N only [Figs. 7(a) and 7(b)]. When approaching the TP, the slow mode corresponding to relaxation along the direction of the edge state emerges due to CSD, and eventually rises to the position of ξ_1 , with ξ_2 (ξ_3) being a quadratic (cubic) function thereof [Figs. 7(c) and 7(d)]. ξ_1 exhibits a strong nonlinear anticorrelation with S_N and a nonlinear positive correlation with S_T (see also Ref. [15]). A good physical observable representing ξ_1 is found by considering polynomial functions in a projected space of a subset of (e.g., two) variables. The best cubic polynomial fit of ξ_1 (estimated at $H = 0.315$) is a function of S_N and S_T . The resulting function shows a monotonic, nonlinear increase from the base state toward the edge state, while being essentially flat in other directions [Fig. 7(e)]. This renders the observable very sensitive to excursions toward the edge state, and thus ideal for EWS.

D. Observables and extrapolation of tipping times

So far we discussed the construction of observables that show a large increase in variance (and also autocorrelation) as a result of CSD, which is a qualitative indication that the system moves toward a TP. One may go further and attempt a quantitative prediction of the expected time of tipping by

extrapolating the CSD signal in observational time series, as was done in the context of real-world climate observations in Refs. [20,24]. For this one needs to assume that an observed time series samples the critical dynamics by obeying the SNB normal form. This implies that such a prediction is sensitive to the choice of observable, as shown in the following.

Consider the general multidimensional system described by the coupled stochastic differential equations (1), where we now assume that the drift $b^y(\mathbf{X}, \mu)$ depends on a control parameter μ and the noise $\sigma_v^y(\mathbf{X}) = \sigma_v^y$ is additive. If the system undergoes an SNB, the noise-driven dynamics is expected to become restricted to the vicinity of a one-dimensional center manifold, and is described by the normal form

$$dx = (x^2 - \mu)dt + \sigma dW_t, \quad (20)$$

where $\mu = 0$ demarcates the bifurcation. Close to the fixed point, the system can be linearized and approximated by the Ornstein-Uhlenbeck process $d\tilde{x}_t = -\lambda\tilde{x}dt + \sigma dW_t$. Crucially, the linear restoring rate λ is related to the control parameter with $\lambda = 2\sqrt{\mu}$. Data sampled at small time intervals Δ_t can be approximated by an AR(1) process

$$X_{k+1} = e^{-\lambda\Delta t} X_k + \epsilon_k, \quad (21)$$

where ϵ_k are Gaussian random variables with variance $\sigma^2(2\lambda)^{-1}(1 - e^{-2\lambda\Delta t})$. For this process the autocorrelation at lag 1 is given by $\rho_1 = e^{-\lambda\Delta t}$. Since $\lambda = 2\sqrt{\mu}$, we can reconstruct μ from data by

$$\mu = \left(\frac{\ln \rho_1}{2\Delta t} \right)^2. \quad (22)$$

Thus, in a sliding window one can estimate ρ_1 as a function of time, and, assuming a linear trend in μ , estimate with a linear fit to the function on the right-hand side at what time the control parameter will cross zero, which is exactly when the autocorrelation tends to 1 at the SNB.

If given data from an arbitrary scalar observable of the system, this extrapolation to the time of tipping can fail since the observable need not obey the SNB normal form, or may only approximately do so when arbitrarily close to the bifurcation. In fact, even for univariate bistable systems (where the question of observable is obsolete), the scaling $\mu \propto [\ln(\rho_1(X))]^2$ according to the saddle-node normal form only applies when close to the bifurcation. For the 1D double-well (16) under a linear change in time of the control parameter μ the function $[\ln(\rho_1(X))]^2$ is convex with respect to time (Fig. 8). Thus, the TP, which occurs at $t \approx 420$, would be predicted too early at $t \approx 390$ [red line in Fig. 8(b)] when the linear extrapolation is performed based on data not close enough to the bifurcation. An estimation of the autocorrelation in a moving window introduces a further bias toward a later estimated tipping, because the autocorrelation is underestimated by removing some correlation during the necessary step of detrending within each window.

As more dimensions are involved the prediction depends on the choice of observable. For the 2D double-well (19), the variable x can give an accurate prediction when data are available close enough to the TP [Fig. 9(b)]. In contrast, y initially shows a quasilinear relation of μ and $[\ln(\rho_1(X))]^2$, but then a much steeper relationship closer to the bifurcation




FIG. 8. Simulations of the 1D double well (16) with $\sigma = 0.05$ and a linear increase of the control parameter within 500 time units from $\beta = -0.7$ past the bifurcation point $\beta_c = \sqrt{8/27}$ to $\beta = 0.7$. (a) Time series of one simulation overlaid on the bifurcation diagram. (b) Reconstructed control parameter (22) obtained by the lag-1 autocorrelation ρ_1 , estimated at each time step ($\Delta t = 0.05$ time units) for an ensemble of 15 000 simulations (black trajectory) until a cutoff time where 98% of the ensemble members have not tipped yet (evaluated by crossing a threshold of $x = 0.2$). At the noise level $\sigma = 0.05$, this is only very shortly before the bifurcation is reached. The 2% of realizations that tipped are removed. The solid lines are linear fits using different segments of the data, with the red (blue) line using the first half (last sixth) of the data. The crosses as well as dotted and dashed lines are for ρ_1 estimated in a moving window of length 40.

[Fig. 9(c)]. Extrapolating of the initial slope would lead to a tipping time estimate that is far too late. Especially unsuited observables exist, such as $O(x, y) = x - y$, where no tipping can be predicted before a noise-induced transitions would occur [Fig. 9(d)]. In contrast, $O(x, y) = x + y$ is very closely aligned with the direction of the edge state [15] and permits to predict the time of tipping accurately for data sufficiently close to the bifurcation [Fig. 9(e)]. The nonlinear observable obtained from the DM approximation to the first subdominant backward eigenfunction [Fig. 6(f)] is most accurate, even when evaluated at data far from the bifurcation [Fig. 9(f)].

E. Application to tipping points in a global ocean model

As the final result, we show our method is capable of successfully detecting TPs in a high-dimensional system. We consider the global ocean model Veros [50], which shows a TP of the AMOC from its present-day state to a collapsed




FIG. 9. Same as Fig. 8, but for simulations with the 2D double-well (19) with $\sigma_x = \sigma_y = 0.05$, where the parameter e is increased linearly past the bifurcation. Shown in panel (a) is a single example trajectory overlaid on the bifurcation diagram. Panels (b)–(f) show the evolution of the quantity $[\ln(\rho_1(O))]^2(2\Delta t)^{-2}$ for different observables $O(x, y)$, along with linear fits to different parts of the time series.

state as a result of increasing meltwater input to the North Atlantic. Veros is a primitive-equation finite-difference ocean model forced with a fixed atmospheric climatology, and discretized on a grid of 40 latitudinal and 90 longitudinal grid points, as well as 40 depth levels. This is a coarse-resolution setup, but it enables long steady-state simulations to get good statistics beneficial for our feasibility study. As a dynamical system, the model possesses almost one million degrees of freedom. For more details on the model, see Refs. [50–52]. The meltwater input F is the control parameter, and the stability landscape with respect to F (computed in Ref. [52]) is shown in Fig. 10. There are several branches of attractors with an AMOC similar to present-day, but these collapse at a high freshwater forcing of $F \approx 0.36$. After this TP, there remain only attractors with a collapsed AMOC.

We use four 33 000-year-long equilibrated simulations performed at four fixed values of F leading up to the TP, and sampled as 5-year averages of the state variables. These are referred to as simulations I to IV [see Fig. 10]. Additive surface temperature and salinity noise drives fluctuations of the system around its deterministic attractors (for more details see Ref. [15]), which otherwise feature relatively small-amplitude chaotic oscillations [52]. The system without noise forcing has been investigated previously to determine an edge state on the separatrix of the present-day and collapsed AMOC regimes [53]. By analyzing its mean climatology it was found that the edge state distinguishes itself most strongly from the mean states on the attractors in terms of its fresh and cold deep Atlantic. This is the “fingerprint” of the edge state. Subsequently, it was argued that increased fluctuations toward

the edge state as a result of CSD should be most prominent in variables quantifying this fingerprint [15]. Indeed, only a very small subset of all d.o.f., coinciding exactly with the variables describing the deep ocean fingerprint, shows sig-




FIG. 10. Bifurcation diagram of the Veros ocean model without noise forcing (obtained in Ref. [52]), with the maximum AMOC strength as order parameter, and the freshwater forcing F as control parameter. All individual solid lines correspond to different branches of attractors, and the dotted lines as well as the arrows indicate the transition path of the system as a given attractor loses stability. The edge state at $F = 0.3472$ (computed in Ref. [53]) is marked by the green triangle.




FIG. 11. (a)–(h) Example time series of two variables in the four Veros data sets used here. Shown is the north Atlantic deep ocean salinity (a)–(d) at the four different values of the control parameter F , as well as the south Atlantic deep ocean density (e)–(h). The time series are shown as anomalies with the mean removed. (i)–(l) Time series of an observable constructed by projection onto the leading diffusion coordinate ξ_1 in the subspace of deep ocean salinity. The eigenmode was estimated with the DM algorithm from data sampled at $F = 0.3557$ Sv.

nificant variance increase prior to the AMOC collapse [15]. The variable that was found to exhibit the largest increase in variance is shown in Figs. 11(a)–11(d) across the four data sets. In the following, we show that similar (if not better) results can be obtained with our the DM method which does not require prior knowledge of the edge state or a brute-force search across all d.o.f. (risking false positives). Instead, only observational data close to the TP is required.

It would be feasible to compute the DM distance Kernel in the full space of the three-dimensional fields, perhaps after a weighting of the different physical units (temperature, salinity, density and velocity). But for simplicity we perform an initial dimensionality reduction, by averaging the salinity, temperature and density fields over boxes covering the entire ocean at different depths, and by summarizing the strength of the main ocean currents in terms of the spatial maxima of the meridional and barotropic stream functions (see Ref. [53]). This yields time series of 83 variables covering most important aspects of the model state. Figures 11(a)–11(h) shows time series of two of these variables at four different values of F , with the mean removed. Before applying the DM algorithm, we normalize all variables to have unit variance. A bandwidth of $\epsilon = 17$ for the Kernel (11) was found to be optimal to resolve the data manifold at all parameter values F without being influenced by outliers.

The physical meaning of the first two inferred diffusion coordinates is summarized in Table I, where the five physical variables with the highest correlation to ξ_1 and ξ_2 are listed. As the control parameter is changed toward the TP, we can see that the expected critical mode emerges. When far from the TP, ξ_1 is best correlated with deep ocean density in the Indo-Pacific, South Atlantic, and Southern Ocean. The next mode, represented by ξ_2 , is best explained by variability in the tropical subsurface ocean. When increasing the control parameter to $F = 0.3515$, it is replaced by a mode correlated with temperature and salinity in the deep northern and tropical Atlantic. Increasing F further to $F = 0.3557$ shortly before the TP, this becomes the leading mode ξ_1 . The most important variables in this mode (upper right column in Table I) are exactly those that make up the fingerprint of the edge state [53] and feature the largest increase in variance [15].

To design a physical observables for EWS from ξ_1 , the first option is again to use the Nyström extension based on all variables and interpolate the function ξ_1 to observations further back in time, to find evidence for increased fluctuations in the critical mode. One may also fit ξ_1 as a linear or nonlinear function of the variables. In high dimensions it is sensible to only consider a subset of variables to find a parsimonious observable with the best signal-to-noise ratio when applied for EWS. We leave a treatment of this statistical

TABLE I. Spearman correlation of the top five Veros variables with the first two diffusion coordinates ξ_1 and ξ_2 , where the ξ_n were estimated at the four different values of the control parameter F . A shorthand notation encodes the meaning of the variables. For example, “rho deep TP” is deep ocean density in the tropical Pacific. The first part (“rho,” “salt,” or “temp”) refers to density, salinity or temperature, respectively. The second part (“subs” or “deep”) refers to the subsurface ocean averaged until 1000 m depth or the deep ocean below 1000 m. The last part is geographical location, which can be the northern, tropical or southern Atlantic or Pacific (NA, TA, SA, NP, TP, SP), or the Southern Ocean (SO) and the Indian Ocean (IO). In bold text are key variables related to the mode of cold and fresh excursions in the deep north and tropical Atlantic directed toward the edge state. In italic are key variables related to the mode of fast cold excursions in the Southern ocean, which strongly increase the density in large parts of the deep ocean.

	$F = 0.2957$		$F = 0.3472$		$F = 0.3515$		$F = 0.3557$	
	Variable	r_S	Variable	r_S	Variable	r_S	Variable	r_S
ξ_1	rho deep TP	0.854	<i>rho deep SA</i>	−0.874	<i>rho deep SA</i>	−0.886	salt deep NA	−0.873
	rho deep TA	0.848	rho deep TA	−0.870	rho deep TA	−0.868	salt subs NA	−0.834
	<i>rho deep SP</i>	0.835	<i>rho deep IO</i>	−0.842	<i>rho deep IO</i>	−0.860	temp deep NA	−0.789
	<i>rho deep IO</i>	0.822	<i>rho deep SP</i>	−0.803	<i>rho deep SP</i>	−0.827	salt deep TA	−0.735
	<i>rho deep SA</i>	0.813	<i>rho deep SO</i>	−0.794	<i>rho deep SO</i>	−0.821	temp deep IO	0.730
ξ_2	temp subs TP	0.627	temp subs TA	0.624	salt deep NA	−0.775	<i>rho deep SO</i>	−0.782
	rho subs TP	−0.607	rho subs TA	−0.615	salt subs NA	−0.765	<i>rho deep SA</i>	−0.763
	rho subs TA	−0.597	rho subs TP	−0.567	temp deep NA	−0.690	<i>rho deep SP</i>	−0.723
	temp subs TA	0.582	salt subs NA	−0.563	temp subs TA	0.667	salt deep SO	−0.701
	temp deep TP	0.518	temp subs TP	0.554	rho subs TA	−0.643	rho deep TP	−0.681

optimization problem for future research, and consider here three simple examples. First, simply take the variable with highest correlation to ξ_1 as observable. This is north Atlantic deep ocean salinity [Figs. 11(a)–11(d)], which indeed has the highest increase in variance of all individual features [15], increasing by a factor of 12.23 when going from $F = 0.2957$ to $F = 0.3557$. Second, fit ξ_1 to a linear combination of the two best features in Table I (at $F = 0.3557$). This yields a variance increase by a factor of 14.21. Third, fit ξ_1 to a cubic polynomial of two variables, which are chosen as the two features among the best four (Table I) that represent two different ocean sectors, thus giving some degree of spatial independence. This observable captures the directionality of the edge state well (Fig. 12), and shows an increase in variance by a factor of 12.60.

Finally, ξ_1 can be mapped back into the physical space of full dimension, whereafter spatiotemporal anomalies from different time periods can be projected onto the mode. In particular, we propose to obtain the physical representation of this mode by averaging over the time points where ξ_1 is extremized. For instance, one can choose the data points with the top 5% largest and 5% smallest values of ξ_1 . Averaging independently over these two sets of data points defines a positive and a negative phase of the mode. By taking the difference of the positive and negative patterns we obtain a pattern that describes the mode as a whole and that we can project data onto. This may be viewed as linear approximation that interpolates the physical mode linearly as a function of the value of ξ_1 . Figure 13 shows the first subdominant mode extracted in this way, comparing results far from [Figs. 13(a) and 13(d)] and close to the TP [Figs. 13(b) and 13(e)]. The modes are projected down to the two-dimensional physical space of vertically averaged deep ocean of temperature and salinity. Far from the TP, the mode is characterized by a global cooling of the deep ocean initiated by abrupt (decadal-scale) cooling events in the Southern Ocean (seen as spikes of density increase in Figs. 11(e)–11(h), which are excited by

noise in the multistability regime, as discussed in Ref. [15]. The resulting dense deep ocean water spreads throughout the deep ocean, before the signal decays on a multicentennial timescale. There is no significant salinity signal.

Close to the TP, the new physical mode that has emerged as subdominant eigenfunction is characterized by a cold and fresh anomaly of the deep northern and tropical Atlantic [Figs. 13(b) and 13(e)], with a spatial pattern that very closely resembles the anomaly of the edge state [Figs. 13(c) and 13(f)]. A scalar observable is created by projecting the snapshots of the data fields (at each time step, and as anomalies




FIG. 12. Two-dimensional observable (color map and contours) created from a cubic polynomial fit of two variables of the Veros data (north and tropical Atlantic deep ocean salinity at parameter value $F = 0.3557$) to the values of ξ_1 obtained from the diffusion map. The point cloud depicts the data with color coding according to corresponding the value of the diffusion coordinate ξ_1 . Shown are anomalies with respect to the mean state of the model. The green line is a vector pointing from this mean state to the edge state [15].




FIG. 13. (a), (b), (d), (e) Anomaly maps of deep ocean (vertical average below 1000 m) temperature (a), (b) and salinity (d), (e) describing the first subdominant spatiotemporal mode ξ_1 of the Veros model obtained by the DM algorithm at the control parameter values $F = 0.3472$ (a), (d) and $F = 0.3557$ (b), (e). (c), (f) Corresponding anomaly map of the edge state with respect to the mean state on an attractor with active AMOC (state II in Fig. 10), as estimated from the deterministic version of the model in Ref. [53].

with respect to a mean state) onto this pattern via the scalar product of the two fields. Comparing different observational slices (such as the simulations I to IV) then reveals the changes in variability of this critical mode, which serves as EWS. The variance increases by a factor of 21.05 when going from states I to IV, and the time series of this observable are shown in Figs. 11(i)–11(l).

While these are encouraging results, in practice there is a risk of false positives and negatives, since here the critical physical mode only appears in its correct place at ξ_1 when already quite close to the TP. One needs to verify that the leading mode captured by ξ_1 is likely a critical mode. The guiding principle should be to look how strongly fluctuations in the leading diffusion coordinate, estimated at the current time slice, have increased compared to data slices back in time, and then set a level of statistical significance based on a reference period. In the Veros data one can in this way rule out the initially leading mode [Figs. 13(a) and 13(d)], because its variability does not keep increasing toward the TP [Figs. 11(f)–11(h)], and the associated excursions do not last longer. But in principle there remains a chance for false positives if a new timescale separation would arise upon change of a control parameter for reasons not related to a TP.

IV. DISCUSSION

Here we propose a method to obtain observables for detecting CSD before bifurcation-induced TPs [28] in multistable systems via a data-driven approximation of the FP operator. The scenario of bifurcation-induced tipping implies low noise, where the first k FP eigenfunctions ψ_k are very slowly decaying modes related to rare noise-induced escape between metastable states [54]. The subsequent eigenfunctions describe probability density patterns in phase space that relax slowest toward a quasistationary distribution centered around a metastable state. We consider without loss of generality the bistable situation with $k = 1$, where at the TP a base state loses stability and the system transitions to its alternative state. In practical applications, the system prior to the TP is only

observed in the basin of the base state and the contribution of ψ_1 remains quasiconstant (until infinitesimally close to the TP). In this case, the first nontrivial eigenfunction that can be observed in data is ψ_2 , which describes the slowest relaxation mode toward the quasistationary distribution around the base state within its basin of attraction. As the TP is approached, due to CSD this mode will eventually represent the slowing relaxation along the critical d.o.f.

We suggest to approximate the corresponding eigenfunction of the backward FP operator by the first diffusion coordinate ξ_1 , obtained as scaled eigenvector of the DM Markov matrix. From ξ_1 , one can obtain a physical observable—e.g., by projecting onto patterns obtained as average over data points that extremize ξ_1 —that shows a monotonic increase from the base state along the critical d.o.f. toward the edge state. With several examples of low-dimensional bistable models we demonstrated that such an observable shows increases in noise-driven fluctuations that provide robust statistical EWS of the CSD associated with the impending TP. We also showed that measuring CSD in the correct observable is crucial when attempting to predict the time of tipping by extrapolating the scaling of variance or autocorrelation of a scalar time series based on the SNB normal form. The method was successfully applied to simulation data from a high-dimensional global ocean model that features a TP of the AMOC. A critical mode was extracted that is in excellent agreement with the mode that would be expected from knowledge of the edge state [15,53], and a scalar observable was derived from ξ_1 that shows highly significant increases in variance that are useful as statistical EWS.

A general caveat for this and other methods aiming to measure CSD from high-dimensional systems is that the critical mode may only emerge as the first subdominant eigenmode when already very close to the TP. If there are competing noncritical modes that are very slow, it may not be possible to identify the critical mode, since one would need a very long observational time horizon while being close enough to the TP where the critical mode is finally the slowest. Indeed, it was seen for the Veros model that the correct physical mode

takes on the role of ξ_1 , but only as the system was quite close to the TP. Unless one knows beforehand which d.o.f. should be measured, for instance by knowledge of the edge state [15] or robust physical considerations, it may only be possible to issue a reliable warning when very close to the TP, and potentially only after the probability of noise-induced transitions has become substantial.

There are other previously proposed dimensionality reduction methods aiming to extract a scalar observable that can be used to detect CSD. These include variance-based techniques [8,55–57], where the first principle component is identified by empirical orthogonal functions—as originally proposed to obtain the critical mode for EWS [8]—or principal oscillation patterns. An autocorrelation-based method has also been proposed, where the directions of maximum variance of the first differences of multivariate time series are found, which gives components of high autocorrelation that should indicate directions of lowest resilience [58]. Other methods look for an SNB in the full set of (observed) variables via the eigenvalues of a reconstructed Jacobian, which is determined by fitting a multivariate autoregressive model [59] or a multidimensional Langevin equation [44].

Our approach is distinguished by combining several attributes. It yields an observable derived from the first subdominant backward FP eigenfunction that is designed to represent the critical mode displaying CSD, based on the flattening of the quasipotential expected for a broader class of TPs. This mirrors the reasoning for a natural tipping observable recently proposed in Ref. [29], and it is also supported by other operator-theoretic work on the topic [60–64]. The DM algorithm that we propose to use is a nonlinear dimensionality reduction method and thus allows for observables to be nonlinear functions of the state variables. It can be deployed for relatively high dimensional systems, since the quality of its approximation for a given sample size is not dependent on the full state space dimension, but on the intrinsic dimension of the underlying data manifold, which may be much lower. An equal time spacing of data points, or any time ordering at all, is not required, thus increasing its applicability for instance to paleoclimate data [65]. The method furthermore allows for a qualitative assessment of changes in the dominant physical modes, by observing the functional relation between eigenfunctions, and the relation of eigenfunctions and physical variables. This is useful for determining an emerging timescale separation before the TP, and it may help to prevent false positives for instance due to changing amplitude or correlation of the driving noise. Estimating the eigenfunctions with the generalized DM method presented in Ref. [36] also allows treating multiplicative noise and nongradient dynamics, which can alleviate the issues that unknown changes in oscillatory modes cause for statistical EWS [52]. This method gives equivalent results when applied to the ocean model (see Appendix A).

Future work should address how systems driven far away from a steady state need to be treated. In many cases of real-world relevance, such as tipping of the polar ice sheets, the change in the control parameter is fast compared to the critical relaxation mode ψ_2 and perhaps many other modes. This means that the system is not in a quasistationary state, as was assumed here, and it is likely that the critical mode is

not displayed before crossing the TP. An extended framework based on nonautonomous dynamical systems theory can hopefully yield useful insights on whether EWS before the de-facto TP exist in this case.

ACKNOWLEDGMENTS

G.A.G. acknowledges valuable discussions with Matheus Mankato de Castro and Caroline Wormell. We thank R. Nuterman and the Danish Center for Climate Computing for supporting the simulations with the Veros ocean model. J.L. has received support from Danmarks Frie Forskningsfond (Grant No. 2032-00346B) and research Grant No. VIL59164 from VILLUM FONDEN.

DATA AVAILABILITY

The data that support the findings of this article are openly available [67]; embargo periods may apply.

APPENDIX A: EIGENFUNCTIONS IN NONGRADIENT SYSTEMS WITH NONADDITIVE NOISE ESTIMATED BY LOCAL KERNEL DIFFUSION MAPS

Here we present how the Fokker-Planck and backward Kolmogorov eigenfunctions can be estimated without assuming gradient dynamics or additive noise, and demonstrate that similar results are obtained for data from the Veros model as with the standard diffusion map method presented in the main text. It can be achieved by the method introduced in Ref. [36], which generalizes diffusion maps to apply to SDEs of the form of Eq. (1). In addition to the isotropic Kernel (11), an anisotropic Kernel with a second bandwidth ϵ_l , as introduced in Ref. [66], is used:

$$K_{\epsilon_l}(\mathbf{x}, \mathbf{y}) = \exp(-(4\epsilon_l)^{-1}(\mathbf{x} - \mathbf{y} + \epsilon_l \hat{\mathbf{b}}(\mathbf{x}))^\top (\hat{\mathbf{A}}(\mathbf{x}))^{-1} \times (\mathbf{x} - \mathbf{y} + \epsilon_l \hat{\mathbf{b}}(\mathbf{x}))). \quad (\text{A1})$$

This requires estimates of the drift $\hat{\mathbf{b}}(\mathbf{x})$ and diffusion matrix $\hat{\mathbf{A}}(\mathbf{x})$ locally at each data point $\mathbf{x}$. To do so we use Kramers-Moyal expansion estimates

$$\hat{\mathbf{b}}(\mathbf{x}) = \lim_{\tau \rightarrow 0} \frac{1}{\tau} \mathbf{E}[X_\tau - X_0 | X_0 = \mathbf{x}], \quad (\text{A2})$$

$$\hat{\mathbf{A}}(\mathbf{x}) = \lim_{\tau \rightarrow 0} \frac{1}{2\tau} \mathbf{Cov}[X_\tau - X_0 | X_0 = \mathbf{x}]. \quad (\text{A3})$$

Ideally, to calculate expectation values, one has many realizations with short integration time τ at each data point as initial condition $\mathbf{x}$. But in the case of the Veros model, we only have one long realization (for each parameter value), and thus we estimate drift and diffusion at each data point $\mathbf{x}$ by computing expectation and covariance over the set of increments $\{Y_{n_i+\tau} - Y_{n_i}\}_{i=1}^n$, where $\{Y_{n_i}\}_{i=1}^n$ are the n nearest neighbors of $\mathbf{x}$. While we cannot take the limit $\tau \rightarrow 0$, the sample time τ of 5 years is short compared to the long timescales of the deep ocean.

After computing the anisotropic Kernel $(K_{\epsilon_l})_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$ for all pairs of data points, it is normalized using the row sum of the isotropic Kernel $p_\epsilon(\mathbf{x}_i) = \sum_{j=1}^N K(\mathbf{x}_i, \mathbf{x}_j)$, i.e., a Kernel density estimate of the data sample. Thus, with the diagonal matrix $(D_\epsilon)_{ii} = p_\epsilon^{-1}(\mathbf{x}_i)$, form $\tilde{K}_{\epsilon_l} = K_{\epsilon_l} D_\epsilon$. This




FIG. 14. (a)–(d) Anomaly maps of deep ocean (vertical average below 1000 m) temperature (a), (b) and salinity (c), (d) describing the first subdominant spatiotemporal mode ξ_1 of the Veros model obtained by the local Kernel DM algorithm [36] at the control parameter values $F = 0.3472$ (a), (c) and $F = 0.3557$ (b), (d). The modes are constructed in the same way as described in Sec. III E. (e), (f) Spectrum of the local Kernel diffusion map (higher eigenvalues discarded). Note that the reported eigenvalues of the matrix L_{ϵ_l} correspond to e^{λ_n} of the approximated generator.

normalization removes additional drift terms in the reconstructed operator, which are induced by the nonuniform sampling density. Finally, to conserve probability, the matrix is normalized by its row sums, i.e., $L_{\epsilon_l} = D_{\epsilon, \epsilon_l}^{-1} \tilde{K}_{\epsilon_l}$ with the diagonal matrix $(D_{\epsilon, \epsilon_l})_{ii} = \sum_{j=1}^N (\tilde{K}_{\epsilon_l})_{ij}$. As shown in Ref. [36], in the limits $N \rightarrow \infty$, $\epsilon \rightarrow 0$, and $\epsilon_l \rightarrow 0$ the matrix $(L_{\epsilon_l} - I)/\epsilon_l$ converges at every point to the generator (backward FP operator) $\mathcal{L}^*$ of the general SDE (1), and we can approximate $\{\phi_n\}$ by the eigenfunctions of L_{ϵ_l} . The forward FP operator can be constructed in a similar fashion [36].

For the analysis of the Veros data, we chose $\epsilon_l/\tau = 5$ and $n = 100$. As explained in Ref. [36], the bandwidth may be interpreted as a characteristic timescale, and is chosen here to match the average displacement in phase space after one sample time τ . The estimate of the inverse diffusion matrix is regularized by replacing $\hat{A}(\mathbf{x})$ with $\hat{A}(\mathbf{x}) + \eta I$ in Eq. (A1). A large regularization parameter of $\eta = 0.3$ had to be chosen due to the small sample size. In Fig. 14 the results are summarized. We focus on the parameter very close to the bifurcation ($F = 0.3557$), as well as one further away ($F = 0.3472$), just as presented in Fig. 13. The evolution of the eigenfunctions is very similar to the isotropic diffusion map method. Since the generalized method also picks up nongradient dynamics, there are oscillatory modes, but they relax more quickly and

thus appear higher up in the spectrum [Figs. 14(e) and 14(f)]. The initially dominant mode of cold excursions initiated in the Southern ocean is picked up more clearly with the generalized method [Figs. 14(a) and 14(c)]. The critical mode close to the TP [Figs. 14(b) and 14(d)] is again very similar to the fingerprint of the edge state [Figs. 13(c) and 13(f)] and thus allows strong EWS by projecting time series data onto it.

APPENDIX B: EQUATIONS FOR THE FIVE-BOX OCEAN MODEL

In this Appendix, the equation and parameter values of the five-box ocean model, originally published in Ref. [49], are described. The boxes, labeled by $X = N, S, T, IP, B$, are coupled unidirectionally by the thermohaline overturning circulation q , and bidirectionally by the wind-driven circulation. The dynamical equation for box B can be eliminated by salt conservation. The remaining boxes forced by an atmospheric freshwater flux F_X multiplied by the reference salinity $S_0 = 0.035$, which is then modulated by HA_X to emulate the effect of climate change, where H is the control parameter. The varying strength of the overturning q is proportional to the density difference in the northern and southern boxes, and the temperatures T_X are fixed everywhere except in the northern

box, where it is assumed that $T_N = \mu q + T_0$, with a global reference temperature T_0 . This yields

$$q = \lambda \frac{\alpha(T_S - T_0) + \beta(S_N - S_S)}{1 + \lambda\alpha\mu}. \quad (\text{B1})$$

In the model, $q > 0$ corresponds to an AMOC “ON” state, and it is assumed that in case of a reversed circulation $q < 0$ the unidirectional coupling by the overturning flow is reversed. This yields different dynamics for positive and negative q , and a nonsmooth system of four ODEs, using the Heaviside function $\Theta(\cdot)$:

$$V_N \frac{dS_N}{dt} = |q|[\Theta(q)(S_T - S_B) + S_B - S_N] + K_N(S_T - S_N) - (F_N + HA_N)S_0, \quad (\text{B2a})$$

$$V_T \frac{dS_T}{dt} = |q|[\Theta(q)(\gamma S_S + (1 - \gamma)S_{IP} - S_N) + S_N - S_T] + K_S(S_S - S_T) + K_N(S_N - S_T) - (F_T + HA_T)S_0, \quad (\text{B2b})$$

$$V_S \frac{dS_S}{dt} = \gamma |q|[\Theta(q)(S_B - S_T) + S_T - S_S] + K_{IP}(S_{IP} - S_S) + K_S(S_T - S_S) + \eta(S_B - S_S) - (F_S + HA_S)S_0, \quad (\text{B2c})$$

$$V_{IP} \frac{dS_{IP}}{dt} = (1 - \gamma)|q|[\Theta(q)(S_B - S_T) + S_T - S_{IP}] + K_{IP}(S_S - S_{IP}) - (F_{IP} + HA_{IP})S_0. \quad (\text{B2d})$$

Time is rescaled by $\tau_Y = 3.15 \times 10^7$ to go from seconds to years, and the remaining parameter values are listed in

TABLE II. Parameter values used for the five-box model, adapted from the FAMOUS1xCO2 calibration in Ref. [49]. $\alpha = 0.12$ (thermal coefficient) and $\beta = 790$ (haline coefficient) define a linear equation of state for the density of sea water. V_i is box volume, F_i the freshwater fluxes, T are temperatures, K_i are wind fluxes, and A_i determine the distribution of freshwater forcing. η is a mixing parameter between the S and B boxes, γ determines the proportion of water which takes the cold-water path, λ and μ are constants. Subscripts indicate box labels, $i \in \{N, T, S, IP, B\}$, and “0” indicates a global reference value.

Parameter	Value	Parameter	Value
V_N (m ³)	3.683×10^{16}	F_N (m ³ s ⁻¹)	0.375×10^6
V_T (m ³)	5.151×10^{16}	F_T (m ³ s ⁻¹)	-0.723×10^6
V_S (m ³)	10.28×10^{16}	F_S (m ³ s ⁻¹)	1.014×10^6
V_{IP} (m ³)	21.29×10^{16}	F_{IP} (m ³ s ⁻¹)	-0.666×10^6
V_B (m ³)	88.12×10^{16}	F_B (m ³ s ⁻¹)	0
A_N	0.194	η (m ³ s ⁻¹)	66.061×10^6
A_T	0.597	γ	0.1
A_S	-0.226	λ (m ⁶ kg ⁻¹ s ⁻¹)	2.66×10^7
A_{IP}	-0.565	μ (°Cm ⁻³ s)	7.0×10^{-8}
K_N (m ³ s ⁻¹)	5.439×10^6	T_S (°C)	5.571
K_S (m ³ s ⁻¹)	3.760×10^6	T_0 (°C)	3.26
K_{IP} (m ³ s ⁻¹)	89.778×10^6		

Table II. Additive noise is included to yield stochastic differential equations of the form

$$dS_X = f_X(S_X, H)dt + \sigma_X dW_X, \quad (\text{B3})$$

with $X \in \{N, T, S, IP\}$, $\sigma_X = 10^{-6}$. The drift f_X represents the deterministic model (B2) and W_X are standard independent Wiener processes.

- [1] J. Ladyman, J. Lambert, and K. Wiesner, What is a complex system?, *Eur. J. Philos. Sci.* **3**, 33 (2013).
- [2] C. Kuehn, A mathematical framework for critical transitions: Bifurcations, fast-slow systems and stochastic dynamics, *Physica D* **240**, 1020 (2011).
- [3] H. Haken, *Synergetics*, 2nd ed. (Springer, Berlin, Heidelberg, New York, 1980).
- [4] G. R. North, R. F. Cahalan, and J. J. A. Coakley, Energy balance climate models, *Rev. Geophys.* **19**, 91 (1981).
- [5] C. Wissel, A universal law of the characteristic return time near thresholds, *Oecologia* **65**, 101 (1984).
- [6] T. J. Crowley and G. R. North, Abrupt climate change and extinction events in earth history, *Science* **240**, 996 (1988).
- [7] T. Kleinen, H. Held, and G. Petschel-Held, The potential role of spectral properties in detecting thresholds in the earth system: Application to the thermohaline circulation, *Ocean Dyn.* **53**, 53 (2003).
- [8] H. Held and T. Kleinen, Detection of climate system bifurcations by degenerate fingerprinting, *Geophys. Res. Lett.* **31**, 2004GL020972 (2004).
- [9] A. Morr and N. Boers, Detection of approaching critical transitions in natural systems driven by red noise, *Phys. Rev. X* **14**, 021037 (2024).
- [10] E. Knobloch and K. A. Wiesenfeld, Bifurcations in fluctuating systems: The center-manifold approach, *J. Stat. Phys.* **33**, 611 (1983).
- [11] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara, Early-warning signals for critical transitions, *Nature (London)* **461**, 53 (2009).
- [12] R. Kubo, The fluctuation-dissipation theorem, *Rep. Prog. Phys.* **29**, 255 (1966).
- [13] M. Hairer and A. J. Majda, A simple framework to justify linear response theory, *Nonlinearity* **23**, 909 (2010).
- [14] R. Graham, A. Hamm, and T. Tél, Nonequilibrium potentials for dynamical systems with fractal attractors or repellers, *Phys. Rev. Lett.* **66**, 3089 (1991).
- [15] J. Lohmann, A. B. Hansen, A. Lovo, R. Chapman, F. Bouchet, and V. Lucarini, The role of edge states for early-warning of tipping points, *Proc. R. Soc. A* **481**, 20240753 (2025).
- [16] C. Diks, C. Hommes, and J. Wang, Critical slowing down as an early warning signal for financial crises? *Empir. Econ.* **57**, 1201 (2019).
- [17] I. A. van de Leemput, M. Wichers *et al.*, Critical slowing down as early warning for the onset and termination of depression, *Proc. Natl. Acad. Sci. USA* **111**, 87 (2014).

- [18] M. Wichers, P. C. Groot *et al.*, Critical slowing down as a personalized early warning signal for depression, *Psychother. Psychosom* **85**, 114 (2016).
- [19] C. Meisel, A. Klaus, C. Kuehn, and D. Plenz, Critical slowing down governs the transition to neuron spiking, *PLoS Comput. Biol.* **11**, e1004097 (2015).
- [20] N. Boers and M. Rypdal, Critical slowing down suggests that the western Greenland ice sheet is close to a tipping point, *Proc. Natl. Acad. Sci. USA* **118**, e2024192118 (2021).
- [21] C. Boulton, T. M. Lenton, and N. Boers, Pronounced loss of Amazon rainforest resilience since the early 2000s, *Nat. Clim. Change* **12**, 271 (2022).
- [22] N. Boers, Observation-based early-warning signals for a collapse of the Atlantic meridional overturning circulation, *Nat. Clim. Change* **11**, 680 (2021).
- [23] S. L. L. Michel, D. Swingedouw, P. Ortega, G. Gastineau, J. Mignot, G. McCarthy, and M. Khodri, Early warning signal for a tipping point suggested by a millennial Atlantic multidecadal variability reconstruction, *Nat. Commun.* **13**, 5176 (2022).
- [24] P. Ditlevsen and S. Ditlevsen, Warning of a forthcoming collapse of the Atlantic meridional overturning circulation, *Nat. Commun.* **14**, 4254 (2023).
- [25] M. C. Boerlijst, T. Oudman, and A. M. de Roos, Catastrophic collapse can occur without early warning: Examples of silent catastrophes in structured ecological models, *PLOS One* **8**, e62033 (2013).
- [26] C. Kuehn, A mathematical framework for critical transitions: Normal forms, variance and applications, *J. Nonlinear Sci.* **23**, 457 (2013).
- [27] A. Morr, N. Boers, and P. Ashwin, Internal noise interference to warnings of tipping points in generic multidimensional dynamical systems, *SIAM J. Appl. Dyn. Syst.* **23**, 2793 (2024).
- [28] P. Ashwin, S. Wieczorek, R. Vitolo, and P. Cox, Tipping points in open systems: Bifurcation, noise-induced and rate-dependent examples in the climate system, *Philos. Trans. R. Soc. A* **370**, 1166 (2012).
- [29] V. Lucarini and Mickael D. Chekroun, Detecting and attributing change in climate and complex systems: Foundations, Green's functions, and nonlinear fingerprints, *Phys. Rev. Lett.* **133**, 244201 (2024).
- [30] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps, *Proc. Natl. Acad. Sci. USA* **102**, 7426 (2005).
- [31] R. R. Coifman and S. Lafon, Diffusion maps, *Appl. Comput. Harmon. Anal.* **21**, 5 (2006).
- [32] B. Nadler, R. R. Coifman, S. Lafon, and I. G. Kevrekidis, Diffusion maps, spectral clustering and reaction coordinates of dynamical systems, *Appl. Comput. Harmon. Anal.* **21**, 113 (2006).
- [33] A. Bittracher, P. Koltai, S. Klus, R. Banisch, M. Dellnitz, and C. Schütte, Transition manifolds of complex metastable systems, *J. Nonlinear Sci.* **28**, 471 (2018).
- [34] M. Lücke, S. Winkelmann, J. Heitzig, N. Molkenhuth, and P. Koltai, Learning interpretable collective variables for spreading processes on networks, *Phys. Rev. E* **109**, L022301 (2024).
- [35] G. Margazoglou, T. Grafke, A. Laio, and V. Lucarini, Dynamical landscape and multistability of a climate model, *Proc. R. Soc. A* **477**, 20210019 (2021).
- [36] R. Banisch, Z. Trstanova, A. Bittracher, K. S., and P. Koltai, Diffusion maps tailored to arbitrary non-degenerate Itô processes, *Appl. Comput. Harmon. Anal.* **48**, 242 (2020).
- [37] K.-J. Engel and R. Nagel, *A Short Course on Operator Semigroups*, Universitext (Springer, New York, 2006), pp. x+247.
- [38] M. Hairer, On Malliavin's proof of Hörmander's theorem, *Bull. Sci. Math.* **135**, 650 (2011).
- [39] M. W. G. Metafun and D. Pallara, Compactness properties of feller semigroups, *Studia Math.* **153**, 179 (2002).
- [40] K. Hasselmann, Stochastic climate models. Part 1: Theory, *Tellus* **28**, 473 (1976).
- [41] G. Gottwald, D. Crommelin, and C. Franzke, Stochastic climate theory, in *Nonlinear and Stochastic Climate Dynamics*, edited by C. L. E. Franzke and T. J. O'Kane (Cambridge University Press, Cambridge, UK, 2017), pp. 209–240.
- [42] R. R. Coifman, I. G. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler, Diffusion maps, reduction coordinates, and low dimensional representations of stochastic systems, *Multiscale Model. Simul.* **7**, 842 (2008).
- [43] D. Proverbio, A. Skupin, and J. Goncalves, Systematic analysis and optimization of early warning signals for critical transitions using distribution data, *iScience* **26**, 107156 (2023).
- [44] A. Morr, K. Riechers, L. R. Gorjão, and N. Boers, Anticipating critical transitions in multidimensional systems driven by time- and state-dependent noise, *Phys. Rev. Res.* **6**, 033251 (2024).
- [45] A. Singer, From graph to manifold Laplacian: The convergence rate, *Appl. Comput. Harmonic Anal.* **21**, 128 (2006).
- [46] T. Berry and J. Harlim, Variable bandwidth diffusion kernels, *Appl. Comput. Harmonic Anal.* **40**, 68 (2016).
- [47] J. S. Chang and G. Cooper, A practical difference scheme for Fokker-Planck equation, *J. Comput. Phys.* **6**, 1 (1970).
- [48] L. Kikuchi, R. Singh, M. E. Cates, and R. Adhikari, Ritz method for transition paths and quasipotentials of rare diffusive events, *Phys. Rev. Res.* **2**, 033208 (2020).
- [49] R. Wood *et al.*, Observable, low-order dynamical controls of thresholds of the Atlantic meridional overturning circulation, *Clim. Dyn.* **53**, 6815 (2019).
- [50] D. Häfner, R. L. Jacobsen, C. Eden, M. R. B. Kristensen, M. Jochum, R. Nuterman, and B. Vinter, Veros v0.1—A fast and versatile ocean simulator in pure python, *Geosci. Model Dev.* **11**, 3299 (2018).
- [51] J. Lohmann and P. D. Ditlevsen, Risk of tipping the overturning circulation due to increasing rates of ice melt, *Proc. Natl. Acad. Sci. USA* **118**, e2017989118 (2021).
- [52] J. Lohmann, H. A. Dijkstra, M. Jochum, V. Lucarini, and P. D. Ditlevsen, Multistability and intermediate tipping of the Atlantic Ocean circulation, *Sci. Adv.* **10**, eadi4253 (2024).
- [53] J. Lohmann and V. Lucarini, Melancholia states of the Atlantic meridional overturning circulation, *Phys. Rev. Fluids* **9**, 123801 (2024).
- [54] W. Huisinga, S. Meyn, and C. Schütte, Phase transitions and metastability in Markovian and molecular systems, *Ann. Appl. Probab.* **14**, 419 (2004).
- [55] S. Bathiany, M. Claussen, and K. Fraedrich, Detecting hotspots of atmosphere–vegetation interaction via slowing down—Part 1: A stochastic approach, *Earth Syst. Dyn.* **4**, 63 (2013).
- [56] F. Kwasniok, Detecting, anticipating, and predicting critical transitions in spatially extended systems, *Chaos* **28**, 033614 (2018).

- [57] J. Prettyman, T. Kuna, and V. Livina, Generalized early warning signals in multivariate and gridded data with an application to tropical cyclones, *Chaos* **29**, 073105 (2019).
- [58] W. Weijer, W. Cheng, S. Drijfhout, A. V. Fedorov, A. Hu, L. C. Jackson, W. Liu, E. L. McDonagh, J. V. Mecking, and J. Zhang, Stability of the Atlantic meridional overturning circulation: A review and synthesis, *JGR Oceans* **124**, 5336 (2019).
- [59] M. S. Williamson and T. M. Lenton, Detection of bifurcations in noisy coupled systems from multiple time series, *Chaos* **25**, 036407 (2015).
- [60] A. Tantet, V. Lucarini, F. Lunkeit, and H. A. Dijkstra, Crisis of the chaotic attractor of a climate model: A transfer operator approach, *Nonlinearity* **31**, 2221 (2018).
- [61] M. Chekroun, A. Tantet, H. A. Dijkstra, and J. D. Neelin, Ruelle-Pollicott resonances of stochastic systems in reduced state space. Part I: Theory, *J. Stat. Phys.* **179**, 1366 (2020).
- [62] A. Tantet, M. Chekroun, H. A. Dijkstra, and J. D. Neelin, Ruelle-Pollicott resonances of stochastic systems in reduced state space. Part II: Stochastic Hopf bifurcation, *J. Stat. Phys.* **179**, 1403 (2020).
- [63] M. S. Gutiérrez and V. Lucarini, On some aspects of the response to stochastic and deterministic forcings, *J. Phys. A: Math. Theor.* **55**, 425002 (2022).
- [64] N. Zagli, V. Lucarini, and G. A. Pavliotis, Response theory identifies reaction coordinates and explains critical phenomena in noisy interacting systems, *J. Phys. A: Math. Theor.* **57**, 325004 (2024).
- [65] J. Lohmann, Direct test for critical slowing down before Dansgaard-Oeschger events via the volcanic climate response, [arXiv:2510.21539](https://arxiv.org/abs/2510.21539).
- [66] T. Berry and T. Sauer, Local kernels and the geometric structure of data, *Appl. Comput. Harmon. Anal.* **40**, 439 (2016).
- [67] J. Lohmann, Ocean model simulation data underlying the study “Choosing observables that capture critical slowing down before tipping points: A Fokker-Planck operator approach”, Version 1.0, Repository Electronic Research Data Archive repository of the University of Copenhagen (ERDA), Zenodo (2025), doi: [10.5281/zenodo.17593820](https://doi.org/10.5281/zenodo.17593820).