Collective variables between large-scale states in turbulent convection

Priyanka Maity [1], Andreas Bittracher [2], Péter Koltai [3,2] and Jörg Schumacher [1,4]

[1] *Department of Mechanical Engineering, Technische Universität Ilmenau, Postfach 100565, D-98684 Ilmenau, Germany*

[2] *Department of Mathematics, Freie Universität Berlin, Arnimallee 6, D-14195 Berlin, Germany*

[3] *Department of Mathematics, University of Bayreuth, Universitätsstraße 30, D-95440 Bayreuth, Germany*

[4] *Tandon School of Engineering, New York University, New York, New York 11201, USA*

(Received 5 April 2023; accepted 1 July 2023; published 28 July 2023)

The dynamics in a confined turbulent convection flow is dominated by multiple long-lived macroscopic circulation states that are visited subsequently by the system in a Markov-type hopping process. In the present work, we analyze the short transition paths between these subsequent macroscopic system states by a data-driven learning algorithm that extracts the low-dimensional transition manifold and the related new coordinates, which we term collective variables, in the state space of the complex turbulent flow. We therefore transfer and extend concepts for conformation transitions in stochastic microscopic systems, such as in the dynamics of macromolecules, to a deterministic macroscopic flow. Our analysis is based on long-term direct numerical simulation trajectories of turbulent convection in a closed cubic cell at a Prandtl number $\mathrm{Pr} = 0.7$ and Rayleigh numbers $\mathrm{Ra} = 10^6$ and 10^7 for a time lag of 10^5 convective free-fall time units. The simulations resolve vortices and plumes of all physically relevant scales, resulting in a state space spanned by more than 3.5 million degrees of freedom. The transition dynamics between the large-scale circulation states can be captured by the transition manifold analysis with only two collective variables, which implies a reduction of the data dimension by a factor of more than a million. Our method demonstrates that cessations and subsequent reversals of the large-scale flow are unlikely in the present setup, and thus it paves the way for the development of efficient reduced-order models of the macroscopic complex nonlinear dynamical system.

DOI: 10.1103/PhysRevResearch.5.033061

I. INTRODUCTION

Complex nonlinear systems typically incorporate orders of magnitude of relevant dynamical scales. Examples start at the microscopic *stochastic* level where protein macromolecules, which remain in a certain conformation for milliseconds, switch within nanoseconds into a different configuration that leads to a significant change in their chemical functionality [1,2]. All the way up to the macroscopic *deterministic* level, turbulent flows in confined geometries or extended layers can exhibit differently ordered large-scale spatial patterns that are visited for longer transients in a long-term evolution [3–6]. The (rapid) crossover from one configuration to another is triggered by fluctuations of secondary flow structures, smaller eddies, shear layers, or plumes that can affect the turbulent transport of heat or momentum [7–13]. The state or phase space of macroscopic flows is infinite-dimensional or at least extremely high-dimensional and requires drastic dimensionality reductions to model the observed large-scale dynamics effectively [14].

The present study transfers an unsupervised data reduction strategy, the transition manifold framework, from stochastic molecular dynamics [15–17] to deterministic turbulent flows. More precisely, we will analyze a turbulent thermal convection flow, also denoted as Rayleigh-Bénard convection flow, which is driven by buoyancy forces and is confined in a closed cubic cell [18–20]. Despite the simple dynamics and physical-space geometry, this configuration serves as a paradigm for many applications in nature and technology. The present turbulent flow has several similarities with the mentioned microscopic example. It appears in different conformations, here different large-scale circulation (LSC) states [21–26] that occupy different regions of the phase space for many convective time units before rapidly switching from one to another in phase space. On the one hand, one could think of these LSC states to correspond in phase space to states marked by strong similarity in their velocity and temperature fields, hence they can be thought to build concentrated clusters. On the other hand, the transitory dynamics between the LSC states may be arbitrarily complicated, providing a—if not *the*—major obstacle to reduced modeling of this system. It is exactly the transitory dynamics, more precisely the progress between the different LSC states, that we will target here. In particular, we will provide data-based evidence that it is low-dimensional.

The LSC in a closed cubic cell appears in the form of four diagonal discrete circulation roll states that fill the whole cell aligned along the diagonals [27–31]. It was shown recently that this hopping from one long-lived LSC state to another, which proceeds via short-lived LSC states, can be approximately described as a Markov process based on an analysis of long-lived LSC lifetimes and transition probabilities [32].

The analysis of the fast transition events of the long-term trajectory will give a minimal set of new coordinates spanning a low-dimensional *surrogate state space* that is sufficient to represent the statistical dynamical behavior of the system. The description by these new coordinates is then connected to the dynamical processes in the turbulent convection flow in the physical space, namely the intermittent interplay of local corner vortices next to the LSC which grow transiently and kick the diagonal circulation roll into the next macrostate. The Koopman eigenfunction framework in [30] was able to isolate the four LSC states. Secondary and tertiary eigenfunctions could be connected with these transition processes, but still gave a complex three-dimensional structure. Our approach reduces exactly this complexity. It combines and extends two established frameworks from fluid mechanics and data-driven analysis and thus opens the door to an efficient reduced dynamical model of the LSC dynamics in this particular high-dimensional application, namely as a Markov process on the surrogate state space described by the new coordinates. Apart from this combination and the resulting insights, our main contribution is the extension of a coordinate-learning framework [16] from stochastic systems to deterministic dynamics.

These new coordinates are often termed in the original chemical context as reaction coordinates [16], we will use the notion of *collective variables* for the present application. The turbulence data stem from a long-term simulation trajectory of 10^5 convective free-fall time units, t_f, at two different Rayleigh numbers $\mathrm{Ra} = 10^6, 10^7$ and a Prandtl number $\mathrm{Pr} = 0.7$; see the next section for the Rayleigh-Bénard flow model of thermal convection and exact definitions of Ra and Pr. Given the spectral resolution in the simulations for the four turbulence fields involved, we count more than 3.5 million degrees of freedom that describe the turbulent convection flows.

Recent years have witnessed a large bloom in designing methods that use (deep) neuronal networks to find surrogate dynamical models and in particular low-dimensional variables in terms of which these models are expressed [33–38]. Their success relies strongly on the ability of such networks to represent coordinates from a large general class. In these approaches, however, the dynamical conditions necessary for them to perform well remain implicit. The methodology presented here relies on very explicit dynamical assumptions, which are validated over the course of the data-driven computation.

Section II discusses the convection flow model, the numerical simulations, and basic flow statistics. Section III is dedicated to the large-scale circulation, in particular to the transition from one macrostate to another. In Sec. IV, the transition manifold framework is laid out including the preprocessing steps. Our results follow in Sec. V. A summary and discussion can be found in Sec. VI, including a brief summary of the data processing pipeline. The Appendixes provide further technical details and mathematical background.

II. TURBULENT RAYLEIGH-BÉNARD FLOW

A. Model equations and simulation method

We simulate the dimensionless three-dimensional incompressible Boussinesq equations of motion representing the Rayleigh-Bénard convection dynamics. They are given by

$$\partial_t \mathbf{u} + (\mathbf{u}\cdot\boldsymbol{\nabla})\mathbf{u} = -\boldsymbol{\nabla} p + \sqrt{\frac{\mathrm{Pr}}{\mathrm{Ra}}}\nabla^2\mathbf{u} + T\hat{\mathbf{z}}, \tag{1}$$

$$\partial_t T + (\mathbf{u}\cdot\boldsymbol{\nabla})T = \frac{1}{\sqrt{\mathrm{RaPr}}}\nabla^2 T, \tag{2}$$

$$\boldsymbol{\nabla}\cdot\mathbf{u} = \mathbf{0}. \tag{3}$$

Here, $\mathbf{u}(\mathbf{r},t)$ is the velocity field of the fluid, $T(\mathbf{r},t)$ is the temperature field, $p(\mathbf{r},t)$ is the pressure field, and $\hat{\mathbf{z}}$ is the unit vector in vertical direction z. The spatial position in the flow volume is given by $\mathbf{r} = (x,y,z) \in V \subset \mathbb{R}^3$. The dimensionless control parameters of the flow are the Rayleigh number Ra, which gives a measure of the strength of driving by buoyancy forces, and the Prandtl number Pr, which is the ratio of momentum to thermal diffusivity. Both numbers are given by

$$\mathrm{Ra} = \frac{\alpha g\,\delta T d^3}{\nu\kappa} \quad \text{and} \quad \mathrm{Pr} = \frac{\nu}{\kappa}, \tag{4}$$

where α represents the isobaric thermal expansion coefficient, g is the acceleration due to gravity, $\delta T = T_{\mathrm{bottom}} - T_{\mathrm{top}}$ is the temperature difference maintained along the fluid layer of thickness d, and ν and κ are the kinematic viscosity and thermal diffusivity of the fluid, respectively. The equations have been made dimensionless by rescaling lengths by the height of the cell d (which is equal here to the two horizontal side lengths), velocities by the free-fall velocity $U_f = \sqrt{\alpha g \delta T d}$, and temperatures by the outer difference δT. The time units for nondimensionalization results then in the free-fall timescale $t_f = \sqrt{d/(\alpha g \delta T)}$, which is the large-scale convective time unit in the present work.

We assume no-slip velocity boundary conditions at all six faces of the cubic cell $V = d^3$ with $\mathbf{u} = \mathbf{0}$. The system is uniformly heated from below at $z = 0$ with $T = T_{\mathrm{bottom}}$ and cooled from above at $z = 1$ with $T = T_{\mathrm{top}}$. We also assume thermally insulated side faces, $\mathbf{n}\cdot\boldsymbol{\nabla}T = \mathbf{0}$, such that all supplied heat at the bottom has to pass through the fluid to the top.

The simulations were performed using the open source code Nek5000 (version 17), which is based on a spectral element method [39,40]. The simulations were performed assuming 16 spectral elements along each space direction and a Lagrangian interpolation polynomial of order 5 (for $\mathrm{Ra} = 10^6$) and 7 (for $\mathrm{Ra} = 10^7$) along each space direction and on each spectral element, which results in 884 736 and 2 097 152 collocation points, respectively. The vertical profiles of the mean kinetic energy dissipation rate were analyzed to verify that this spectral resolution is sufficiently well resolved for analyzing the large-scale circulations. The simulations were carried out for two values of Rayleigh numbers $\mathrm{Ra} = 10^6$ and 10^7 in a fluid with Prandtl number $\mathrm{Pr} = 0.7$, corresponding to an effective Reynolds number of $\mathrm{Re} = \sqrt{\mathrm{Ra/Pr}}\,u_{\mathrm{rms}}$ of 467 and 1350, respectively, with u_{rms} denoting the root-mean-square velocity; see Table I. Starting from a random initial condition, we waited for a lag of 5000 free-fall times for the system to settle to a steady state. The trajectory was further allowed to evolve for another 10^5 free-fall times, and the turbulence fields were output at every free-fall time t_f, such

TABLE I. Summary of basic parameters and turbulence quantities. Comparison of root-mean-square values of velocity ($u_{\rm rms}$), vorticity ($\omega_{\rm rms}$), and temperature ($T_{\rm rms}$) obtained for the two combinations of Rayleigh number Ra and Prandtl number Pr. Here, $\langle\cdot\rangle_{V,t}$ is a combined average over the cubic volume and time. The mean persistence times ($t_{\rm pers}$) of the long-lived LSCs and short-lived LSCs are enlisted in the last two rows in terms of convective free-fall time units t_f. They are obtained from the fit of exponential laws to the persistence time distributions.

(Pr,Ra)	$(0.7, 10^6)$	$(0.7, 10^7)$
Re	467	1350
$u_{\rm rms} = \langle u_i^2\rangle_{V,t}^{1/2}$	0.39	0.36
$\omega_{\rm rms} = \langle \omega_i^2\rangle_{V,t}^{1/2}$	2.86	3.83
$T_{\rm rms} = \langle T'^2\rangle_{V,t}^{1/2}$	0.29	0.29
$t_{\rm pers}$ (SL-LSC)	$4.2t_f$	$8.6t_f$
$t_{\rm pers}$ (LL-LSC)	$12.1t_f$	$26t_f$

that we gather $N_s = 10^5$ full flow snapshots for each case considered here.

B. Velocity and vorticity statistics

Table I lists the velocity, vorticity, and temperature fluctuations in the system for both runs. The vorticity vector field is given by $\boldsymbol{\omega} = \nabla \times \mathbf{u}$. The fluctuations are determined as a root mean square with respect to the flow volume V and time t. Our data analysis revealed that a lower switching frequency between the long-lived LSC states is observed, even though a higher level of vorticity fluctuations in the convection flow exists when comparing the two long-term runs at $\mathrm{Ra} = 10^6$ and 10^7 and $\mathrm{Pr} = 0.7$. As also seen, the root-mean-squared values of the velocity, $u_{\rm rms}$, and temperature fields, $T_{\rm rms}$, remain nearly unchanged. The root-mean-square value of the temperature, $T_{\rm rms}$, is computed using the temperature fluctuations. They are given by $T'(\mathbf{r}, t) = T(\mathbf{r}, t) - \langle T(z)\rangle_{A,t}$, representing the variation of temperature field T from the mean temperature profile $\langle T(z)\rangle_{A,t}$. The latter profile depends on the vertical coordinate z and is obtained by a combined average over horizontal cross sections $A = d^2$ and time t.

The higher magnitude of vorticity fluctuations $\omega_{\rm rms}$ at the higher Rayleigh number implies that the vortical structures are fragmented and thinner for the higher Ra, whereas we observed more coherent vortical structures for the lower one. This is confirmed by the plots of vorticity magnitude isosurfaces in Fig. 1. The absence of spatially coarse coherent vortical structures in the case of $\mathrm{Ra} = 10^7$ might explain the longer lifetime of the LSC structures in one specific macrostate. Our observation is in accordance with previous studies in cylindrical cells at the same aspect ratio [41,42]. Following Ref. [42], the dependence of the mean kinetic energy dissipation rate in turbulent convection at $\mathrm{Pr} = 0.7$ on the Rayleigh number follows the scaling law:

$$\langle\epsilon\rangle_{V,t} \sim \mathrm{Ra}^{-0.2} \tag{5}$$

with the kinetic energy dissipation rate field

$$\epsilon(\mathbf{r}, t) = \sqrt{\frac{\mathrm{Pr}}{4\mathrm{Ra}}}[\nabla\mathbf{u} + (\nabla\mathbf{u})^T]^2. \tag{6}$$

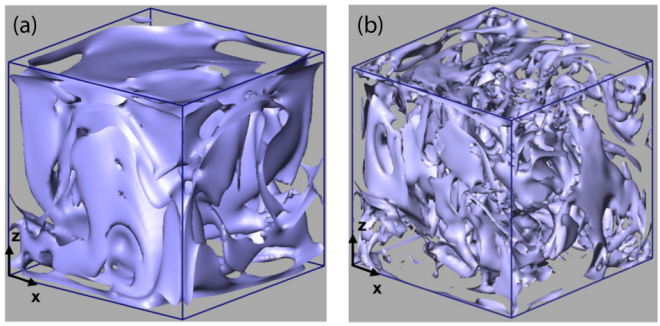

FIG. 1. Isosurfaces of vorticity magnitude $|\boldsymbol{\omega}|$ plotted for the case of fluid $\mathrm{Pr} = 0.7$ at Rayleigh numbers of (a) $\mathrm{Ra} = 10^6$ and (b) $\mathrm{Ra} = 10^7$. Displayed are two time instants of the convection flows. The isosurfaces are plotted at a value of $|\boldsymbol{\omega}| = 1.5$, which corresponds approximately to the most probable vorticity magnitude in both flows.

The mean kinetic energy dissipation rate is in turn to a very good approximation directly proportional to the mean square of the vorticity, as already found in [41,43]. This implies

$$\langle\epsilon\rangle_{V,t} \approx \sqrt{\frac{\mathrm{Pr}}{\mathrm{Ra}}}\omega_{\rm rms}^2 \sim \mathrm{Ra}^{-0.2}, \tag{7}$$

and thus a growth with Ra following $\omega_{\rm rms} \sim \mathrm{Ra}^{0.15}$. Note that this is an exact equality in homogeneous isotropic turbulence; the proportionality constant is the kinematic viscosity.

III. LARGE-SCALE CIRCULATION

A. Detection of large-scale circulation states

In the following, we discuss how the large-scale flow states are obtained from the numerical simulation data. For this analysis, the scalar and vector fields were interpolated spectrally to a uniform Cartesian mesh. The visualization and computation of the time evolution of the angle of orientation of the LSC followed that of [32]. The extraction of LSC orientation is an evolved version of that used in [9].

Our primary aim, after identification of the LSC structure, is to compute the alignment of the LSC within the box. For this purpose, we first extract the vertical velocity component on a uniformly spaced Cartesian grid at half-height. Thereafter, we perform a transformation to polar coordinates to project the vertical velocity component to a circle with a fixed radius of 0.45 and an angle θ which varies in steps of 5°. The angle θ is measured with respect to the y-axis in a clockwise manner. We then perform an azimuthal Fourier transform of the vertical velocity component data interpolated on the circle. The identification of orientation angle θ of the LSCs is conditioned on two quantities: (i) The ratio of energy carried by the largest Fourier mode (k_θ is a wave number), which is given by

$$\delta(t) = \frac{\max_{k_\theta} |\hat{u}_z(k_\theta, t)|^2}{\sum_{k_\theta} |\hat{u}_z(k_\theta, t)|^2}, \tag{8}$$

and (ii) the phase of the largest Fourier mode. During the time intervals with a pronounced LSC structure, the maximum energy will be possessed by this particular Fourier mode, and

we will have a large value of $\delta(t)$. The orientation of the LSC structure corresponds to the phase of the largest Fourier mode.

In the absence of any LSC structure, the kinetic energy will be distributed among all the modes, and we will have a low value of $\delta(t)$. The decoherent states are extremely rare and unstable. For a methodical approach of identifying LSCs, we first calculated the standard deviation σ of the probability density function (PDF) of δ. All flow states that fall below $\mu - 2\sigma$, with μ being the mean of the PDF, i.e., states that have extremely low values of δ, are identified as decoherent states. All remaining states having values greater than $\mu + 2\sigma$ are categorized as one of the four distinct LSC structures with an angle of orientation θ corresponding to the phase of the largest Fourier mode; see also [32].

We identify the six different discrete LSC states of the convection flow heuristically on the basis of the energy content contained in Fourier modes. Note that this classification relies heavily on our prior knowledge of the geometry of the physical space and the already observed LSC configurations. We observe that the system mainly prefers the long-lived large-scale circulation states aligned along the diagonals, termed as $S_{\pi/4}$, $S_{3\pi/4}$, $S_{5\pi/4}$, and $S_{7\pi/4}$. The transitions between the long-lived LSC states occur via the short-lived LSC states termed $S_{n\pi/2}$, which are aligned along the side faces, and the decoherent state S_0. We monitored the large-scale state every convective free-fall time based on collecting a total of 10^5 full three-dimensional convection flow snapshots for each of the two simulation runs. The evolution of the long-term trajectories of the LSC orientation θ for both runs is shown in Figs. 2(a) and 2(b). It is noted that the transition frequency in the case of $\mathrm{Ra} = 10^7$ is lower than for $\mathrm{Ra} = 10^6$. This has also been observed in quasi-two-dimensional flows [44]. The lower level of angular fluctuations results in the impression that there are significantly fewer data points in the bottom panel; the number of data points is $N_s = 10^5$ in both cases.

One such transition between two long-lived LSC states is detailed in Figs. 2(c)–2(j) as a typical example. The switch from the long-lived $S_{\pi/4}$ to the long-lived $S_{7\pi/4}$ via the short-lived $S_{n\pi/2}$ and the decoherent null state S_0 proceeds within a relatively short time period of 18 convective free-fall times. This is about the time it takes for a fluid parcel to circulate within an LSC roll; thus it is considered as a fast process in comparison with the total time lag. Three-dimensional streamline views from the side superimposed with isosurfaces of high-vorticity magnitudes (in gray) are shown in the middle row of the figure. We clearly observe that these intense vorticity structures are always present on both sides of the LSC mean flow; see panels (c)–(f). Their imbalance in terms of the magnitude should thus be responsible for the destabilization of the LSC roll. The significance of high-vorticity regions for large-scale flow cessations is known from quasi-two-dimensional cases [44]; the composition of corner vortices is, however, different in 3D. A visualization of vorticity magnitude contours in the bottom row of the figure in panels (g)–(j) reveals that the high-vorticity regions are a part of the curling arms of the rising plumes. The direction of the alignment of the LSC roll is determined by interactions between three-dimensional up- and down-welling plumes in the system; see, e.g., the rising plume in the center of the left front face and the falling one in the right front face of Figs. 2(d) and 2(h). In Sec. II B, we also show that the magnitude of vorticity fluctuations increases with increasing Rayleigh number and that the vortical structures become more filamented; see Fig. 1 and Table I. This example also demonstrates that a complex fully three-dimensional pattern is incorporated into the switching process, e.g., S_0. A description of this dynamics by a minimal set of coordinates will be essential to build a reduced order model of the large-scale flow behavior.

B. Transition between large-scale circulation states

Our observation suggests that the long-lived LSC states always destabilize to $S_{n\pi/2}$ (short-lived states) or S_0 (decoherent state) before converting into another long-lived LSC state. To substantiate our proposition of the destabilization mechanism of the LSC by the coherent vortices in the corners along the plane perpendicular to the direction of circulation, we compare the heat flux along the flow alignment and the vorticity magnitude triggering the destabilization. For this purpose, we choose subvolumes $V_s = (d/4)^3$, which are centered in the eight corners of the cubic simulation box. Thereafter, we calculate the vorticity centers within each of the eight V_s which are given by

$$\mathbf{r}_s(t) = \frac{\int_{V_s} |\boldsymbol{\omega}(\mathbf{r},t)|\mathbf{r}\, d\mathbf{r}}{\int_{V_s} |\boldsymbol{\omega}(\mathbf{r},t)|\, d\mathbf{r}}. \tag{9}$$

At these eight vorticity centers $\mathbf{r}_s(t)$ (which are the "center of mass" of high-vorticity regions in V_s), we center a further smaller cubic subvolume V_c consisting of 8^3 grid points to calculate the following time series:

$$\langle|\boldsymbol{\omega}|(t)\rangle_s = \frac{1}{V_c}\int_{V_c} |\boldsymbol{\omega}(\mathbf{r}+\mathbf{r}_s(t),t)|\, d\mathbf{r}. \tag{10}$$

Similarly, we proceed for the convective heat flux,

$$\langle u_z T(t)\rangle_s = \frac{1}{V_c}\int_{V_c} u_z(\mathbf{r}+\mathbf{r}_s(t),t)T(\mathbf{r}+\mathbf{r}_s(t),t)\, d\mathbf{r}. \tag{11}$$

Finally, we take an arithmetic average $\bar{(\cdot)}$ of those four $\langle|\boldsymbol{\omega}|(t)\rangle_s$ and $\langle u_z T(t)\rangle_s$, which pass through the two vertical diagonal planes in the cube. Thus, we now have two time series for each of both quantities. When the long-lived LSC is aligned along the diagonals (representing either $S_{7\pi/4}$ or $S_{3\pi/4}$ along one diagonal plane or the $S_{5\pi/4}$ or $S_{\pi/4}$ aligned along the other one), the heat flux along that diagonal plane maximizes.

In a destabilization phase of the LSC, the vortices residing in the plane perpendicular to the LSC plane become intense, and the corresponding vorticity magnitude time-series peaks. Simultaneously, the heat flux along the direction of the LSC drops. This is shown in Fig. 3 for a shorter time window of $1000t_f$. The system undergoes a transition from the $S_{5\pi/4}$ to the $S_{7\pi/4}$ state. It is crucial to note that the transition is not smooth. The system toggles between the $S_{5\pi/4}$ state and the short-lived $S_{n\pi/2}$ or decoherent state before transitioning to the $S_{7\pi/4}$ state. Figure 3 captures the temporal variation of heat flux along the LSC plane and vorticity magnitude in a plane that crosses perpendicularly in the vicinity of one such randomly picked destabilizing event. The onset of the destabilizing mechanism is marked by the drop of heat flux along the LSC alignment when the vorticity magnitude peaks. One such destabilizing event is marked in the rectangular

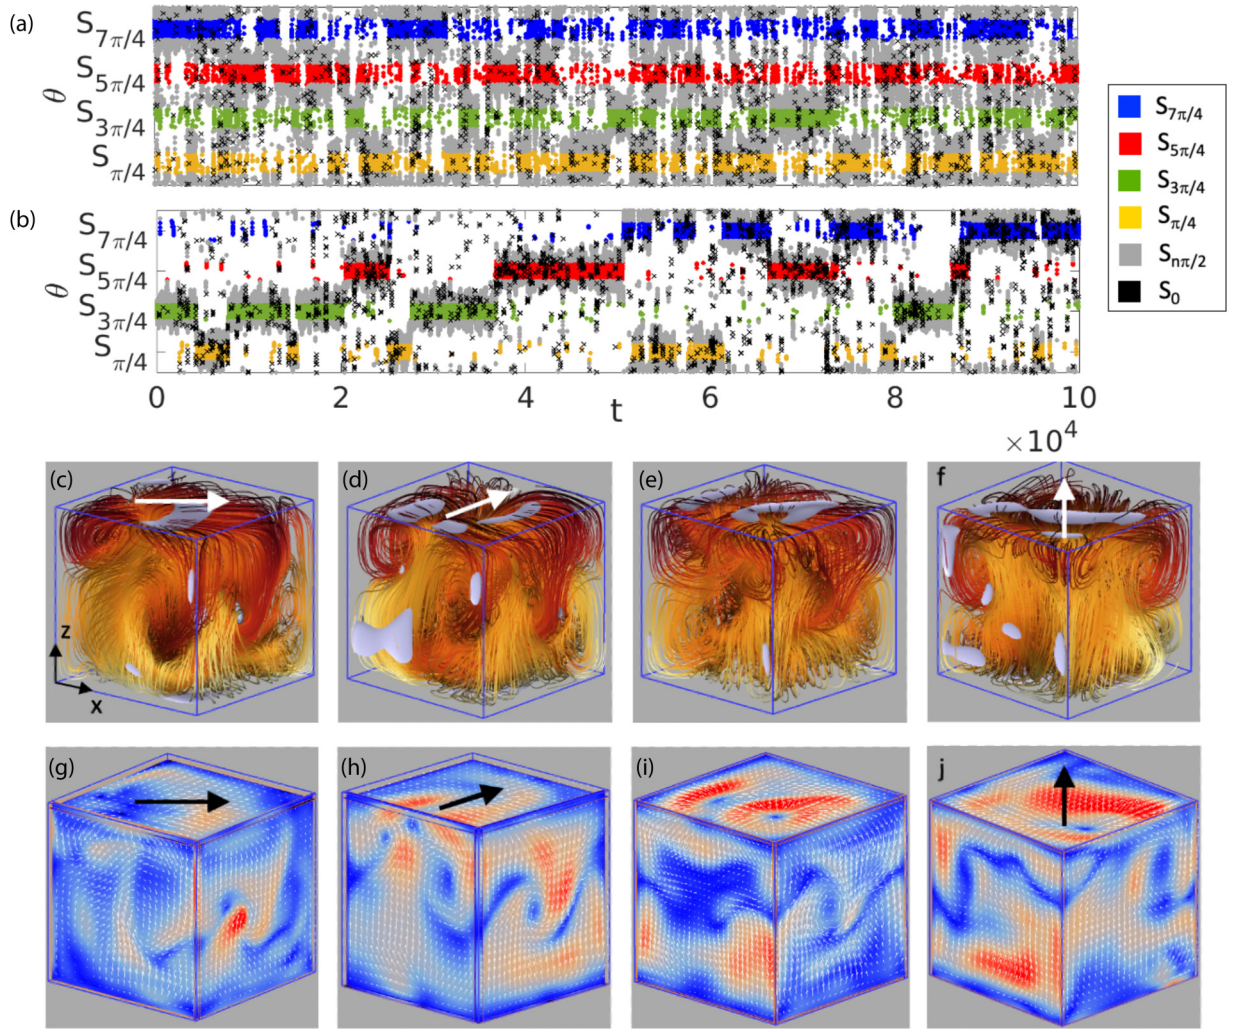

FIG. 2. Large-scale flow behavior of the convection flow in the cubical cell. Temporal evolution of orientation angle θ of the large-scale circulations (LSCs) for Prandtl number Pr = 0.7 and two different Rayleigh numbers. (a) Ra = 10^6. (b) Ra = 10^7. The legend for both panels indicates the six different LSC states. (c)–(j) Volumetric visualization of an example of a fast transition between different LSC states at Ra = 10^6. Panels (c)–(f) represent velocity streamline plots of the LSC states in the course of a transition from the long-lived LSC state $S_{\pi/4}$ in panels (c),(g) to the intermediate short-lived LSC state $S_{n\pi/2}$ state in (d),(h) and via the decoherent null state S_0 in (e),(i) to the new long-lived LSC state $S_{7\pi/4}$ state in (f),(j). The velocity streamlines are colored according to the temperature at that instant with yellow denoting the highest and black representing the lowest temperature, respectively ($0 \leqslant T \leqslant 1$). The direction of the large-scale circulation flow is indicated by white arrows, except for the case of null states. The isosurfaces of the vorticity magnitude with $|\boldsymbol{\omega}| \geqslant 7$ are added. Corresponding velocity vectors are plotted in panels (g)–(j), where blue and red contours represent the minimal, $|\boldsymbol{\omega}| = 0$, and maximal vorticity magnitudes, $|\boldsymbol{\omega}| = 8$, respectively. The direction of alignment of LSCs is represented again by black arrows.

box with a duration of $100t_f$, in Fig. 3, where the system transitions from the long-lived $S_{\pi/4}$ state to the short-lived $S_{n\pi/2}$ state. Within this time window, the Pearson correlation coefficient was obtained to be −0.71, indicating the anticorrelation between heat flux and vorticity magnitude signals. The two signals remain anticorrelated only during the onset of the destabilizing mechanism (as shown in the rectangular box). Just after the transition, the anticorrelation drops down (right-hand side of the rectangular box in Fig. 3). We computed the Pearson coefficient between the two signals for many such destabilizing events and always obtained a negative correlation coefficient (not shown). Away from the transition events, where the long-lived LSCs are stable, the Pearson coefficients are found to be either positive or slightly negative. It is also noted that the transitions are not smooth, and not all destabilization mechanisms lead to an immediate persistent transition to the next long-lived LSC state. The system undergoes multiple oscillations between different states before settling to the next long-lived LSC state (as depicted by several such destabilizing events in Fig. 3).

This establishes our proposition that the intense vortical structures sitting perpendicular to the direction of LSC circulations trigger the destabilizing mechanism, whereby the long-lived LSC destabilizes and transitions to another long-lived LSC state via the short-lived or decoherent states. The present discussion of a single transition event also underlines a dynamics with complex three-dimensional flow structures involved. A separation of these features in a reduced dynamical model is desirable. We thus continue with the data-driven framework in the next section.

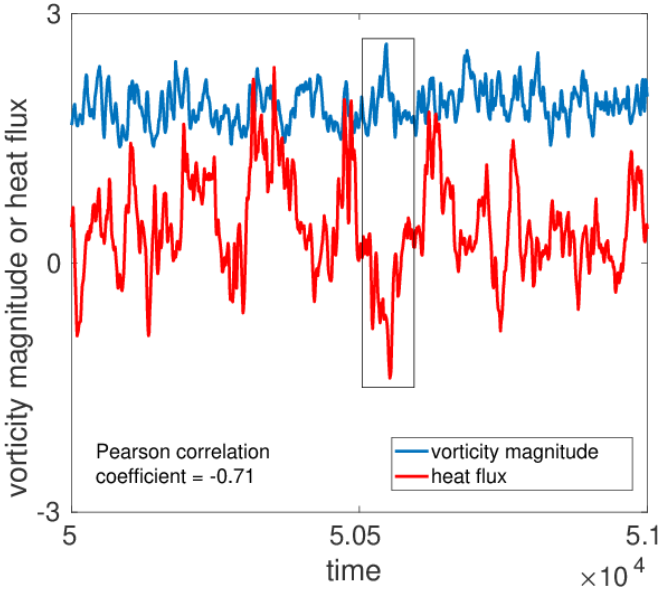

FIG. 3. Transition from LSC state $S_{5\pi/4}$ to $S_{7\pi/4}$ via a short-lived state for $\mathrm{Ra} = 10^7$ monitored by a comparison between subvolume-averaged convective heat flux $\overline{\langle u_z T(t)\rangle}_s$ computed along the direction of alignment of the LSC, and the subvolume-averaged vorticity magnitude $\overline{\langle|\boldsymbol{\omega}|(t)\rangle}_s$ computed in the opposite corner subvolumes where the intense vortical structures reside. The convective heat flux is also scaled up by a multiplicative factor of 50 for better visual comparison with the vorticity magnitude. The destabilization of long-lived LSC occurs when the vorticity peaks and the heat flux minimizes. One such destabilization event occurs in a time window of 100 free-fall times. It is one example that is highlighted by the box. The Pearson correlation coefficient between both time signals in this window is -0.71, which highlights the strong anticorrelation near the transitions.

IV. DATA REDUCTION BY TRANSITION MANIFOLD FRAMEWORK

The transition manifold framework identifies low- but possibly multidimensional collective variables, i.e., observables of the full state space that characterize essential dynamical phenomena. The present algorithm is unsupervised and data-driven, i.e., it does not require any physical knowledge about the turbulent flow. While the framework was designed originally for stochastic processes, we will demonstrate how it is applied to chaotic (here in a common, rather than a strict mathematical sense), but deterministic systems, such as the macroscopic Rayleigh-Bénard convection flow. We have already seen in [32] that the flow acts as a "memoryless" stochastic process on the largest scales, which is demonstrated by the exponential distribution of persistence times. This suggests that there is a coarse-graining level in between on which the system can be well described as a stochastic process, making it amenable to the present framework.

In the following, we will ease the notation for the data-driven framework. The variable $x(t)$ stands for a velocity field state of the turbulent Rayleigh-Bénard convection flow, $\mathbf{u}(\cdot, t)$. Thus, $x(t)$ is one trajectory point in the state space $\mathbb{X}$, the dimension of which depends on the number of spatial grid points—and is usually very high. From now on, also the variables y and z will not denote coordinates in physical space, rather points in the full state space $\mathbb{X}$ or in some latent space.

A. Original framework for stochastic systems

We now outline briefly some basic mathematical relations to keep the manuscript self-contained. Consider a time- and space-continuous, homogeneous (that is, time-invariant) Markov process $\{X_t\}_{t\geqslant 0}$ (or $\{X_t\}$ for short) on a bounded set of states $\mathbb{X} \subset \mathbb{R}^N$, with variable initial condition X_0 and an underlying probability space $(\Omega, \mathcal{F}, \mathrm{Prob})$. The dimension N will be specified further below. Assume that $\{X_t\}$ is ergodic and possesses a unique invariant density $\pi : \mathbb{X} \to \mathbb{R}^+$. We will denote by π also the probability measure induced by π, and the Lebesgue measure by Leb. Under mild assumptions on the regularity of $\{X_t\}$, the statistics of the process are characterized by the family of *transition density functions* $\{p^t\}_{t\geqslant 0} \subset L^1_{\pi\times\mathrm{Leb}}(\mathbb{X}\times\mathbb{X})$, in the sense that

$$\mathrm{Prob}[\, X_t \in A \mid X_0 = x \,] = \int_A p^t(x, y)\, dy. \tag{12}$$

In other words, $p^t(x, \cdot)$ is the distribution of X_t with starting distribution $X_0 \sim \delta_x$, with δ_x being a Dirac distribution. Written as a conditional distribution, $(X_t \mid X_0 = x) \sim p^t(x, \cdot)$. In general, $L^1_\mu(\mathbb{S})$ denotes the linear space of functions $f : \mathbb{S} \to \mathbb{C}$ integrable with respect to the measure μ, that is, for which $\|f\|_{L^1} := \int |f(x)\, d\mu(x) < \infty$.

Consider now some candidate collective variable $\xi : \mathbb{X} \to \mathbb{R}^r$, $r \ll N$, and let $\mathbb{Y} := \xi(\mathbb{X})$. As our interest in the Rayleigh–Bénard flow focuses on the most slowly evolving LSC phenomena and subprocesses of the system, we may informally call ξ a "good" collective variable if on mapping the process by ξ the long-term statistics of the system are preserved. That is, for t larger than some threshold timescale τ, the statistics of $\{\xi(X_t)\}$ will mimic that of $\{X_t\}$ in a certain sense. In the formal setting introduced above, this translates to

$$\|p^t(x, \cdot) - \bar{p}^t(\xi(x), \cdot)\|_{L^1} \leqslant \varepsilon, \quad \forall x \in \mathbb{X}, t \geqslant \tau \tag{13}$$

for some threshold lag time $\tau > 0$, some family of functions $\{\bar{p}^t(\cdot, \cdot)\}_{t\geqslant 0} \subset L^1_{\bar{\pi}\times\mathrm{Leb}}(\mathbb{Y}\times\mathbb{X})$ that are smooth in their first coordinate, and some small $\varepsilon > 0$. Here, $\bar{\pi}$ denotes the pushforward of the measure π by ξ, i.e., $\bar{\pi} = \pi \circ \xi^{-1}$. We call a collective variable that satisfies (13) *ε-consistent*.

Observe that Eq. (13) implies that the set of transition densities

$$\widetilde{\mathbb{M}} := \{p^\tau(x, \cdot) \,|\, x \in \mathbb{X}\} \subset L^1(\mathbb{X}) \tag{14}$$

accumulates ε-closely around the r-dimensional manifold,

$$\mathbb{M} := \{\bar{p}^\tau(z, \cdot) \,|\, z \in \mathbb{Y}\} \subset L^1(\mathbb{X}), \tag{15}$$

the *transition manifold*. We emphasize that the transition manifold does not live in state space, and hence one should not think of it as a low-dimensional manifold connecting certain parts of state space, as heteroclinic orbits or manifolds would do.

Inversely, one can construct ε-consistent collective variables from parametrization of the transition manifold. For this let $\mathcal{E} : \mathbb{M} \to \mathbb{R}^r$ be any parametrization of $\mathbb{M}$, i.e., a one-to-one map between $\mathbb{M}$ and its image, and let $\mathcal{Q} : \widetilde{\mathbb{M}} \to \mathbb{M}$ be

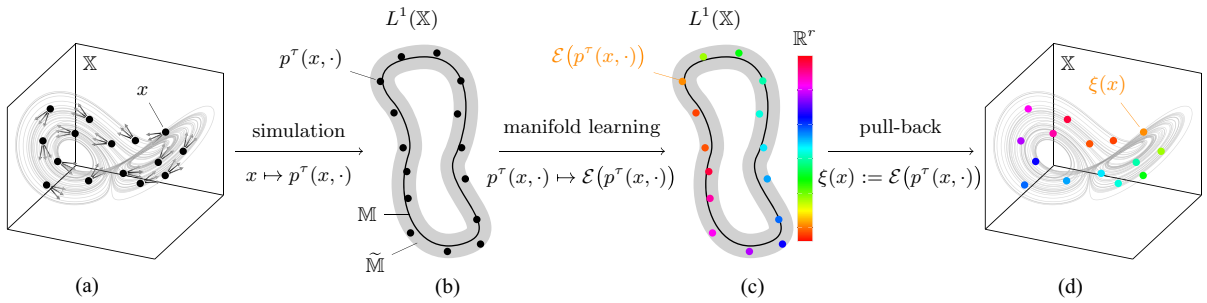

FIG. 4. Illustration of the transition manifold framework for a noisy system. (a) State space with initial states x. The multiple arrows out of the dots in the picture underline the probabilistic nature of the system. (b) Initial states are mapped by their transition densities to $\widetilde{\mathbb{M}} \subset L^1$, where they might accumulate around a low-dimensional manifold $\mathbb{M}$. (c) A low-dimensional parametrization of $\widetilde{\mathbb{M}}$ is found, indicated by the colors. (d) This parametrization is trivially pulled back to the set of initial states to define the collective variable ξ.

any map with the property

$$\|p^\tau(x,\cdot) - \mathcal{Q}(p^\tau(x,\cdot))\|_{L^1} \leqslant \varepsilon. \tag{16}$$

We think of $\mathcal{Q}$ as a map "naturally collapsing" $\widetilde{\mathbb{M}}$ to its low-dimensional approximation $\mathbb{M}$. Under the assumption of (13), such a map exists and it can be shown that the collective variable

$$\xi(x) := \mathcal{E}(\mathcal{Q}(p^\tau(x,\cdot))) \tag{17}$$

satisfies indeed (13). We call ξ the *transition manifold collective variable*. As in the case of the transition manifold, ξ does not necessarily parametrize the time-evolution of every trajectory in state space, rather it parametrizes the *progress of trajectories* during transitory dynamics in a collective manner.

As shown in [16], the process $\{\xi(X_t)\}_{t\geqslant 0}$ admits dominant timescales (decay rates of correlations) which are $O(\varepsilon)$-close to those of the original process $\{X_t\}_{t\geqslant 0}$, if this is (stochastically) reversible. This quantifies how "good" a collective variable ξ is, and it paves the way for a constructive method that finds such collective variables of the transition manifold; cf. Fig. 4. Indeed, the *dynamical* property (13), or equivalently the existence of a low-dimensional transition manifold $\mathbb{M}$, can be validated during the computation. Hence, ε-consistency of a collective variable is a well-accessible property both from theoretical and practical viewpoints. A detailed discussion on the algorithmic realization of collective-variable computation can be found in Appendix B, and in particular, the dimension estimation is discussed in Appendix B 5.

B. Set-based approximation of collective variables for deterministic systems

For a deterministic flow with a map Φ^τ, the transition "densities" $p^\tau(x,\cdot)$ are Dirac distributions $\delta_{\Phi^\tau(x)}(\cdot)$ and hence not in $L^1(\mathbb{X})$. This is a problem for our framework as introduced above. Yet, complex chaotic systems behave in many aspects as genuinely stochastic systems, and their analysis is of a probabilistic or statistical nature. This motivates our application of the transition manifold framework on deterministic chaotic systems such as the present one.

A solution to the above problem is offered by the concept of *small random perturbations* [45]. Intuitively, a family $(p_\varepsilon)_{\varepsilon>0}$ of transition densities (or measures, in general) is a small random perturbation of the map Φ if in the space of measures, $p_\varepsilon(x,\cdot) \xrightarrow{w} \delta_{\Phi(x)}$ in the weak sense as $\varepsilon \to 0$. Such small random perturbations retain certain dynamical properties of the original system in the limit of vanishing perturbation [45,46]. This is a kind of structural stability that we now assume extends to collective variables as well. The idea is to use small random perturbations with transition densities in $L^1(\mathbb{X})$ to transfer the transition manifold framework to deterministic systems.

The following discretization, that we propose to use in the collective variable computation pipeline for a deterministic flow Φ^τ, can be shown to provide for the original system a small random perturbation that has a transition density [47]. More precisely, instead of computing the transition manifold collective variable ξ at a single point, we will evaluate the average of ξ over partition elements $\{A_1,\ldots,A_L\}$ of the state space $\mathbb{X}$:

$$\xi(x) \rightsquigarrow \Xi(A_i) := \mathbb{E}[\,\xi(X) \mid X \sim \pi|_{A_i}]. \tag{18}$$

This will give us (an approximation of) ξ on the whole state space instead of just in isolated points, and at the same time allow its computation from long serial trajectory data, as opposed to simulation "bursts"; cf. Algorithm 1.

Algorithm 1. Setwise computation of the transition manifold collective variables.

Input: Data set pairs $\mathcal{X} = \{x_1,\ldots,x_{N_s}\}$, $\mathcal{Y} = \{y_i,\ldots,y_{N_s}\}$ with $y_i = \Phi^\tau x_i$.

1: Choose a partition $\{A_1,\ldots,A_L\}$ of $\mathbb{X}$ so that the metastable and transition regions are covered evenly.

2: Sort the data set $\mathcal{X}$ into the partition elements, i.e., for each j, compute the index set

$$\mathcal{I}_j := \{i | x_i \in A_j\}.$$

The set

$$\mathcal{Y}_j = \{y_i | i \in \mathcal{I}_j\}$$

is then an empirical approximation of $p^\tau(A_j,\cdot)$.

3: Approximate the distance matrix $D \in \mathbb{R}^{L\times L}$,

$$D \in \mathbb{R}^{L\times L},\ D_{ij} = d(p^\tau(A_i,\cdot), p^\tau(A_j,\cdot))$$

from the sampled $p^\tau(A,\cdot)$.

4: Apply an unsupervised, distance-based manifold learning algorithm to D.

Output: Approximation to transition manifold collective variable ξ, averaged over the sets $\{A_1,\ldots,A_L\}$, i.e.

$$\{\Xi(A_1),\ldots,\Xi(A_L)\}.$$

In our setting, the data set pairs $\mathcal{X}, \mathcal{Y}$ will be constructed from a single long trajectory, which is given by

$$\mathcal{T} := \{x_0, \Phi^\tau x_0, \Phi^{2\tau} x_0, \ldots, \Phi^{N_s \tau} x_0\}, \tag{19}$$

where Φ^τ represents the flow map of the dynamical system, and N_s is the number of snapshots, here $N_s = 10^5$. We use

$$\mathcal{X} = \{x_0, \ldots, \Phi^{N_s - 1} x_0\}, \tag{20}$$

$$\mathcal{Y} = \{\Phi^\tau x_0, \ldots, \Phi^{N_s \tau} x_0\}, \tag{21}$$

where we use the lag time $\tau = 1 t_f$. A trajectory output at this sampling rate is high enough to resolve the turnover process of fluid parcels in a LSC circulation roll.

To obtain a partition of $\mathbb{X}$, we use a distance-based method. As the results of this method can strongly depend on the distance used, especially in very high-dimensional normed spaces, we will perform the partition on the trajectory $\mathcal{T}$ mapped to $z_i = \psi(x_i) \in \mathbb{R}^{\tilde{r}}$ in a lower-dimensional space by an auxiliary observation function ψ. Here, $\tilde{r} \ll N$. The viability of this and the choice of ψ is discussed in the next section. The partition of $\mathcal{T}$ is then induced by the partition found for $\psi(\mathcal{T})$. To partition the set $z_1, \ldots, z_{N_s} \subset \mathbb{R}^{\tilde{r}}$, we use a *Voronoi tessellation* of $\mathbb{R}^{\tilde{r}}$ with center points $\{z^{(1)}, \ldots, z^{(L)}\}$ selected according to some sensible rule. A Voronoi tessellation of $\mathbb{R}^{\tilde{r}}$, with center points $\{z^{(1)}, \ldots, z^{(L)}\}$, is a collection of sets $A_1, \ldots, A_L$ with $\mathbb{R}^{\tilde{r}} = \cup_{i=1}^{L} A_i$, such that $z \in A_i \Leftrightarrow \|z - z^{(i)}\| \leqslant \|z - z^{(j)}\|$ for $j \neq i$. In other words, the sets A_i are formed by the points that are closer to $z^{(i)}$ than to any other center points $z^{(j)}$. We applied here the k-means clustering algorithm with $k = L$ to $\psi(\mathcal{T})$, and we use the resulting cluster centroids as Voronoi centers, as proposed in [48]. This has been demonstrated [17] to provide an even covering of the data points with center points. Further details are deferred to Appendix B 1.

V. COLLECTIVE VARIABLES FOR RAYLEIGH-BÉNARD FLOW

A. Preprocessing by time-lagged independent component analysis

A flow state $x \in \mathbb{X}$ of the turbulent convection flow is originally infinite-dimensional, and even after the spatial discretization that is used for the direct numerical simulation it has millions of dimensions, as discussed in Sec. II. The subsequent down-sampling from the original computational mesh still leaves us with a 12 288-dimensional data space in the present cases; see also the next subsection. Representing general distributions in so many dimensions is a challenging task, and it has proven itself useful to further preprocess the data by some coarse dimension reduction technique. Note that the assumption underlying the transition manifold framework is that the set $\widetilde{\mathbb{M}}$ in (14) is almost a low-dimensional smooth manifold. This means that almost every linear projection onto a sufficient but low-dimensional space is preserving the manifold structure of its core $\mathbb{M}$. [We note that this "projection theorem" is a useful quantitative statement, which is also used to prove Whitney's (weak) embedding theorem.]

We will consider the pushforward of the distributions $p^\tau(x, \cdot)$, $x \in \mathbb{X}$, through a linear function $\psi : \mathbb{X} \to \mathbb{R}^{\tilde{r}}$ with moderate range dimension $\tilde{r}$. More precisely, as $(X_\tau \mid X_0 = x) \sim p^\tau(x, \cdot)$, we will consider the distribution of $(\psi(X_\tau) \mid X_0 = x)$ on $\mathbb{R}^{\tilde{r}}$. In addition to the above, this has the advantage that sampling this latter distribution is more accurate with the same amount of samples than $p^\tau(x, \cdot)$ as $\tilde{r} \ll N$.

The function ψ we choose to be composed by *time-lagged independent component analysis* (TICA). In brief, TICA determines linear coordinates of the state that are ordered according to their so-called "kinetic variances"; see Appendix A for further details on this data analysis algorithm. In this way, dominant TICA coordinates are expected to capture a lot of the dynamic variability and thus suggest themselves to be good observables to resolve the transition densities. Note that they are not expected to be able to parametrize the transitory dynamics—this would indeed be a lot to ask from linear observables for a highly nonlinear (turbulent) system.

B. TICA coordinates for Rayleigh-Bénard case

Since we are interested in macroscopic and mesoscopic dynamical effects at the larger scales of the convective flow, we reduce the data amount and downsample from the original grid resolution of the DNS. The three components of the velocity vector field are therefore spectrally interpolated on a coarser uniform mesh with $16 \times 16 \times 16$ points. Hence, the Rayleigh-Bénard convection flow trajectory is embedded in a phase space with dimension $N = 16^3 \times 3 = 12\,288$.

TICA is then run on this N-dimensional trajectory $(x_i)_{i=1}^{N_s}$ of length $N_s = 10^5$, producing the *TICA coordinates* $\psi_1(x_i), \psi_2(x_i), \ldots, \psi_N(x_i)$ of the flow states $x_i = \mathbf{u}(\cdot, t_i)$. The TICA spectra in Figs. 5(a) and 5(d) exhibit a clear gap after the third eigenvalue, which is shown in panels (b) and (e) of the same figure. This gap indicates twice as much kinetic variance in each of the first three TICA coordinates as in any of the others, and thus suggests to use $\psi = (\psi_1, \psi_2, \psi_3)^\intercal$ as the observation function introduced in the previous section (this means $\tilde{r} = 3$); see panels (c) and (f) of Fig. 5. As detailed in Appendix B 5, for the transition manifold analysis there is no information gain in increasing the number of TICA coordinates, so the three leading ones are sufficient for both cases at hand.

Before proceeding to an analysis that considers the transitory dynamics of the system, let us analyze the data $\{\psi(x_i)\}_{i=1}^{N_s}$ from a static point of view, i.e., without taking information on succession of data points into account. Figure 6 shows histograms of the leading six TICA coordinates, showing a multimodal structure of ψ_1, ψ_2, ψ_3 at $\mathrm{Ra} = 10^6$. This suggests that large-scale structures of the convection flow could be identified as clusters in the data mapped into ψ-space. Indeed, plotting a subset (sampled equidistantly in time with step 50, hence 2000 points) of the trajectory in ψ-space reveals four densely populated regions as seen in Fig. 7. Applying the standard Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [49] with parameters $\varepsilon = 0.12$ and minPts = 15, we identify four clusters, i.e., four sets of trajectory index sets, that are robust to moderate parameter changes in the DBSCAN algorithm. Pulling these four index sets back into physical space, and averaging the corresponding flow fields, we see that they correspond to the four LSC states; see the streamline plots in Fig. 7 on the sides. Results for $\mathrm{Ra} = 10^7$ are similar (not shown).

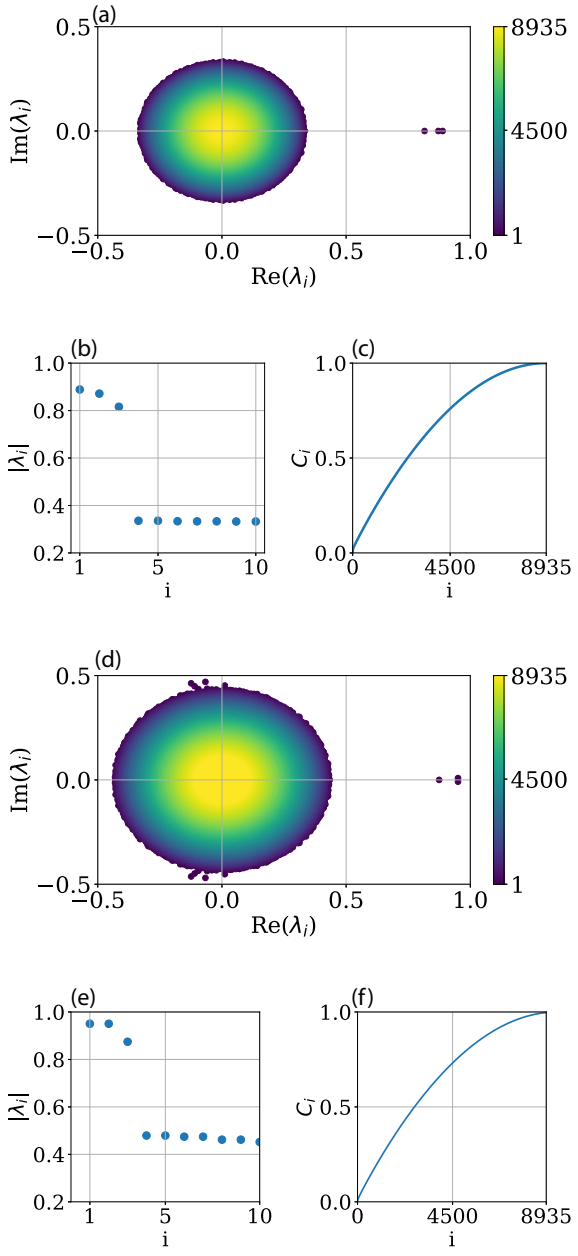

FIG. 5. (a) Real and imaginary parts of the TICA eigenvalues for the run at $\mathrm{Ra} = 10^6$ (a)–(c) and $\mathrm{Ra} = 10^7$ (d)–(f). (a),(d) The coloring corresponds to the eigenvalue index. (b),(e) Absolute value of the leading 10 TICA eigenvalues. A spectral gap is clearly visible. (c),(f) Cumulative kinetic variance c_i, indicating the portion of kinetic variance preserved by projecting onto the first i TICA coordinates.

The TICA-based analysis of the Rayleigh-Bénard flow so far was intrinsically linear. With the following transition manifold framework, we will advance to a nonlinear reduction. As mentioned above, although the transition manifold framework can be applied to arbitrary systems, it has a quantitative performance guarantee only for (stochastically) reversible systems. Rayleigh-Bénard convection is a deterministic system and hence cannot be reversible. The remedy is that both the results in [32] and the realness of the dominant TICA spectrum in Fig. 5(a) indicate reversibility of the large-scale processes.

C. Transition manifold analysis

While a projection onto the dominant three TICA coordinates clearly allows for the identification of the LSC states, it does not provide insight into the transition pathways between them. For this reason, we now apply the set-based transition manifold algorithm (Algorithm 1). It is applied to the trajectory projected onto the leading three TICA coordinates.

The first step consists of the construction of a Voronoi tessellation of $\mathbb{X}$ that, in particular, resolves the transition regions, or equivalently, the selection of points $z^{(1)}, \ldots, z^{(L)}$ that will become the centers of the Voronoi cells. For the reasons detailed in Appendix B 1, we choose the k-means clustering algorithm with a cluster count of $k = L = 2000$ for the center point selection. This results in an average number of 50 trajectory snapshots being designated to each Voronoi cell. In previous studies involving the transition manifold framework, comparable sample numbers have been demonstrated to capture the dynamical properties of interest [17,48]. The selected center points $z^{(1)}, \ldots, z^{(L)}$ slightly undersample the cluster areas, and cover the transition areas more densely, when compared to the equidistant subsampling. Note that, by the nature of k-means, the center points are in general not actual states from the data set (i.e., the trajectory) but instead the centroids of certain subsets of the data set.

Next we apply Algorithm 1 with the Voronoi cells $A_1, \ldots, A_L$ associated with the center points $z^{(1)}, \ldots, z^{(L)}$. For the distance measure between the transition densities we use maximum mean discrepancy (see Appendix B 2), and we use diffusion maps to learn the collective variable ξ from the resulting distances; see Appendix B 3. The resulting analysis (see Appendixes B 4 and B 5) shows that the associated transition manifold has at most dimension $r = 2$, hence we obtain a two-dimensional collective variable: $\xi(x_i) \in \mathbb{R}^2$. The collective variable ξ of the Voronoi centers is shown in Fig. 8(a) as black dots. The data now accumulate in the four corners of this collective-variable space, corresponding to the LSCs.

In contrast to the TICA embedding in Fig. 7, however, the clusters are much more concentrated. This indicates that, while two points x_i, x_j from the same cluster may be considerably distinct in physical or TICA-ψ-space, their *dynamical* distance, i.e., the maximum mean discrepancy between the densities $p^\tau(x_i, \cdot)$ and $p^\tau(x_j, \cdot)$, is much smaller. This allows us to better group states by their future statistics, i.e., *identify states that have the same long-term evolution*.

D. Transitory dynamics in collective variables

As close-by points on the transition manifold have similar evolution, the shortest path in collective-variable space is likely to have distinctive dynamical relevance. To validate this, we consider the shortest paths between the corners in Fig. 8(a); these are the colored curves. More precisely, for two cluster centers $\zeta_1, \zeta_2 \in \mathbb{R}^2$, we consider the interconnecting line

$$\gamma(r) := r\zeta_1 + (1-r)\zeta_2, \qquad r \in [0,1],$$

and we define the discrete pathway from ζ_1 to ζ_2 between sample points of the transition manifold as

$$\Gamma(r) := \mathrm{argmin}_{i=1,\ldots,L} \left\| \gamma(r) - \xi(z^{(i)}) \right\|.$$

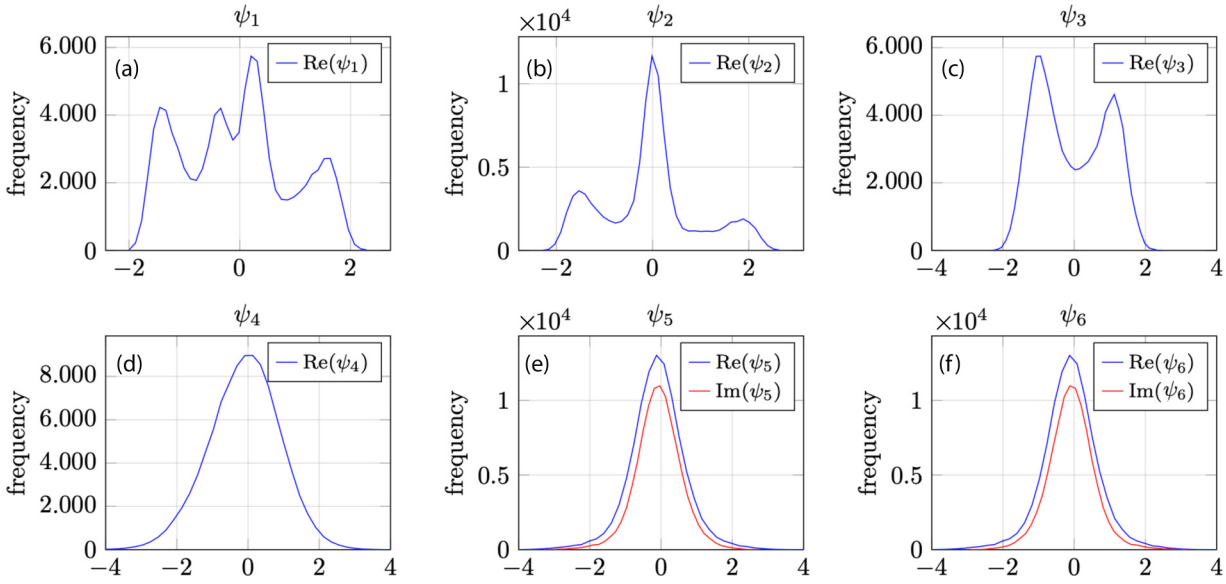

FIG. 6. Histograms of the values of the leading six TICA coordinates over the trajectory. The leading four coordinates are purely real, so only the real parts are plotted in panels (a)–(d); for the fifth and sixth leading coordinate in panels (e),(f), the real and imaginary parts are plotted. We observe a multimodal distribution in ψ_1 to ψ_3, and approximately normal distributions (real or complex) in the higher coordinates. There exist higher purely real TICA coordinates, which again are approximately normal distributed (not shown). Data are for $\mathrm{Ra} = 10^6$.

Note that these paths are not trajectories of the system. Progress along the individual pathways shows strong correlation to the midplane angle θ (recall Fig. 2), as indicated in Fig. 8(c). The imperfect correlation between θ and $\Gamma(r)$ is expected as the classification of circulation states through the midplane angle. It is based (i) on a geometric-physical intuition, and (ii) on solely time-instantaneous flow-field information, in contrast with dynamical information as in the transition manifold framework. Yet, there is a clear correlation between the results of these two classification methods. These transition pathways align well with the heuristic classification from Sec. III A. This is further underlined by the portion of the pathways spent in the respective LSCs, shown in Fig. 8(b). We observe that a pathway connecting two corners associated with two long-lived LSCs spends most of the time in these states, while transitioning through the short-lived LSC and decoherent states and notably spending barely any time in other long-lived LSCs.

Figure 9 illustrates the prevalence, i.e., the frequency of appearance in associated Voronoi cells, of the different LSC states in collective-variable space. LSCs of trajectory snapshots are again computed by the classification method used in Fig. 2. We observe that snapshots of the long-lived LSCs each concentrate in one corner of the transition manifold.

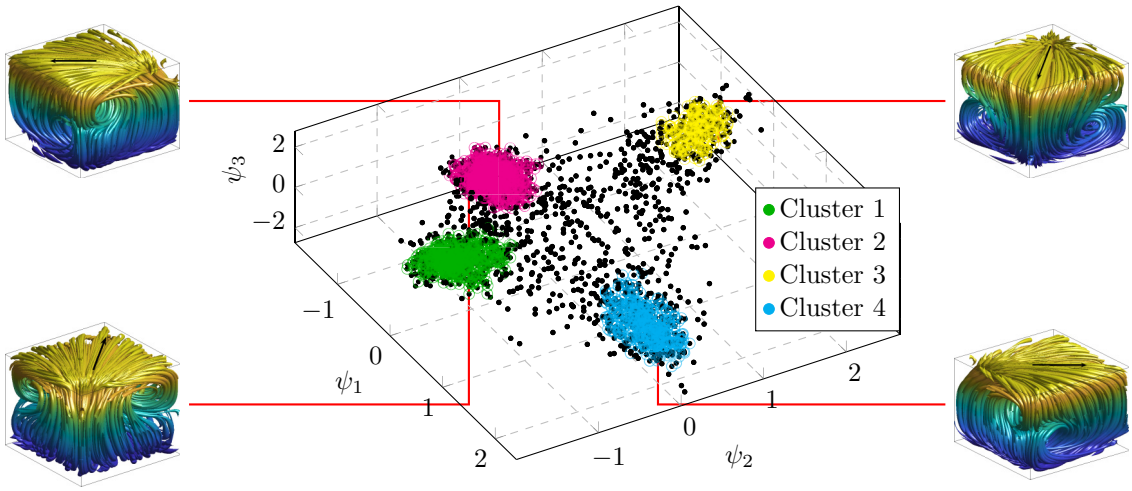

FIG. 7. Large-scale circulation states in the hyperplane which is spanned by the first three coordinates ψ_k determined by time-lagged independent component analysis (TICA). The center of the figure shows a subsample of the trajectory (2000 points) projected onto (ψ_1, ψ_2, ψ_3)-space. The points marked in color indicate the affiliation to four clusters as identified by the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm in this projection. At the sides, time-averaged velocity fields in the form of streamline plots are displayed which correspond to the four clusters. Data are for $\mathrm{Ra} = 10^6$.

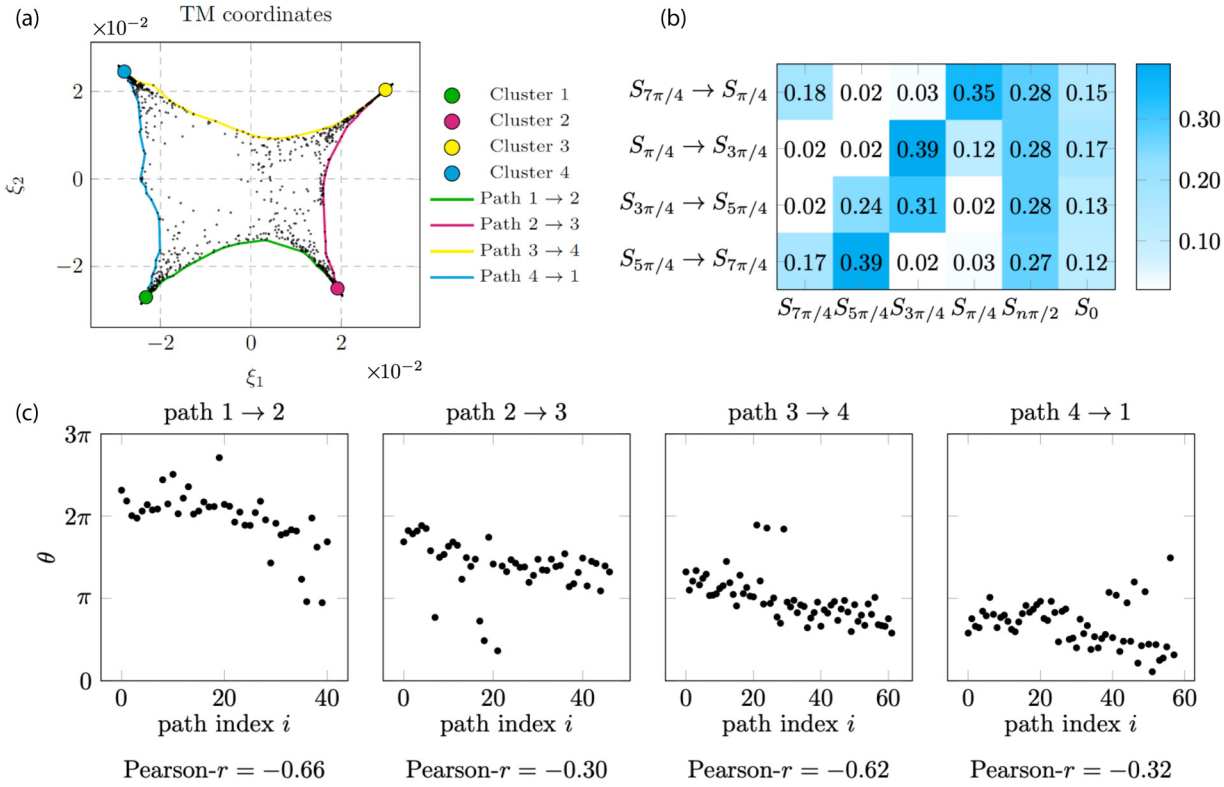

FIG. 8. Transition manifold analysis for the Rayleigh-Bénard flow. (a) Shortest pathways between the cluster centers which were determined in the transition manifold space spanned by the two-dimensional collective variable ξ. (b) Portion of time the transition path spends in or between each of the six large-scale circulation states. (c) Comparison of the progress along the individual transition pathways, listed in the form of the indices of the points along the paths, to the midplane angle θ. One visually observes a strong linear correlation, which is also confirmed by a high Pearson correlation coefficient. Note that, due to the periodic nature of θ, plotting its raw value would not convey its trend along the paths correctly. Hence, in order to obtain these results, a preprocessing step has been conducted to possibly shift each θ_i by $\pm 2\pi$, according to which value gives the closest distance to θ_{i-1}. Data are for $\mathrm{Ra} = 10^6$.

Moreover, the short-lived LSCs tend to be evenly distributed along the edges and corners of the structure, whereas decoherent snapshots have a stronger tendency to fall in the center of the transition manifold. These results demonstrate that the geometry of the transition manifold encodes and highlights the long-term stability and transition dynamics of the system.

Finally, we note that the transition manifold analysis shows also that there are no flow reversals. For this, an alternative parametrization of the transition manifold—and thus slightly distorted collective variables—was considered in Fig. 12(a) in Appendix B 5, which emphasizes small-scale geometric features. This parametrization shows no pathways connecting diagonally opposite corners of the transition manifold, which correspond to long-lived LSCs of reverse circulation direction.

We have repeated the transition manifold framework analysis for the data at $\mathrm{Ra} = 10^7$ (not shown). The pipeline of data processing is the same, resulting first in three TICA coordinates as already shown in Figs. 5(d)–5(f). The transition manifold framework gives us again four distinct clusters. The maximum mean discrepancy histogram, which is shown in Fig. 10 in Appendix B 2, suggests that the data are more strongly clustered, which can also be seen from Fig. 2(b). However, there is a difference from the case at $\mathrm{Ra} = 10^6$. Namely, the structure is not quite two-dimensional anymore. It is instead three-dimensional, or maybe even more. The third dimension, however, does not give the transition manifold any additional interpretable structure, it only seems to "thicken" the manifold, i.e., as if it were merely noise. It can be expected that with increasing Rayleigh number, the dimension of the transition manifold will increase, an aspect that we have to leave for future work.

VI. SUMMARY AND DISCUSSION

This work presented a data-driven description of the crossover dynamics from one large-scale state into another in a turbulent Rayleigh-Bénard convection flow—a paradigm for a deterministic macroscopic and complex nonlinear dynamical system. The Rayleigh-Bénard convection flow in a closed cubic cell is known to rest in macrostates in the form of large-scale circulation rolls along the diagonals for hundreds of convective time units before switching rapidly into another diagonal flow state. The specifics of these transitions were in the focus of our study. They are characterized by a complex interplay of the thermal plumes which build up the large-scale circulation and fluctuating high-amplitude vortices at the sides of the LSC, as we discussed in Secs. II and III in detail. See also Figs. 1 and 2, which displayed the transition. This dynamics is fully three-dimensional and includes the majority of the degrees of freedom of our dynamical system at hand.

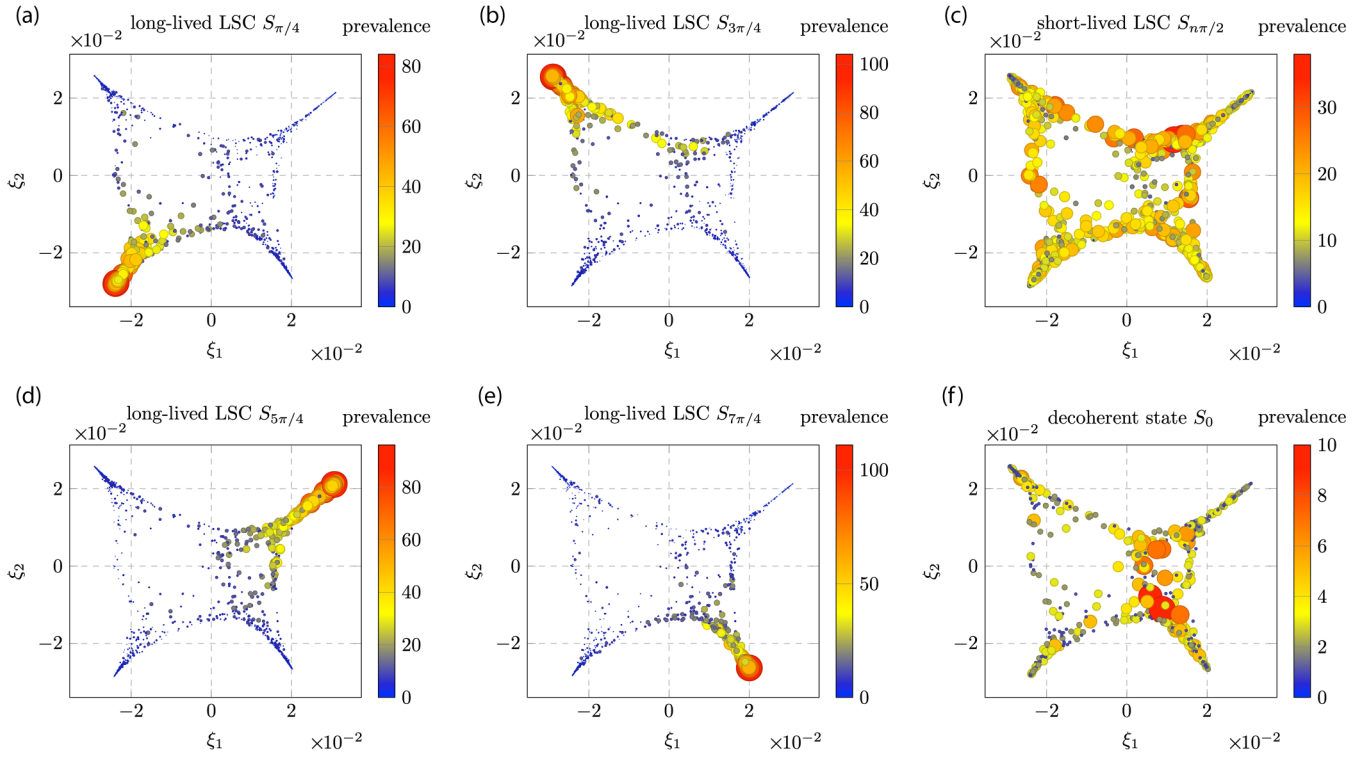

FIG. 9. Prevalence of the four long-lived (a),(b),(d),(e), the short-lived (c), and the decoherent (f) states among all trajectory snapshots for the case of $\mathrm{Ra} = 10^6$. Both the color and the size of a point indicate the number of trajectory snapshots that fall within the Voronoi cell corresponding to the point, and it belongs to one of the six states, as identified by the method used in Fig. 2 and described in detail in Sec. III A. Data are for $\mathrm{Ra} = 10^6$.

We applied the transition manifold framework to the present system, which was originally developed for stochastic microscopic systems, such as for conformational transitions of macromolecules. In our macroscopic case, a significant reduction to two coordinates, denoted as collective variables, was achieved. The initial computational grid for the four turbulent fields $(\mathbf{u}, T) = (u_x, u_y, u_z, T)$ with a total of at least $N_{\mathrm{full}} \gtrsim 3.5 \times 10^6$ degrees of freedom was first downsampled in both runs to $3 \times 16^3 = 12\,288$ degrees of freedom for $\mathbf{u} = (u_x, u_y, u_z)$. This forms the starting point of our reduced description of the transition dynamics.

The subsequent data-driven calculation of the two coordinates that span the nonlinear transition manifold was preconditioned by a time-lagged independent component analysis (TICA) that reduces the 12 288-dimensional data space. TICA determines linear coordinates which allowed us to identify LSC states as shown in Fig. 7. On the TICA-preprocessed time series, we then ran the transition manifold analysis to find that actually two coordinates describe its transition dynamics. The projection of the dynamics on these collective variables improved the distinction of the clusters significantly, which can be seen when comparing Figs. 7 and 9. The LSC states are clearly assigned with the cusps of the transition manifold. Furthermore, Fig. 9 demonstrates clearly that the short-lived LSCs are lined up along the boundary of the manifold, connecting the cusps, and that the decoherent state is found in the interior; i.e., all six macrostates can be properly separated in the plane which is spanned by the two collective variables. The new collective variables also demonstrate that flow reversals are not observed for the chosen parameter sets, thus we can conclude that they are very rare in the present three-dimensional convection case.

In future work, one could use these reduced coordinates to identify a surrogate dynamics of the large-scale processes in the system. A further interesting extension of this framework can be obtained by the study of convection flow configurations in closed cells with a somewhat larger aspect ratio. Then multiple LSC rolls will fill the simulation domain. The first step would comprise a verification of the Markov property of the large-scale flow dynamics as done for the present case in Ref. [32]. It is expected that more long-lived macrostate configurations are possible then and that in turn more than two collective variables are required to describe the transitions within the transition manifold framework.

A. Data processing pipeline summary

Several specifics of the applied methods are outlined in detail in the Appendixes, including additional figures and algorithmic elaborations. They contain technical details on the transition manifold framework and the intermediate reduction by time-lagged independent component analysis. Here, we briefly summarize the data processing pipeline of learning the collective variables once more in five subsequent steps. These are as follows:

(i) Downsampling of the grid from $N_{\mathrm{full}} \approx 3.54 \times 10^6$ and 8.39×10^6 degrees of freedom for $\mathrm{Ra} = 10^6$ and 10^7, respectively, to $N = 12\,288$ degrees of freedom. This is justified by

our focus on processes that show up on mesoscopic scales and larger.

(ii) Computation of the time-lagged independent component analysis coordinates $\psi_1, \ldots, \psi_N$ and setting of $\psi := (\psi_1, \psi_2, \ldots, \psi_{\tilde{r}})^\intercal$ with $\tilde{r} = 3$ as observables of the high-dimensional original system. This reduction follows the assumption that variations in the transition densities are low-dimensional and that density estimation is more accurate for $\tilde{r} \ll N$.

(iii) Subsequent Voronoi tessellation of $\{\psi(x_i) \mid i = 1, \ldots, N_s\}$ to find Voronoi cells $A_i \subset \mathbb{R}^{\tilde{r}}$, $i = 1, \ldots, L = 2000$. This is justified by the fact that complex turbulent flows behave like stochastic systems. Branching off from the main pipeline, we applied DBSCAN to the cell center points and associated the clusters with LSCs of the convection flow.

(iv) Approximate the ψ-pushforward of transition densities $p^\tau(x_i, \cdot)$ by the samples $\{\psi(x_{j+\tau}) \mid \psi(x_j) \in A_i\}$ and use maximum mean discrepancy to compute the distance matrix $D \in \mathbb{R}^{L\times L}$ between them. The feasibility of this step was demonstrated in [17], where the maximum mean discrepancy turns out to allow for a "quasi-embedding."

(v) Compute transition manifold by applying diffusion maps to D to obtain the estimated dimension $r = 2$ and the associated collective variable $\xi = (\xi_1, \xi_2)^\intercal$.

ACKNOWLEDGMENTS

The work of P.M. is supported by Grants No. SCHU 1410/29-1 and No. SCHU 1410/30-1 of the Deutsche Forschungsgemeinschaft (DFG). A.B. is supported by DFG through Grant No. CRC 1114 "Scaling Cascades in Complex Systems," Project No. 235221301, and Project B03 "Multilevel coarse graining of multiscale problems." P.K. has been partially supported by the DFG through Grant No. CRC 1114 "Scaling Cascades in Complex Systems," Project No. 235221301, Project A08 "Characterization and prediction of quasi-stationary atmospheric states," and under Germany's Excellence Strategy–The Berlin Mathematics Research Center MATH+ (EXC-2046/1 project ID: 390685689). The comprehensive long-term simulations were conducted at the compute cluster of the University Computing Centre at Technische Universität Ilmenau (Germany). We acknowledge support for the publication costs by the Open Access Publication Fund of the Technische Universität Ilmenau.

APPENDIX A: TIME-LAGGED INDEPENDENT COMPONENT ANALYSIS

The time-lagged independent component analysis (TICA) is applied to preprocess the simulation data. TICA provides a reduction of the data to a low-dimensional hyperplane in the state space that will be the starting point for the subsequent transition manifold analysis to determine the collective variables.

1. TICA algorithm

The evolution of the LSCs proceeds in a high-dimensional configuration space (or phase space) and consists of both slowly evolving or metastable long-lived LSC states and fast-evolving transition states, which includes short-lived LSC and decoherent states. Due to the high dimensionality of the data, it is demanding to extract the basic configuration and transition states of the system. This suggests employing dimension reduction techniques for efficacious handling of the generated data.

Perhaps the most widely used linear dimension-reduction technique is principal component analysis. While a principal component analysis finds high-variance linear combinations of the input degrees of freedom, TICA [50] is a linear transformation method based on a variational approach, which transforms a set of high-dimensional input coordinates to a set of low-dimensional output coordinates having maximal autocorrelation. It can thus be seen as a dynamic version of a principal component analysis. TICA has been developed in signal processing [50] and more recently applied in molecular dynamics [15,51,52]. Its use has two objectives, namely (i) to reduce the dimensionality of the configuration space for a faster processing of the data, and (ii) to extract the slow order parameters which can give us an idea about the slowly evolving states.

Another method frequently applied in fluid dynamics is dynamic mode decomposition (DMD) [53,54], where a linear model is fitted to the observed time series in a least squares sense. It should be noted that TICA and DMD are essentially equivalent on an algorithmic level, as shown in [55] (Sec. 3.2). However, while DMD is designed and used to visualize and forecast evolving fluid flows by a linear model, TICA is utilized to produce a set of (reduced linear) coordinates that explain most of the dynamical variability in the system. We shall thus employ the TICA-terminology in the following.

We begin with an N-dimensional dynamical system trajectory $\mathbf{x}(t) = (x_1(t), \ldots, x_N(t))^\intercal \in \mathbb{R}^N$. Here, Cartesian coordinates are taken as the three components of the velocity vector field are spectrally interpolated on a uniform $16 \times 16 \times 16$ mesh (coarser than the computational one). Hence, the convection flow trajectory is embedded in a phase space with dimension $N = 16^3 \times 3 = 12\,288$. Note that we do not include the temperature field in the analysis. The input data should obey a zero mean, thus we will use

$$\mathbf{x}'(t) = \mathbf{x}(t) - \langle \mathbf{x}(t) \rangle_t, \tag{A1}$$

where $\langle\cdot\rangle_t$ denotes time averaging. Thereafter, we compute the autocovariance matrix $C(\tau)$ at various lag times τ, which is given by

$$C_{ij}(\tau) = \langle x_i'(t) x_j'(t+\tau) \rangle_t \in \mathbb{R}^{N\times N}. \tag{A2}$$

The matrix elements of C are evaluated by

$$C_{ij}(\tau) = \frac{1}{N_s - \tau - 1} \sum_{t=1}^{N_s-\tau} x_i'(t) x_j'(t+\tau), \tag{A3}$$

where t and τ are integer-valued multiples of the output time interval, $1t_f$. Times are thus associated with integers that run from 1 to $N_s = 10^5$ in the present case. The TICA eigenvectors $\mathbf{\Psi}_k \in \mathbb{R}^N$ can be obtained as the solution of the generalized eigenvalue problem

$$C(\tau)\mathbf{\Psi}_k = \lambda_k C(0)\mathbf{\Psi}_k, \tag{A4}$$

with the kth TICA eigenvalue λ_k. The *TICA coordinates* (or TICA modes) $\psi_1(x'(i)), \psi_2(x'(i)), \ldots$ are then the coefficients

of the $x'(i)$ represented in the basis composed of the TICA eigenvectors $\mathbf{\Psi}_k$. We note that with the data matrices

$$X := \begin{bmatrix} | & & | \\ x'(1) & \cdots & x'(N_s-\tau) \\ | & & | \end{bmatrix} \quad \text{and}$$
$$Y := \begin{bmatrix} | & & | \\ x'(1+\tau) & \cdots & x'(N_s) \\ | & & | \end{bmatrix}, \tag{A5}$$

one often writes

$$C(0) = \frac{1}{N_s-\tau-1} XX^\intercal \quad \text{and}$$
$$C(\tau) = \frac{1}{N_s-\tau-1} XY^\intercal. \tag{A6}$$

2. Coordinate reduction

We compute the TICA coordinates for a lag time of $\tau = 50$ using the PyEMMA package [56]. Figures 5(a) and 5(d) shows the eigenvalues λ_i, $i \geqslant 1$ of the autocovariance matrix, which will be called *TICA eigenvalues* in the following.

It is of separate interest, but we note that TICA (and DMD) provide approximations to the composition operator associated with the dynamics, named after Koopman; see [55] and references therein. As such, the TICA spectrum provides an admittedly coarse approximation of the system's spectrum, and hence of the decay of correlations in the dynamics.

We observe that the majority of the spectrum is confined to a complex disk of radius smaller than 1. This is a typical situation for quasicompact transfer (and Koopman) operators that arise for complicated dynamics exhibiting a certain degree of chaoticity or mixing [57] (sections 1.3 and 2.3 in particular). Furthermore, the dominant (by absolute value) TICA eigenvalues are purely real, which suggests that on long timescales associated with these eigenvalues, the system behaves in a statistically reversible manner. Indeed, transitions from one LSC to other ones do not seem to have a cyclic tendency. The absolute value of a TICA eigenvalue is also known as the *kinetic variance* of its corresponding eigenvector, and it can be an indicator for how much "dynamical information" is preserved by projecting onto this coordinate [58]. The kinetic variances of the leading 10 coordinates are shown for both data sets in Figs. 5(b) and 5(e). The spectral gap after λ_3 is immediately apparent, and it hints at the disproportional dynamical significance of the first three eigenvectors.

The *cumulative kinetic variance* is now defined as the sum of the kinetic variances up to a certain TICA eigenpair, relative to the sum of all kinetic variances:

$$c_i := \frac{\sum_{k=1}^{i} |\lambda_k|}{\sum_{k=1}^{N} |\lambda_k|}. \tag{A7}$$

The course of c_i in i is shown in Figs. 5(c) and 5(f). We observe a flattening of the curve, indicating a greater impact of lower TICA coordinates (which is clear as the eigenvalues are ordered by decreasing absolute value), but not a visible jump or kink after c_3. Moreover, the curve exhibits a considerably shallower slope than what one is accustomed to in the analysis of, for example, many molecular dynamical data sets [51]. This tells us that a large portion of the kinetic variance is generated by medium- and small-scale processes and not overly concentrated in the global large-scale processes, which are well represented by the leading three TICA coordinates (see Fig. 7 in the main text). Hence, while λ_i, $i = 1, 2, 3$, may contain a disproportional amount of information, a much greater number of TICA coordinates would be required to resolve the dynamics of the system by a linear model more or less completely. Also, note that in the single precision floating point arithmetic used by PyEMMA, c_{8935} is equal to 1, hence the data set can be reduced by $N - 8935 = 3353$ dimensions without measurable loss of kinetic variance, where $N = 12\,288$ is the original system dimension.

Figure 6 visualizes the leading six TICA coordinates by the histograms of their values along our trajectory. Corresponding to the three leading purely real eigenvalues, the leading three coordinates are also purely real. Moreover, the coordinates belonging to the well-separated, dominant part of the spectrum possess a multimodal distribution, whereas the coordinates belonging to eigenvalues in the aforementioned complex disk are approximately normal-distributed.

The multimodal structure of $\psi_{1,2,3}$ raises hope that it allows the identification of large-scale structures of the Rayleigh-Bénard convection flow. Indeed, plotting a subset (sampled equidistantly in time with step 50, hence 2000 points) of the trajectory in (ψ_1, ψ_2, ψ_3)-space reveals four densely populated regions; see again Fig. 7. Applying the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [49] with parameters $\varepsilon = 0.12$ and minPts $= 15$, we identify four clusters, i.e., four sets of trajectory index sets, that are robust to moderate parameter changes in the DBSCAN algorithm. Pulling these four index sets back into physical space, and averaging the corresponding flow fields, we see that they correspond to the four LSC states (see again Fig. 7, on the sides).

APPENDIX B: ALGORITHMIC REALIZATION OF THE TRANSITION MANIFOLD FRAMEWORK FOR STOCHASTIC SYSTEMS

In practice, neither the transition manifold $\mathbb{M}$ nor the maps $\mathcal{E}, \mathcal{Q}$ are known, and $p^\tau(x, \cdot)$ is only known empirically, e.g., by starting a large number of simulations runs with different random seeds in x up to time τ. The algorithmic strategy, therefore, consists of *learning* the parametrization $\mathcal{E}$ and the projection $\mathcal{Q}$ simultaneously from a finite subsample of the set $\mathbb{M}$. A detailed description of the steps required to perform the algorithm is given in the following subsections.

1. State space sampling

The evaluation points $\{x_1, \ldots, x_L\}$—alternatively, the sets A_i in Algorithm 1, or equivalently, their center points $\{z^{(1)}, \ldots, z^{(L)}\}$ if they arise as Voronoi tessellation—should cover the dynamically relevant regions of $\mathbb{X}$, typically the metastable sets and the transition regions between them, evenly, in order for ξ to capture the relevant dynamics globally. Various strategies exist, each with different advantages and disadvantages. While *uniform sampling* of $\mathbb{X}$ may be the easiest to realize, it heavily oversamples the dynamically

irrelevant regions, i.e., those with small π-measure. This holds true especially in high dimensions, where π tends to be concentrated in a (Lebesgue) small portion of $\mathbb{X}$. Hence, a large number of samples is required in order to cover the relevant regions with sufficient granularity.

On the other hand, *sampling from the invariant density* π solves this problem, but it introduces a bias of oversampling the metastable regions (in which π is heavily concentrated), while the transition regions are neglected. To alleviate this bias, the sampling of π can undergo a second, subsampling step, through which the distribution is adjusted to yield more evenly distributed samples. Two subsampling strategies were proposed in Ref. [48], namely the *Poisson disk* and *k-means subsampling*. Each of both strategies minimizes a specific approximation error of ξ. Details on these algorithms can also be found in [48]. In the later data analysis, we will use k-means, mainly due to more readily available and robust numerical implementations.

2. Density distance matrix

The assembly of the distance matrix D serves only to prepare the data for the application of the manifold learning algorithm (see the next subsection), but it contains the crucial choice of the density metric $d : L^1(\mathbb{X}) \times L^1(\mathbb{X}) \to \mathbb{R}_{\geqslant 0}$, which determines both the numerical feasibility of the algorithm and the quality of its output. While

$$d(u, v) = \|u - v\|_{L^1}$$

may seem like the obvious choice, and the transition manifold framework is formulated with respect to this metric, it is notoriously hard to realize numerically (involving integrals over the N-dimensional space $\mathbb{X}$), and it cannot be estimated directly from samples of u, v (an approximate analytical expression of either u or v must be derived). Due to these difficulties, in [16,59], several alternative statistical distances were proposed, including the *maximum mean discrepancy* [60], which will be used in the later data analysis.

The maximum mean discrepancy depends on the choice of a positive-semidefinite kernel function $k : \mathbb{X} \times \mathbb{X} \to R^+$, and it is analytically defined as

$$d_k(u, v) := \|\mu(u) - \mu(v)\|_{\mathscr{H}(k)},$$

where $\|\cdot\|_{\mathscr{H}(k)}$ is the norm induced by the inner product. Here, μ denotes the *kernel mean embedding* [61] into $\mathscr{H}(k)$. Its concrete form, as that of $\mathscr{H}(k)$, is irrelevant, since one has the following representation of inner products in $\mathscr{H}(k)$, readily giving a formula to compute the induced norm:

$$\langle \mu(u), \mu(v) \rangle_{\mathscr{H}(k)} = \int_{\mathbb{X}} \int_{\mathbb{X}} u(x) v(y) k(x, y) d\pi(x) d\pi(y).$$

The Hilbert space $\mathscr{H}(k)$ induced by the inner product $\langle \cdot, \cdot \rangle_{\mathscr{H}(k)}$ is known as the *reproducing kernel Hilbert space* associated with k. There exists extensive literature on reproducing kernel Hilbert spaces and their functional analytical properties [62,63]. Moreover, some of the authors have shown in [59] that the map $\mathbb{M} \mapsto \mu(\mathbb{M})$ is one-to-one, with favorable distortion properties, i.e., it preserves central topological properties of $\mathbb{M}$. Importantly, the estimation of $\langle \cdot, \cdot \rangle_{\mathscr{H}(k)}$ and hence $d_k(\cdot, \cdot)$ from samples is straightforward: for

$$\langle u, v \rangle^M_{\mathscr{H}(k)} := \frac{1}{M^2} \sum_{k,\ell=1}^{M} k\big(y_i^k, y_j^\ell\big),$$

where $y_i^k \sim u$, $y_j^\ell \sim v$ holds,

$$\langle \cdot, \cdot \rangle^M_{\mathscr{H}(k)} \to \langle \cdot, \cdot \rangle_{\mathscr{H}_k} \quad (M \to \infty).$$

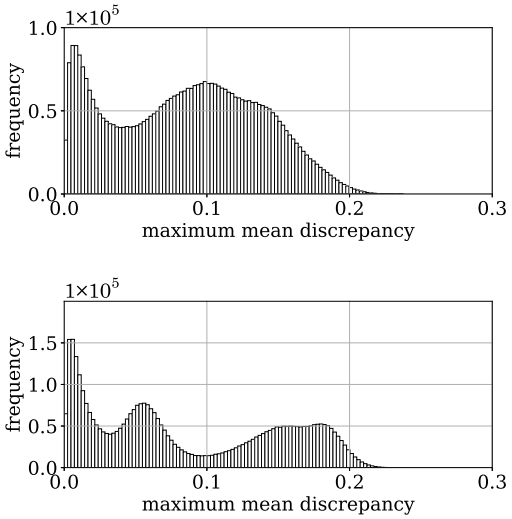

FIG. 10. Distribution of the maximum mean discrepancies between the embedded transition densities. For $\mathrm{Ra} = 10^6$, we observe two accumulations of distances, one close to zero and one around 0.1, corresponding to intra- and intercluster distances, respectively. Data are for $\mathrm{Ra} = 10^6$, (top) and 10^7 (bottom).

As described in the main text, we will compare the transition densities $p^\tau(x, \cdot)$ through their pushforwards by the observation function $\psi = (\psi_1, \psi_2, \psi_3)^\intercal$, where ψ_i denotes the ith TICA coordinate. In other words, we need to compute the maximum mean discrepancies for the densities of $\psi(X)$ and $\psi(Y)$, where $X \sim p^\tau(x, \cdot)$ and $Y \sim p^\tau(y, \cdot)$, for different pairs $x, y \in \mathbb{X}$. For this we chose the kernel $k(\xi, \eta) = \exp(-\|\xi - \eta\|^2/\sigma)$ with $\sigma = 7$ and $\xi, \eta \in \mathbb{R}^3$ (since ψ maps into $\mathbb{R}^3$). The distribution of maximum mean discrepancies for both data sets is shown in Fig. 10. The data for $\mathrm{Ra} = 10^6$ display a bimodal structure of the maximum mean discrepancy histogram, which is going to show up in the discussion about bandwidth selection for the manifold learning method as well. The data for $\mathrm{Ra} = 10^7$ display a trimodal structure. Here we find that the first smallest scale is due to the strong clustering and disconnects the data set (cf. the discussion in Appendix B 5); even the second peak in the histogram indicates length scales that do not capture global properties of the data. It is the third local maximum of the maximum mean discrepancy histogram that provides the right embedding.

3. Manifold learning

We now apply a manifold learning method to the distance matrix D. Conceptually, this corresponds to learning an approximation to the combined parametrization and projection map $\mathcal{E} \circ \mathcal{Q}$ in Eq. (17) in the main text. By pulling this map

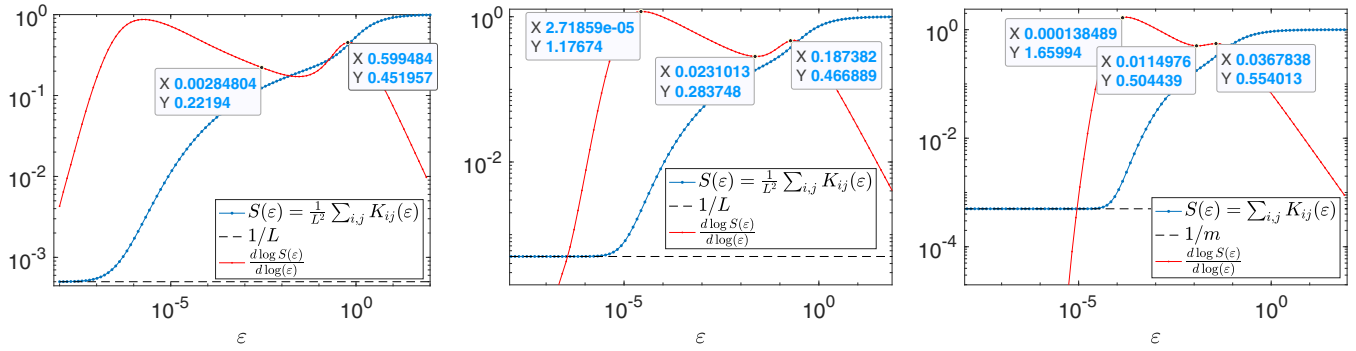

FIG. 11. The quantity $S(\varepsilon)$ and its double-logarithmic derivative for maximum mean discrepancy distance matrices obtained from the first three (left), five (middle), and ten (right) TICA coordinates. The dashed line shows $1/L$, a lower bound for $S(\varepsilon)$. The ticks in the plots show some values, for reference.

back to the starting points x_i, this yields the final collective variable ξ.

Our manifold learning method of choice is the *diffusion maps algorithm* [64], as it has demonstrated good performance in practical scenarios in the past [17,48], although other methods such as multidimensional scaling [65], local linear embeddings [66], or many others could be used as well.

We will not go into detail on the analytical derivation of diffusion maps here, but instead state only the algorithm. From the distance matrix $D \in \mathbb{R}^{L\times L}$, one first constructs a Markov transition matrix $M \in \mathbb{R}^{L\times L}$ via

$$M_{ij} = \frac{K_{ij}}{s_i}, \tag{B1}$$

where $K_{ij} = \exp(-D_{ij}^2/\varepsilon)$ is a *similarity matrix* with some bandwidth parameter $\varepsilon > 0$ [which should not be mixed with the kinetic energy dissipation rate field which is given in (6)] and $s_i = \sum_j K_{ij}$. Being a Markov matrix, the leading eigenvalue of M is 1 with corresponding constant eigenvector $\phi_1 = (1,\ldots,1)^\intercal$. The *diffusion map*, hence our collective variable ξ evaluated at the sample points $\{x_1,\ldots,x_L\}$, is now given by the following subdominant eigenvectors $\phi_2,\ldots\phi_{r+1}$. Here, the number r can typically be determined by a gap in the spectrum of M or by plotting the dominant ϕ_i against one another and discarding the ones that do not carry additional geometric information, sometimes called higher-order harmonics.

We note that diffusion maps such as those developed in [64] have a parameter α that is responsible for factoring out biases due to nonuniformity in the data sample density. The present simplified description corresponds to the case of $\alpha = 0$. However, all our implementations use $\alpha = 1$, which is the value giving no dependence on sampling density in the limit $L \to \infty$. The coordinates ϕ_i depend then on the geometric features of the underlying data manifold only.

4. Bandwidth selection in diffusion maps

Let us briefly recapitulate a method for automatically determining a "good" kernel bandwidth parameter value ε in the diffusion maps method. The procedure stems from [67], and it was later refined in [68]. A summary can also be found in [69] (Appendix A 2).

Recall that the diffusion maps approach first turns pairwise distances in the data set $\{x_i\}_{i=1}^L$ into a similarity matrix $K \in \mathbb{R}^{L\times L}$ with entries $K_{ij}(\varepsilon) = \exp(-\varepsilon^{-1}D_{ij}^2)$. Note that we use the Gaussian kernel throughout this work. By averaging the entries of the similarity matrix, we define

$$S(\varepsilon) := \frac{1}{L^2}\sum_{i,j} K_{ij}(\varepsilon). \tag{B2}$$

We note the two limiting behaviors:

(i) As $\varepsilon \to 0$, we have $K_{ij}(\varepsilon) \to \delta_{ij}$, the Kronecker delta. Thus, $S(\varepsilon) \to \frac{1}{L}$.

(ii) As $\varepsilon \to \infty$, we have $K_{ij}(\varepsilon) \to 1$, thus $S(\varepsilon) \to 1$.

According to theoretical considerations, in between these two extremes there should be a region of affine-linear growth of $\log(S(\varepsilon))$ in $\log(\varepsilon)$. This is suggested to be determined by maximizing the derivative

$$\frac{d\log(S(\varepsilon))}{d\log(\varepsilon)}, \tag{B3}$$

with respect to ε. The idea is that such an ε is neither too small (compared with the data point density) nor too large in comparison with the diameter of the data point cloud. Such bandwidth parameters can then resolve the manifold structure of the data, if there is any, and represent it sufficiently well in the diffusion maps approach. In contrast, too small bandwidths tend to essentially disconnect the data set, while too large bandwidths disregard local geometric features by "blurring" them.

We note that the procedure also gives an estimate of the dimension of the data manifold. It follows as twice the value in (B3) for bandwidths ε in the proper range. For strongly varying data density, this procedure can give multiple "optimal" bandwidths ε for different length scales, and one needs to compromise between resolution and connectivity (or otherwise break up the data in connected components and consider them one at a time). For even more irregular data, one might need to use a *variable bandwidth kernel*, as done in [68].

5. Bandwidth selection for the Rayleigh-Bénard convection data

To assess whether more TICA coordinates can capture more information about the geometric features of the transition dynamics of the convection flow, we ran the

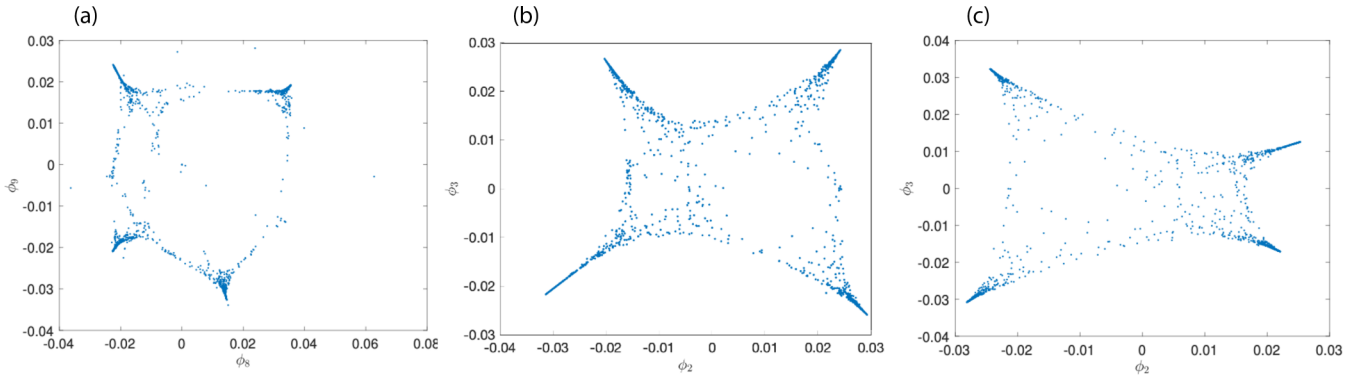

FIG. 12. Diffusion map parametrizations of the transition manifold for different bandwidths: $\varepsilon = 10^{-4}$ in (a), 3×10^{-3} in (b), and 2×10^{-2} in (c). The underlying maximum mean discrepancy distance matrix was computed from the three dominant TICA coordinates. Note that since the diffusion map coordinates are computed as eigenvectors, their sign is not unique and may vary between different computations. Hence, to interpolate between the different panels, one might need to reflect one of the axes or both around zero.

bandwidth-tuning analysis from Appendix B 4 for distance matrices obtained from maximum mean discrepancy computed from the first three, five, and ten TICA coordinates, respectively. The associated values of $S(\varepsilon)$ and the quantity (B3) are shown in Fig. 11 for $\mathrm{Ra} = 10^6$.

If the distance matrices would indicate an at most slowly varying data density, the logarithmic derivatives (red curves) would yield a single plateau, and every bandwidth ε associated with the plateau would be an equally good choice for the representation of the data. The slight bimodal structure of these curves, however, indicates that the data set is somewhat irregular, and there are multiple proximity scales highlighting different features of the data set:

(i) A smaller scale around the left-hand-side maximum resolving local variability, but potentially disconnecting the data.

(ii) A larger scale around the right-hand-side maximum that views the data set as a connected manifold.

We conducted a detailed analysis for the data set obtained for the three dominant TICA coordinates. The reasons are that (i) essentially all three cases estimate the *dimension of the data set to be at most two* (cf. Appendix B 4), and that (ii) on suitable scales ε, the parametrizations of the three data sets are essentially equivalent (not shown). This indicates that the first three TICA coordinates already capture sufficient information on the transition statistics of the system, so that the inclusion of further TICA coordinates cannot add to this (at least not for the presently available turbulent convection data). Moreover, auxiliary algorithms like k-means for the Voronoi center selection are well-known to show deteriorating performance in high dimensions.

Figure 11 left, corresponding to the distances from three TICA coordinates (ψ_1, ψ_2, ψ_3), suggests to consider bandwidths $10^{-6} < \varepsilon < 0.5$. We observed that for $10^{-6} < \varepsilon < 10^{-4}$ the data set is essentially disconnected, i.e., the parametrization is dominated by single outliers, and no structure involving multiple data points can be extracted. For $\varepsilon = 10^{-4}$, the number of outliers reduces to a handful, and the diffusion map eigenvectors ϕ_8, ϕ_9 already show some structure; see also Fig. 12(a).

Increasing the bandwidth, connectivity of the "data manifold" is achieved for $\varepsilon = 3 \times 10^{-3}$ [see Fig. 12(b)], hence this is the bandwidth parameter that we use for further analysis in the main text. Figure 12(c) shows the diffusion map for $\varepsilon = 0.02$ to show how a further increase in the bandwidth starts to blur local features and highlight global geometric aspects of the data set. This does not change if we would increase up to $\varepsilon = 0.5$ (not shown). Beyond this threshold, i.e., for very large bandwidths, the data set starts to look more and more pointlike.

[1] S. E. Jackson, How do small single-domain proteins fold? Folding Des. **3**, R81 (1998).
[2] G. D. Rose, P. J. Fleming, J. R. Banavar, and A. Maritan, A backbone-based theory of protein folding, Proc. Natl. Acad. Sci. (USA) **103**, 16623 (2006).
[3] E. A. Spiegel, Convection in stars I. Basic Boussinesq convection, Annu. Rev. Astron. Astrophys. **9**, 323 (1971).
[4] R. Markson, Atmospheric electrical detection of organized convection, Science **188**, 1171 (1975).
[5] J. Marshall and F. Schott, Open-ocean convection: Observations, theory, and models, Rev. Geophys. **37**, 1 (1999).
[6] J. Schumacher and K. R. Sreenivasan, Colloquium: Unusual dynamics of convection in the Sun, Rev. Mod. Phys. **92**, 041001 (2020).
[7] H.-D. Xi and K.-Q. Xia, Flow mode transitions in turbulent thermal convection, Phys. Fluids **20**, 055104 (2008).
[8] E. P. van der Poel, R. J. A. M. Stevens, and D. Lohse, Connecting flow structures and heat flux in turbulent Rayleigh-Bénard convection, Phys. Rev. E **84**, 045303(R) (2011).
[9] S. Weiss and G. Ahlers, Turbulent Rayleigh-Bénard convection in a cylindrical container with aspect ratio $\tau = 0.50$ and Prandtl number $\mathrm{Pr} = 4.38$, J. Fluid Mech. **676**, 5 (2011).
[10] S. G. Huisman, R. C. A. van der Veen, C. Sun, and D. Lohse, Multiple states in highly turbulent Taylor–Couette flow, Nat. Commun. **5**, 3820 (2014).
[11] Y. Bao, J. Chen, B.-F. Liu, Z.-S. She, J. Zhang, and Q. Zhou, Enhanced heat transport in partitioned thermal convection, J. Fluid Mech. **784**, R5 (2015).

[12] L. Zwirner, A. Tilgner, and O. Shishkina, Elliptical Instability and Multiple-roll Flow Modes of the Large-Scale Circulation in Confined Turbulent Rayleigh-Bénard Convection, Phys. Rev. Lett. **125**, 054502 (2020).

[13] F. Schindler, S. Eckert, T. Zürner, J. Schumacher, and T. Vogt, Collapse of Coherent Large Scale Flow in Strongly Turbulent Liquid Metal Convection, Phys. Rev. Lett. **128**, 164501 (2022).

[14] D. Ruelle, *Turbulence, Strange Attractors and Chaos* (World Scientific, Singapore, 1996).

[15] S. Schultze and H. Grubmüller, Time-lagged independent component analysis of random walks and protein dynamics, J. Chem. Theor. Comput. **17**, 5766 (2021).

[16] A. Bittracher, P. Koltai, S. Klus, R. Banisch, M. Dellnitz, and C. Schütte, Transition manifolds of complex metastable systems: Theory and data-driven computation of effective dynamics, J. Nonlin. Sci. **28**, 471 (2018).

[17] A. Bittracher, J. Moschner, B. Koksch, R. Netz, and C. Schütte, Exploring the locking stage of NFGAILS amyloid fibrillation via transition manifold analysis, Eur. Phys. J. B **94**, 195 (2021).

[18] S. Chandrasekhar, *Hydrodynamic and Hydromagnetic Stability* (Dover Publications, Mineola, NY, 1981).

[19] G. Ahlers, S. Grossmann, and D. Lohse, Heat transfer and large scale dynamics in turbulent Rayleigh–Bénard convection, Rev. Mod. Phys. **81**, 503 (2009).

[20] F. Chillá and J. Schumacher, New perspectives in turbulent Rayleigh–Bénard convection, Eur. Phys. J. E **35**, 58 (2012).

[21] J. J. Niemela, L. Skrbek, K. R. Sreenivasan, and R. J. Donelly, The wind in confined thermal convection, J. Fluid Mech. **449**, 169 (2001).

[22] K. R. Sreenivasan, A. Bershadskii, and J. J. Niemela, Mean wind and its reversal in thermal convection, Phys. Rev. E **65**, 056306 (2002).

[23] A. Parodi, J. von Hardenberg, G. Passoni, A. Provenzale, and E. A. Spiegel, Clustering of Plumes in Turbulent Convection, Phys. Rev. Lett. **92**, 194503 (2004).

[24] B. A. Puthenveettil and J. H. Arakeri, Plume structure in high-Rayleigh-number convection, J. Fluid Mech. **542**, 217 (2005).

[25] E. Brown and G. Ahlers, Rotations and cessations of the large-scale circulation in turbulent Rayleigh–Bénard convection, J. Fluid Mech. **568**, 351 (2006).

[26] H.-D. Xi and K.-Q. Xia, Cessations and reversals of the large-scale circulation in turbulent thermal convection, Phys. Rev. E **75**, 066307 (2007).

[27] N. Foroozani, J. J. Niemela, V. Armenio, and K. R. Sreenivasan, Influence of container shape on scaling of turbulent fluctuations in convection, Phys. Rev. E **90**, 063003 (2014).

[28] K. Bai, D. Ji, and E. Brown, Ability of a low-dimensional model to predict geometry-dependent dynamics of large-scale coherent structures in turbulence, Phys. Rev. E **93**, 023117 (2016).

[29] N. Foroozani, J. J. Niemela, V. Armenio, and K. R. Sreenivasan, Reorientations of the large-scale flow in turbulent convection in a cube, Phys. Rev. E **95**, 033107 (2017).

[30] D. Giannakis, A. Kolchinskaya, D. Krasnov, and J. Schumacher, Koopman analysis of the long-term evolution in a turbulent convection cell, J. Fluid Mech. **847**, 735 (2018).

[31] A. Vasiliev, P. Frick, A. Kumar, R. Stepanov, A. Sukhanovskii, and M. K. Verma, Transient flows and reorientations of large-scale convection in a cubic cell, Int. Commun. Heat Mass Transf. **108**, 104319 (2019).

[32] P. Maity, P. Koltai, and J. Schumacher, Large-scale flow in a cubic Rayleigh–Bénard cell: Long-term turbulence statistics and Markovianity of macrostate transitions, Philos. Trans. R. Soc. A **380**, 20210042 (2022).

[33] S. Brandt, F. Sittel, M. Ernst, and G. Stock, Machine learning of biomolecular reaction coordinates, J. Phys. Chem. Lett. **9**, 2144 (2018).

[34] C. Wehmeyer and F. Noé, Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics, J. Chem. Phys. **148**, 241703 (2018).

[35] B. Lusch, J. N. Kutz, and S. L. Brunton, Deep learning for universal linear embeddings of nonlinear dynamics, Nat. Commun. **9**, 4950 (2018).

[36] A. Mardt, L. Pasquali, H. Wu, and F. Noé, VAMPnets for deep learning of molecular kinetics, Nat. Commun. **9**, 5 (2018).

[37] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, Data-driven discovery of coordinates and governing equations, Proc. Natl. Acad. Sci. (USA) **116**, 22445 (2019).

[38] S. E. Otto and C. W. Rowley, Linearly recurrent autoencoder networks for learning dynamics, SIAM J. Appl. Dyn. Syst. **18**, 558 (2019).

[39] P. F. Fischer, An overlapping Schwarz method for spectral element solution of the incompressible Navier–Stokes equations, J. Comput. Phys. **133**, 84 (1997).

[40] J. D. Scheel, M. S. Emran, and J. Schumacher, Resolving the fine-scale structure in turbulent Rayleigh-Bénard convection, New J. Phys. **15**, 113063 (2013).

[41] J. D. Scheel and J. Schumacher, Global and local statistics in turbulent convection at low Prandtl numbers, J. Fluid Mech. **802**, 147 (2016).

[42] J. D. Scheel and J. Schumacher, Predicting transition ranges to fully turbulent viscous boundary layers in low Prandtl number convection flows, Phys. Rev. Fluids **2**, 123501 (2017).

[43] A. Pandey, D. Krasnov, R. Samtaney, J. Schumacher, and K. R. Sreenivasan, Similarities between characteristics of convective turbulence in confined and extended domains, Physica D **442**, 133537 (2022).

[44] K. Sugiyama, R. Ni, R. J. A. M. Stevens, T. S. Chan, S.-Q. Zhou, H.-D. Xi, C. Sun, S. Grossmann, K.-Q. Xia, and D. Lohse, Flow Reversals in Thermally Driven Turbulence, Phys. Rev. Lett. **105**, 034503 (2010).

[45] Y. Kifer, General random perturbations of hyperbolic and expanding transformations, J. Anal. Math. **47**, 111 (1986).

[46] R. Z. Khas'minskii, Principle of averaging for parabolic and elliptic differential equations and for Markov processes with small diffusion, Theor. Probab. Appl. **8**, 1 (1963).

[47] G. Froyland, Finite approximation of Sinai–Bowen–Ruelle measures for Anosov systems in two dimensions, Random Comput. Dyn. **3**, 251 (1995).

[48] A. Bittracher, R. Banisch, and C. Schütte, Data-driven computation of molecular reaction coordinates, J. Chem. Phys. **149**, 154103 (2018).

[49] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining* (U.S. Department of Energy, Washington, D.C., 1996), pp. 226–231.

[50] L. Molgedey and H. G. Schuster, Separation of a Mixture of Independent Signals Using Time Delayed Correlations, Phys. Rev. Lett. **72**, 3634 (1994).

[51] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, Identification of slow molecular order parameters for Markov model construction, J. Chem. Phys. **139**, 015102 (2013).

[52] F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé, Variational approach to molecular kinetics, J. Chem. Theor. Comput. **10**, 1739 (2014).

[53] J. Sesterhenn and P. Schmid, Dynamic mode decomposition of experimental data, Bull. Amer. Phys. Soc. **61**, 208 (2008).

[54] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, J. Fluid Mech. **656**, 5 (2010).

[55] S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte, and F. Noé, Data-driven model reduction and transfer operator approximation, J. Nonlin. Sci. **28**, 985 (2018).

[56] M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Pérez-Hernández, M. Hoffmann, N. Plattner, C. Wehmeyer, J. H. Prinz, and F. Noé, PyEMMA 2: A software package for estimation, validation, and analysis of Markov models, J. Chem. Theor. Comput. **11**, 5525 (2015).

[57] V. Baladi, *Positive Transfer Operators and Decay of Correlations* (World Scientific, Singapore, 2000), Vol. 16.

[58] F. Noé and C. Clementi, Kinetic distance and kinetic maps from molecular dynamics simulation, J. Chem. Theor. Comput. **11**, 5002 (2015).

[59] A. Bittracher, S. Klus, B. Hamzi, P. Koltai, and C. Schütte, Dimensionality reduction of complex metastable systems via kernel embeddings of transition manifolds, J. Nonlin. Sci. **31**, 3 (2020).

[60] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola, A Kernel Two-Sample Test, J. Mach. Learn. Res. **13**, 723 (2012).

[61] A. Smola, A. Gretton, L. Song, and B. Schölkopf, A Hilbert space embedding for distributions, in *International Conference on Algorithmic Learning Theory* (Springer, New York, 2007), pp. 13–31.

[62] T. Hofmann, B. Schölkopf, and A. J. Smola, Kernel methods in machine learning, Ann. Stat. **36**, 1171 (2008).

[63] A. Berlinet and C. Thomas-Agnan, *Reproducing Kernel Hilbert Spaces in Probability and Statistics* (Springer Science & Business Media, New York, 2011).

[64] R. R. Coifman and S. Lafon, Diffusion maps, Appl. Comput. Harmon. Anal. **21**, 5 (2006).

[65] F. W. Young, *Multidimensional Scaling: History, Theory, and Applications* (Psychology Press, New York, 2013).

[66] S. T. Roweis and L. K. Saul, Nonlinear dimensionality reduction by locally linear embedding, Science **290**, 2323 (2000).

[67] R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer, Graph Laplacian tomography from unknown random projections, IEEE Trans. Image Proc. **17**, 1891 (2008).

[68] T. Berry and J. Harlim, Variable bandwidth diffusion kernels, Appl. Comput. Harmon. Anal. **40**, 68 (2016).

[69] P. Koltai and S. Weiss, Diffusion maps embedding and transition matrix analysis of the large-scale flow structure in turbulent Rayleigh–Bénard convection, Nonlinearity **33**, 1723 (2020).