Heterogeneity of the attractor of the Lorenz '96 model: Lyapunov analysis, unstable periodic orbits, and shadowing properties

Physica D 457 (2024) 133970

Contents lists available at ScienceDirect

Physica D

journal homepage: www.elsevier.com/locate/physd

Heterogeneity of the attractor of the Lorenz '96 model: Lyapunov analysis, unstable periodic orbits, and shadowing properties

Chiara Cecilia Maiocchi [a,b], Valerio Lucarini [a,b,c,*], Andrey Gritsun [d], Yuzuru Sato [e,f,g]

[a] Department of Mathematics and Statistics, University of Reading, Reading, UK
[b] Centre for the Mathematics of Planet Earth, University of Reading, Reading, UK
[c] School of Systems Science/Institute of Nonequilibrium Systems, Beijing Normal University, Beijing, China
[d] Marchuk Institute of Numerical Mathematics, Russian Academy of Sciences, Moscow, Russian Federation
[e] RIES, Hokkaido University, N12 W7, Kita-ku, Sapporo 060-0812, Japan
[f] Department of Mathematics, Hokkaido University, Sapporo, Japan
[g] London Mathematical Laboratory, 8 Margravine Gardens, London, W6 8RH, UK

ARTICLE INFO

Keywords:
Unstable Periodic Orbits
Hyperbolicity
Lyapunov Analysis
Lorenz '96 Model
Transfer Operator
Shadowing

ABSTRACT

It is well known that the predictability of weather and climate is strongly state-dependent. Special, easily recognisable, and extremely relevant atmospheric states like blockings are associated with anomalous instability. This reflects the general property that the attractors of chaotic systems can feature considerable heterogeneity in terms of dynamical properties, and specifically, of their instability. The attractor of a chaotic system is densely populated by unstable periodic orbits (UPOs) that can be used to approximate any forward trajectory through the so-called shadowing. Dynamical heterogeneity can lead to the presence of UPOs with different number of unstable dimensions. This phenomenon – unstable dimensions variability – is a serious breakdown of hyperbolicity and has considerable implications in terms of the structural stability of the system and of the possibility to model accurately its behaviour through numerical models. As a step in the direction of better understanding the properties of high-dimensional chaotic systems, here we provide an extensive numerical investigation of the variability of the dynamical properties across the attractor of the much studied Lorenz '96 model. By combining the Lyapunov analysis of the tangent space with the study of the shadowing of the chaotic trajectory performed by a very large set of UPOs, we show that the observed variability in the number of unstable dimensions is associated with the presence of a substantial number of finite-time Lyapunov exponents that fluctuate about zero also when very long averaging times are considered. The transition between regions of the attractor with different degrees of instability is associated with a significant drop of the quality of the shadowing. By performing a coarse graining based on the shadowing UPOs, we are able to characterise the slow fluctuations of the system between regions featuring, on the average, anomalously high and anomalously low instability. In turn, such regions are associated, respectively, with states of anomalously high and low energy, thus providing a clear link between the microscopic and thermodynamical properties of the system.

1. Introduction

In Chapter 4 of his book Science and Méthode [1], Poincaré proposed for the first time the concept of sensitive dependence of the evolution of a system on its initial conditions, making also explicit reference to the relevance of this issue in the context of Meteorology [2]. After a long hiatus, the theoretical and practical relevance of the phenomenon of sensitive dependence on initial conditions and its compatibility with the presence of orbits contained in a compact set became apparent arguably through the seminal contributions by Lorenz [3], Ruelle and Takens [4], and Li and Yorke [5]. Since then, there has been a great effort in creating sophisticated mathematical frameworks for chaotic systems able to include, at the same time, phenomenology of practical relevance in science and engineering. In what follows, we provide a short summary of some key developments.

At first, uniform hyperbolicity was conjectured to be the default condition for systems featuring sensitive dependence on initial conditions [6–8]. The tangent space of uniformly hyperbolic systems can be split into contracting - stable - and expanding - unstable - subspaces, plus the direction of the flow in the case of time-continuous systems. Exponential divergence of initially infinitesimally nearby orbits is indeed compatible with a dynamics occurring in a compact manifold. This idea stems from the fact that hyperbolic dynamics is robust with respect to perturbations [9]. Smale and others conjectured that any chaotic

* Corresponding author at: Department of Mathematics and Statistics, University of Reading, Reading, UK.
E-mail address: v.lucarini@reading.ac.uk (V. Lucarini).

https://doi.org/10.1016/j.physd.2023.133970
Received 7 June 2023; Received in revised form 4 October 2023; Accepted 22 October 2023
Available online 31 October 2023

dynamical system could be transformed into a hyperbolic system by applying an appropriate perturbation. As a result of this conjecture, it would be possible to identify any chaotic behaviour emerging in a mathematical model of a physical process with a uniformly hyperbolic system. Uniformly hyperbolic systems are structurally stable [10] and amenable to analysis via linear response theory [11] but are, unfortunately, *very atypical* [12].

1.1. Beyond uniform hyperbolicity

Pesin theory has the great advantage of relaxing the definition of hyperbolicity by removing the requirement of being uniform [13] and focuses on studying the rates of expansion or decay (and corresponding modes) of perturbations with respect to a background trajectory, which are defined in terms of the properties of the so-called tangent linear operator. For nonuniformly hyperbolic systems, such asymptotic rates of expansion or contraction — the Lyapunov exponents (LEs) [14] are bounded away from zero, except for the direction associated with the flow (in the case of continuous time systems) [15]. The sum of the first p LEs (ordered from the largest to the smallest) describe the asymptotic rate of growth (or decay) of the p−volume of infinitesimal p−parallelepiped, and the sum of all the LEs gives the time-averaged growth rate of the phase space volume, which is negative – contractive – for dissipative systems [16]. The presence of a positive largest LE can be taken as a proof that a system has sensitive dependence on initial conditions. The notion of LEs, which are global quantities, has been extended via their local version, constructed by considering finite time horizon (finite-time LEs - FTLEs) [17–19], and by considering finite scale, rather than infinitesimal, perturbations with respect to the background trajectory (finite size LEs - FSLEs) [20]. In the last few decades, the so-called Lyapunov analysis has grown into a portfolio of very powerful methods for studying rather general complex systems [21,22].

A separate way to relax the notion of uniform hyperbolicity entails introducing a nontrivial centre manifold in the tangent space where expansion or contraction are extremely slow, thus removing hyperbolicity. The so-defined partial hyperbolic systems can be further generalised by allowing for nonuniformity outside the centre manifold: the continuous-time nonuniform partially hyperbolic systems feature more than one zero Lyapunov exponents [23].

A very influential attempt at creating a powerful paradigm of chaos beyond uniform hyperbolicity has been presented by Bonatti et al. [24], According to such a paradigm, chaos can originate from different mechanisms, such as heterodimensional cycles or homoclinic tangencies. As compared to uniformly hyperbolic systems, these more general systems are less understood and advances have been mainly made on discrete dynamical systems, because of the simpler structure of the phase space [6], hence numerical approaches play an important role in fueling rigorous mathematical investigations [25].

1.2. Unstable periodic orbits

In order to explain the variability of the local properties of the attractor of a chaotic system, the investigation of the tangent space via Lyapunov analysis can be complemented by a different strategy based on the study of special periodic solutions, the so-called Unstable Periodic Orbits (UPOs) [26–28]. In fact, UPOs, true nonlinear modes of the flow, provide a rigid structure hidden in the chaos of the dynamics. For uniformly hyperbolic systems without continuous symmetries, UPOs are dense in the attractor [29]. This implies that it is always possible to find a UPO that is arbitrarily near to a chaotic trajectory, allowing for a reconstruction of the trajectory up to arbitrary accuracy. We can then think of the chaotic trajectory as being continuously scattered between neighbourhoods of different UPOs, because of their instability. In general, even though the shadowing property has not yet been formalised for more general system, it is widely assumed. Evidence of turbulent trajectories being shadowed by UPOs has been found in forced two-dimensional flows [30–32], isotropic turbulence [33], plane Couette flow [34–36], Kolmogorov flow in two [37] and three dimensions [38]. UPOs can be used to approximate any chaotic trajectory with an arbitrary accuracy. While for Axiom A systems a rigorous theory that allows to reconstruct statistical properties of the system as sum with well-defined weights over the UPOs has been developed [28], extensions of this approach have been proposed for more general chaotic systems [39]. UPOs-based analysis has been successfully applied for decomposing and extracting information from the dynamical structure of chaotic flows in many different contexts [28]. Kawhara and Kida showed in their seminal work on a numerical simulation of a plane Couette flow that using just one UPO one manages to capture in a surprisingly accurate way the statistics of the turbulent flow [40]. Chandler and Kerswell [31] identified 50 UPOs of a turbulent fluid at a moderate Reynolds number and used them to reproduce the energy and dissipation probability density functions of the system as dynamical averages over the orbit. More recently, Yalniz and Budanur [38] proposed a coarse grained state space decomposition of the dynamics in terms of a few periodic orbits for both the three dimensional Rossler flow and a discretisation of the Kuramoto–Sivashinksky equation. Page et al. [41] provided evidence that the statistics of a fully developed turbulent flow can be reconstructed in terms of a set of UPOs.

Very interestingly, the investigation of the UPOs of a chaotic system allows one also to identify violations of hyperbolicity in a relatively simple manner. Indeed, if one detects e.g. two UPOs immersed in the attractor that feature a different number of positive LEs, hyperbolicity is broken through the mechanism of so-called unstable dimensions variability (UDV), which establishes the presence of a fundamental heterogeneity in the attractor and hinders the existence of an actual trajectory of the systems that stays uniformly close to a numerical one for long time intervals [42]. UDV is typically associated with the presence of large fluctuations of certain FTLEs between positive and negative values [43–45].

1.3. Unstable dimension variability in geophysical fluids

The study of the tangent space is a key aspect of the science and technology related to geophysical fluids [46]. In this context, it is well known that the predictability of a system, far from being in any sense uniform, is dramatically state-dependent: certain regions of the attractor feature larger instability than others [47,48], and this has great impact on data assimilation strategies [49,50]. In turn, the skill of data assimilation exercises can be used to infer the instability of the underlying system [51]. The state-dependent predictability results into substantial fluctuations in the value of the individual FTLEs [52–55]. These fluctuations can be accurately quantified, when considering sufficiently long time horizons, using large deviation laws [56], in agreement with the general theory presented in [57,58]. One also finds that the number of positive FTLEs fluctuates across the attractor, which is clearly indicative of a violation of the condition hyperbolicity and is associated with the UDV mentioned above [59]. The UDV can be particularly problematic for the efficiency of otherwise very powerful data assimilation schemes [51].

The use of UPOs for studying geophysical flows was introduced by Gritsun [60,61], who proposed using an expansion over UPOs to reconstruct the statistics of a simple atmospheric model based on the barotropic vorticity equation of the sphere. Later, Lucarini and Gritsun [59] used UPOs for clarifying the nature of blocking events in a baroclinic model of the atmosphere. Blockings are rare and persistent large scale deviations in the mid-latitudes from the approximately zonal flow. They are most commonly found in either the Atlantic or in the Pacific sector of the Northern Hemisphere. Specifically, they found that the atmospheric model was characterised by very large variability in the number of unstable dimensions (UDs), thus suggesting that the dynamics of the atmosphere is far from being hyperbolic. Additionally,

it was found that blocked states are associated with conditions of higher instability of the atmosphere, in basic agreement with a separate line of evidence obtained through Lyapunov analysis [62] and through the study of recurrent patterns of the atmosphere based on EVT [63].

1.4. This work: Concept and main results

In a previous work [64] we have studied the classical version of the 3-dimensional (3D) Lorenz 1963 (L63) model [3] using a rather extensive set of UPOs, following [65], and covering up to period 14 in symbolic dynamics. We have been able to investigate accurately the process of shadowing. We have also elucidated how by studying the statistics transitions between the neighbourhood of the various UPOs one can construct a finite-state Markov chain able to represent accurately the statistical and dynamical properties of the system, including its almost-invariant sets [66]. The attractor of the L63 model is extremely heterogeneous in terms of predictability, and features specific regions where return of skill is observed [67]. The detected UPOs do differ in terms of their dynamical characteristics, and specifically in the value of the first LE, thus providing a global counterpart of the heterogeneity of the properties of the tangent space. Nonetheless, in a 3D chaotic flow, by construction, all UPOs feature one positive, one negative, and one vanishing LE. Hence, if we want to investigate the heterogeneity of the attractor of a chaotic flow and possibly relate it to the presence of variability of the UDs number, one needs to consider higher dimensional systems.

With this work we would like to characterise and explain the heterogeneity of the attractor of the very popular Lorenz '96 (L96) model [68,69] in a chaotic regime. We show that the system features clear signature of UDV and that this is accompanied by a substantial number of FTLEs whose value fluctuates about zero also when very long averaging times are considered. By combining the information derived from the analysis of an extensive set of UPOs with Lyapunov analysis, we find that anomalously unstable UPOs preferentially populate regions of the attractor where, applying Lyapunov analysis to the tangent space, one gets, coherently, anomalously high instability indicators. This bridges the gap between global and local properties of the system.

Detecting UPOs in high dimensional, highly chaotic systems is well-known to be extremely challenging [28,60,70]. We show that, while we are able to identify over 3×10^5 UPOs up to period $T \approx 22$, the longer-period UPOs are (a) vastly underrepresented and (b) significantly skewed towards low instability. Such longer-period UPOs occupy preferentially a specific region of the phase space. This provides further support to the heterogeneity of the attractor in terms of instability and clarifies why restricting our analysis to more thoroughly detected, shorter-period UPOs leads to significant loss of information on the system.

We then propose two finite state Markov chain representations of the dynamics based on the shadowing process of the orbit performed by the UPOs. In first instance, UPOs are grouped in states according to their number of UDs and transitions are recorded at the point when the closest shadowing UPO changes and the new shadowing UPO has a different UDs number. In second instance, we consider a much larger space, where each state corresponds to the trajectory being shadowed by a UPO chosen among a suitable defined subset of the whole database. In both cases, by studying the subdominant eigenvectors and eigenvalues of the stochastic matrix, we are able to characterise the relatively slow fluctuations of the systems between regions with prevalence of anomalously unstable and anomalously stable UPOs, and between regions of typical vs atypical UPOs. We also provide a thermodynamical, energetic interpretation of the results, thus linking microscopic and macroscopic properties of the system.

The paper is structured as follows. Section 2 provides a description of the L96 model and of some of its basic properties in the configuration chosen for this study. In Section 3 we present the database of detected UPOs and discuss their accuracy in reproducing the dynamics of the system. We also present evidence of the heterogeneity of the attractor in terms of stability properties. In Section 4 we supplement the UPOs-based analysis with the Lyapunov analysis. We link local and global properties on the attractor, investigate the transitions between regions of the attractor characterised by different stability, and we emphasise how the breakdown of hyperbolicity and the associated UDV emerges according to these two viewpoints. The statistical angle on the problem is then discussed in Section 5, where we study the relaxation of a generic initial ensemble to the invariant measure by extracting two separate finite-state Markov chain from the dynamics. The relaxation can be seen as a mixing associated with transitions between regions of the attractor featuring anomalously high and anomalously low instability, respectively. Finally, in Section 6 we discuss the main results of our study and present perspectives of future investigations. In Appendix A we report the main mathematical concepts used throughout the paper for the benefit of the reader. In Appendix B one can find the link to the Supplementary Material that complements the content of the manuscript and provides access to some key datasets and codes.

2. The Lorenz '96 model: A toy model for spatio-temporal chaos

The L96 model, while not corresponding to a truncated version of any known fluid dynamical system, was developed as a prototype for the midlatitude atmosphere, with the scope of investigating problems of predictability in weather forecasting [68,69]. Each variable of the model corresponds to an atmospheric quantity of interest at a discrete location on a periodic lattice, representing a latitude circle on the sphere. The variables are spatially coupled, and their equation of motion include nonlinear (quadratic) terms to simulate advection, linear terms representing dissipation and constant terms representing external forcing. While the model only shares only such basic characteristics with more complete geophysical fluid dynamical models, it has emerged as an important testbed for different applications, including the study of bifurcations [71,72], of parametrizations [73–75], of data-driven and machine learning techniques, [76–78], of extreme events [79–81], of data assimilation schemes [82,83], of ensemble forecasting techniques [84,85], to develop new tools for investigating predictability [86,87], and for addressing basic issues in mechanics and statistical mechanics [88–92]. The evolution equations of the model are:

$$\dot{X}_j = (X_{j+1} - X_{j-2})X_{j-1} - \alpha X_j + F, \quad j = 1, \dots, J \tag{1}$$

where

$$X_{-1} = X_{J-1}, \quad X_0 = X_J, \quad X_{J+1} = X_1. \tag{2}$$

impose the periodicity conditions, $F \in \mathbb{R}^+$ is a constant forcing, and $\alpha \in \mathbb{R}^+$ modulates the dissipation. The three free parameters of the model are J, F, and α. For large values of F/α and J the model exhibits extensive chaos [91]. In the inviscid case – when the dissipation and the forcing are removed by setting $\alpha = F = 0$ – the so-called energy $E = 1/2 \sum_{j=1}^{J} X_j^2$ is conserved and one can recognise a quasi-symplectic structure in the dynamics [92]. The chaotic behaviour of the systems is associated with the presence of unstable waves that grow, accumulate energy, and then decay, with ensuing dissipation [68,69].

In this work we have considered $J = 20$, $F = 5$, and $\alpha = 1$, which leads to the onset of a chaotic steady state for the system; see some relevant definitions in Appendix A.1. In all the simulations the model is integrated with a Runge–Kutta second order midpoint time scheme with fixed time step $dt = 0.01$. The choice of this (suboptimal but sufficiently accurate) integrator is motivated by the UPOs detecting algorithm used. For such choice of the parameters the model is well within the chaotic regime. It features $n = 4$ positive LEs with the leading one being $\lambda_1 \approx 0.54$ (Fig. 1(a)). The fifth LE vanishes and corresponds to the direction of the flow. The characteristic Lyapunov time of the system is then $\tau_1 = 1/\lambda_1 \approx 1.85$, the Kolmogorov–Sinai

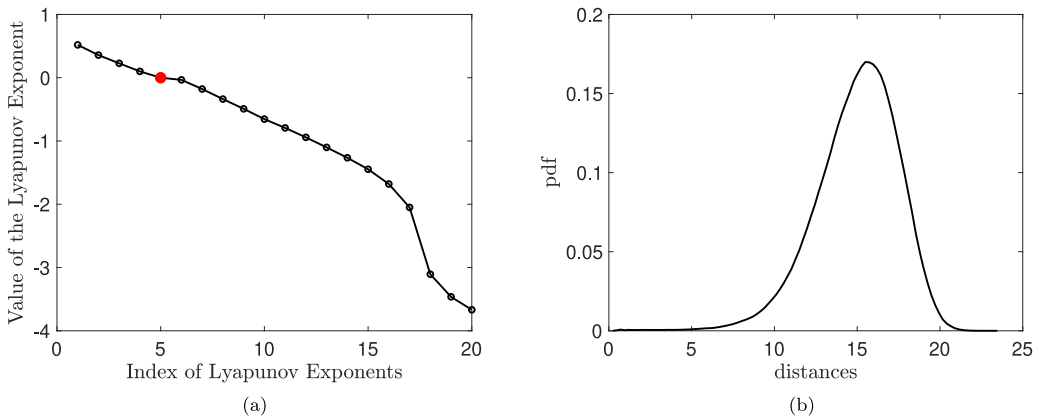

Fig. 1. Panel (a) Lyapunov Exponents of the system; $\lambda_5 = 0$, which corresponds to the direction of the flow, is highlighted in red. Panel (b): Distribution of pairwise distances between points of the attractor.

entropy can be approximated as $h_{KS} = \sum_{\lambda_i>0} \lambda_i \approx 1.23$ and the Kaplan–Yorke dimension of the attractor is $d_{KY} \approx 9.25$. See Appendix A.2 for clarifications on the mathematical terminology.

The diameter of the attractor is approximately 25, with point-to-point distances being distributed as shown in Fig. 1(b). The relatively high dimensionality of the attractor is apparent from the very low prevalence of nearby points [93]. The mean speed over the attractor is ≈ 38, meaning that on average the trajectory spans a distance of 0.38 for a time-step of $dt = 0.01$.

3. Study of the attractor via unstable periodic orbits

3.1. Database of the unstable periodic orbits

The numerical extraction of UPOs from a chaotic system - see definitions and some essential information in Appendix A.3 - is one of the greatest challenges in the applications of periodic orbit theory [60,94,95]. The problem of finding UPOs can be reduced to the solution of the periodicity condition $S^T x_0 = x_0$, where S^T is the evolution operator associated with the flow given in Eq. (1) that acts for a time T and x_0 is an initial condition on the attractor. Hence, one must solve a system of nonlinear equations with respect to x_0 and the period T of the UPO. Even for simple nonlinear systems this represents a complex numerical problem, with a computational cost that grows exponentially with the dimension of the system. The choice of the algorithm and initial condition is thus crucial (see [60] for details). The equations are symmetric with respect to a cyclic permutation of the variables, so that each time an orbit is detected, the other $J-1 = 19$ can be automatically obtained by simply considering all the possible cyclic permutations of variables. In this work we constructed a database of 15019 fundamental UPOs (i.e. none of these orbits can be obtained from another orbit of the database through cyclic symmetry) immersed in the attractor with period ranging from a minimum of ≈ 1.5 ($\approx 0.8/\lambda_1$) to a maximum of ≈ 22.8 ($\approx 12.3/\lambda_1$). This corresponds to lengths ranging from ≈ 2 to ≈ 35 diameters of the attractor.

In a chaotic system one expects to find that the number of UPOs with prime period smaller or equal than T grows as $\propto \exp(h_{top}T)$, where h_{top} is the topological entropy [28]. As opposed to our previous study [64], it is clear that our set of UPOs is far from being complete, as longer-period orbits are clearly underrepresented, see the curve referring to the complete database in Fig. 2(a). The difficulty in computing long-period UPOs has been widely discussed in the literature [28,60,70]. While clearly incomplete, as further discussed below, the dataset of detected UPOs can provide extremely valuable information on the properties of the L96 model. The detected UPOs are characterised by vastly different instabilities properties, which provide a clear evidence of the heterogeneity of the attractor of the L96 model. The UDs number varies from 2 to 9 across the UPOs (Fig. 2(b)): this indicates a very serious violation of hyperbolicity via UDV [42–45,59]. Additionally, the Kolmogorov–Sinai entropy varies between ≈ 0.5 and ≈ 10.0 (Fig. 2(c)) and the first LE varies between ≈ 0.3 and ≈ 1.8 (Fig. 2(d)).

Figs. 2(a)–3(a) show that the periods of the detected UPOs are approximately integer multiples of a fundamental period $T_f \approx 2.1$, which is associated with the least unstable UPO. Fig. 3(a) additionally shows that the detected longer-period UPOs we find are significantly skewed towards low instability; it is extremely difficult to detect long-period, highly unstable UPOs. Indeed, we find that $\lambda_1^{max,T} \times T \approx const.$, where $\lambda_1^{max,T}$ is the largest first Lyapunov exponent detected among the UPOs with period T. This can be explained by considering that the UPOs detection algorithm is aimed at controlling errors that grow $\propto \exp(\lambda_1 T)$ for a UPO with period T[60]. The scatter plot of the UPOs period vs. h_{KS} is qualitatively similar; see Fig. 1(a) in the Supplementary Material. Indeed – compare Figs. 2(c) and 2(d) with Figs. 3(c) and 3(d), respectively – if we restrict the statistics of λ_1 to the orbits with $T \leq 6.4$, where the cutoff corresponds to the first peak of the distribution in Fig. 2(a), which marks the departure from the exponential growth of the number of UPOs with respect to their period, we obtain distributions of h_{KS} and λ_1 that are shifted towards the right with respect to what has been obtained using the whole dataset of UPOs. Even when considering such set of low-period UPOs one finds considerable UDV; see Fig. 3(b). In agreement with the previous observations, the distribution of the UDs number is skewed to higher values with respect to what found when considering the whole dataset, compare with Fig. 2(b).

It is worth exploring whether one find a clear localisation in phase space of the anomalously stable and anomalously unstable UPOs. In this regard, it is helpful to provide a visual representation of the UPOs. This can be achieved by considering the 3D projected space over the first three normalised time-dependent moments C_1, C_2, C_3, defined as:

$$C_k = \frac{\left(\sum_{j=1}^{J} X_j^k\right)^{1/k}}{\left(\langle \sum_{j=1}^{J} X_j^2 \rangle\right)^{1/2}}, \quad k = 1, 2, 3. \tag{3}$$

where $\langle \bullet \rangle$ indicates the expectation value computed according to the invariant measure of the system.

Fig. 4(a) represents all UPOs of the database in such a projected space. We will come back to this representation in Section 5. One notes that less unstable UPOs are localised in the part of the projected space closest to the origin. Such visual impression is further supported by Figs. 4(b), 4(c), and 4(d), which portray the statistics of the C_k's stratified according to the UD of the UPOs. The average value of C_k computed over UPOs with a given number of UD increases monotonically with UD. In the case of C_2, this corresponds to the intuition that

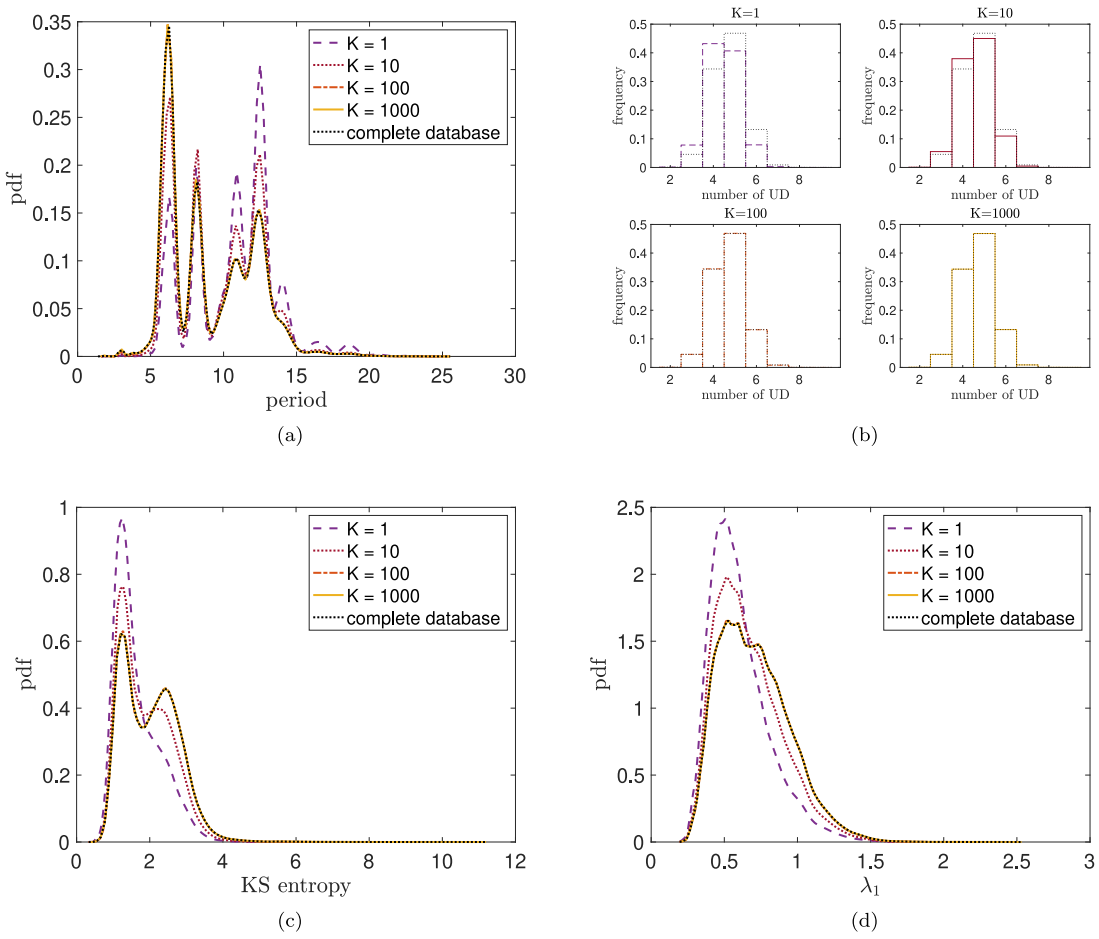

Fig. 2. Heterogeneity of the instability properties of the shadowing UPOs. Each distribution represent the properties of the UPOs that shadows the trajectory at least once in the first $K = 1, 10, 100, 1000$ tiers (see Section 3.2) and the complete database. Panel (a): Distribution of the periods of the shadowing UPOs that shadows - see Section 3.2 - the trajectory at least once in the first $K = 1, 10, 100, 1000$ tiers. Panel (b): Frequency of the number of UD of the shadowing UPOs for different values of K. Panel (c): Distribution of the KS entropy of the UPOs. Panel (d): Distribution of λ_1 across the UPOs of the database. The colour code is the same for all panels.

higher instability is associated with higher energy [68,69,92]. Indeed, one can interpret the periodic variations in the value of the C_k's as describing the life-cycle of the unstable waves defined by the individual UPOs.

Additionally, we have that the region of the phase space covered by lower-period – and thus anomalously unstable – UPOs misses the part of the attractor whose projection is closer to the origin in the (C_1, C_2, C_3) space, compare Fig. 4(e) with Fig. 4(a). Further evidence is shown in Figs. (1b–1d) of the Supplementary Material, where we show that mean values of the C_k's stratified according to the UD of the UPO are shifted to higher values when only lower-period UPOs are considered. Fig. 4(f) portrays the projection of the invariant measure of the system in the (C_1, C_2) and shows the region where the presence of lower-period UPOs is scarce. Hence, when the chaotic trajectory is in the region of phase space corresponding to the black circle in Fig. 4(e) or the rectangular region in Fig. 4(f), there will be only few and nearby lower-period UPOs, so that the best local approximation to the trajectory will necessarily come from higher-period ones.

3.2. Shadowing by unstable periodic orbits

As discussed in the introduction, UPOs can be used to approximate the forward trajectory of chaotic dynamical systems. We say that a UPO is shadowing a chaotic trajectory if the UPO is sufficiently close to the chaotic trajectory and co-evolves with the trajectory for some period of time. Here we proceed as [64] and first introduce the concept of "rank shadowing".

Let $\mathcal{U} = \{U_k\}_{k=1}^{N_{UPO}}$ be the set of UPOs of the database where the kth UPO is defined as $U_k = \{u_k(s)\}_{s=1}^{T_k/dt}$, with T_k being its prime period and $dt = 0.01$ the time step. We have a total of $N_{UPO} \approx 3 \times 10^5$ UPOs. We consider a chaotic trajectory $\mathcal{X}_{chaotic}$ consisting of the set of points $\mathcal{X}_{chaotic} = \{x(t)\}_{t=1}^{N_{max}}$ with output given every dt and with length $T_{max} = N_{max} \cdot dt = 3 \times 10^4$. We assume that the chaotic trajectory lives on the attractor of the system, *i.e.* transients have been discarded. We define as distance between the UPO $U_{\bar{k}}$ and the chaotic trajectory $\mathcal{X}_{chaotic}$ at time t as $d_k(t) = \min_s |u_{\bar{k}}(s) - x(t)|$. At each time t we can rank the UPOs according to their distance form $x(t)$. U_l is the tier 1 UPO at time t if $l = \arg\min_{k=1,\dots,N_{UPO}}(d_k(t))$. Furthermore, U_p is the tier K UPO at time t if $d_p(t)$ is the Kth smallest value among all $d_k(t)$'s, $k = 1, \dots, N_{UPO}$.

When constructing the time-dependent ranking of the UPOs in terms of proximity to the chaotic trajectory, we discover that only around 26% of the total is selected at least once as first tier orbit. Instead, the fraction of UPOs that belong at least once to the first $K = 10, 100, 1000$ tiers increases up to $79\%, 99.6\%$, and 99.9%, respectively, of the complete database. The four panels of Fig. 2 show that the $K = 1$ tier shadowing orbits are preferentially of lower instability (and, hence, feature typically a longer period, see Fig. 3(a)) compared to the whole database because they repel less intensely nearby trajectories. As we include orbits belonging to higher tiers of proximity, the statistical properties rapidly converge to those of the whole database. If we consider only the (much more homogeneous) lower-period UPOs, there is little dependence of the level of instability of the UPOs on their tier of proximity, see Figs. 3(b) and 3(d). As confirmed below, even orbits

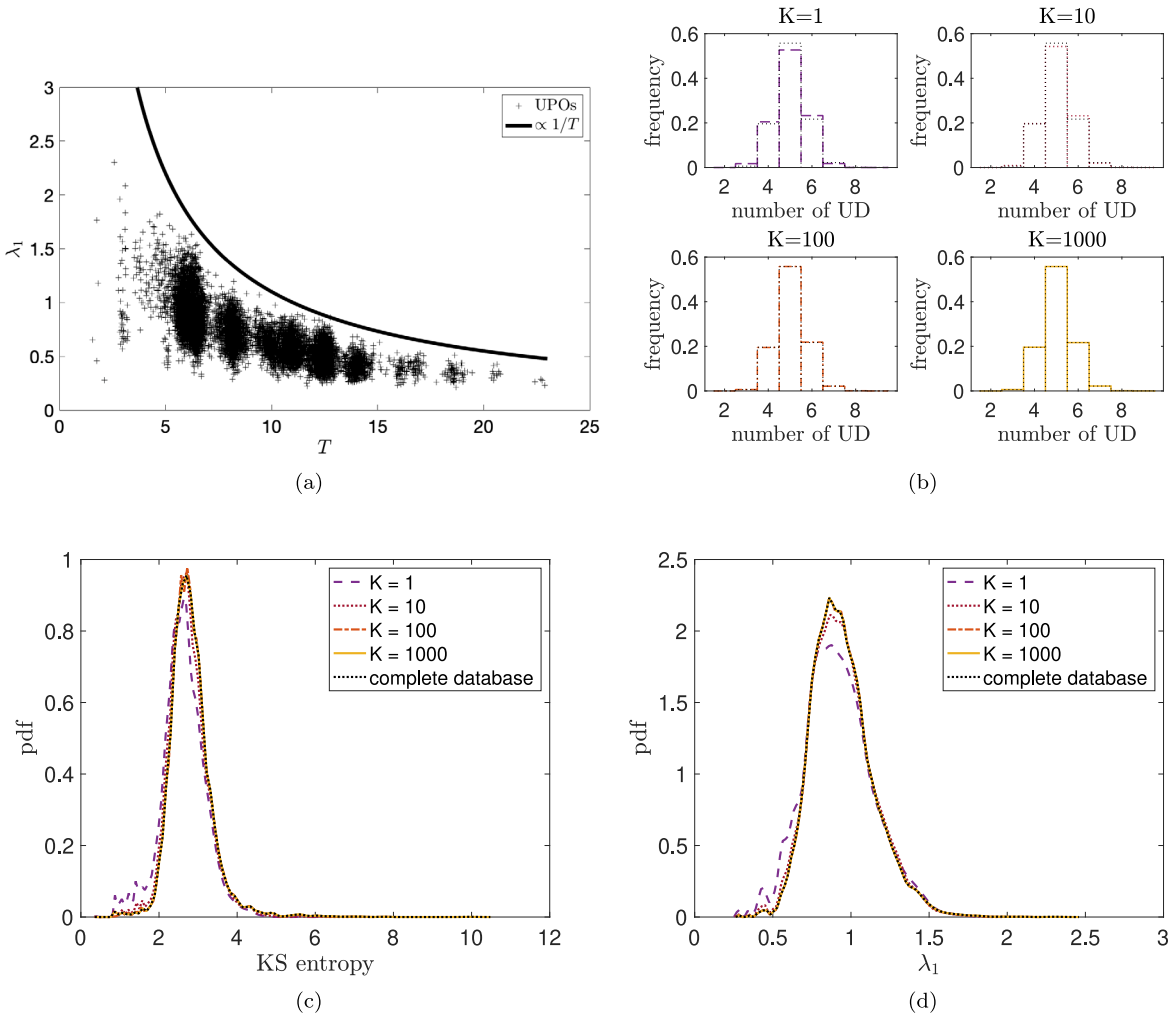

Fig. 3. Subsampling the set of detected UPOs. Panel (a): Scatter plot period vs. λ_1 for the detected UPOs. Panels (b), (c), and (d): Same as Figs. 2(b), 2(c), and 2(d), respectively, but restricted to UPOs with $T \leq 6.4$.

belonging to the $K = 10, 100, 1000$ tiers are very often rather close to the chaotic trajectory.

We will consider two definitions of shadowing orbits. The first is a stricter definition, which considers as shadowing UPOs only the sequence of tier 1 UPOs. We also consider a looser definition, which allows one to slightly prioritise persistence over proximity. In fact, at each time step, a UPO might still provide a very good local approximation to the trajectory even if it is not anymore the nearest one. In particular, if U_l is the closest UPO to the trajectory at time t, we say that U_l ceases to shadow the trajectory at time $t + pdt$ if at that time U_l is not anymore one of the K closest UPOs, or, in other terms, it does belongs to one of first K tiers. We then collect the time series of the distances $d_l(t + jdt)$, $j = 0, \ldots, p - 1$ and we say that U_l's shadowing duration (or U_l's persistence) is pdt. At time $t + pdt$ the tier 1 UPO is selected as the next shadowing UPO. The strict definition of shadowing is obtained by setting $K = 1$.

An example of shadowing (case $K = 1$) is presented in Fig. 5. Panel (a) shows how a short portion of duration $t_1 + t_2 = 2.65$ of the chaotic trajectory is subsequently shadowed by two UPOs featuring different period and different UD. The juxtaposition of the two pieces of the shadowing UPOs of duration $t_1 = 1.78$ and $t_2 = 0.87$, gives a time-dependent field that is visually very similar to the considered portion of the chaotic trajectory. Panel (b) shows the time evolution of the distance between the trajectory and the two UPOs. At time $t = t_1$ there is a transition in the shadowing, as the second UPO becomes the closest one to the chaotic trajectory.

By construction, choosing larger values of K in the definition of shadowing leads to an increase in the average distance between the orbit and the shadowing UPOs, and, at the same time, to an increase in the persistence of each shadowing UPOs. Figs. 6(a) and 6(b) show the distribution of the distance d and of the persistence π of the shadowing UPOs when considering the strict (tier 1, in black) and the looser definition of shadowing, with $K \in \{10, 100, 1000\}$. The average distance increases from 1.81 to 2.70 as K increases from 1 to 1000 (see Table 1). One notices that the quality of the shadowing is in general rather good and extremely similar for $K = 1$ and $K = 10$: the 95% quantile of shadowing distances is approximately 2.99 and 3.21, which is well within the 0.5% quantile of the typical distances distribution over the attractor (See Fig. 1(b)) The mean persistence, in turn, increases from 0.22 to 1.69 time units. These values correspond to average rectified distances ranging from 8 to 63. By comparing these numbers with the size of the attractor, the typical distances over the attractor (Fig. 1(b)), and the average distance between the chaotic trajectory and the shadowing UPOs confirms that there is clear evidence of co-evolution. Assessing co-evolution is instrumental for constructing a framework that allows an accurate statistical and dynamical description of the chaotic flow [96].

Restricting our analysis to the lower-period UPOs unavoidably leads to an increase in the distance between the shadowing UPOs and the chaotic trajectory, compare Fig. 6(d) and Fig. 6(b), and look at Table 1. Note that, instead, the statistics of the log persistence changes negligibly, compare Fig. 6(c) and Fig. 6(a), which indicates that the procedure

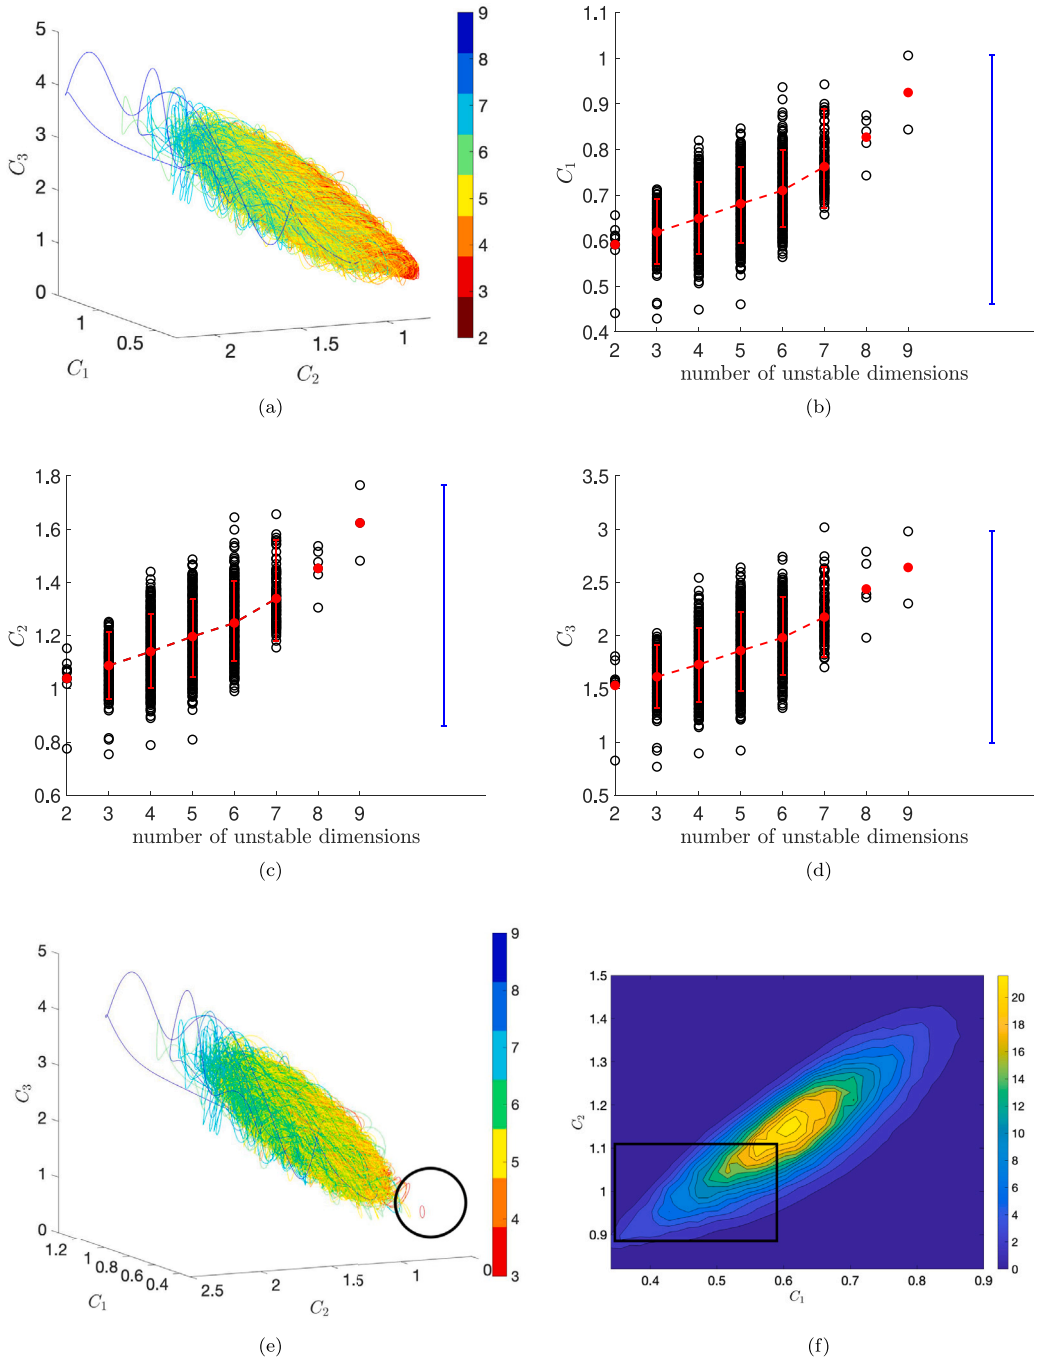

Fig. 4. Panel (a): UPOs of the system projected on the normalised moments C_1, C_2, C_3. The colour indicates the dimension of the unstable subspace. Panel (b): scatter plot of the first moment C_1 averaged along all the UPOs with the same number of UD vs their number of UD (black dots). The bars indicate the range between the 5th and 95th percentile for each UDs value. In blue the range between 5th and 95th percentile value of the corresponding statistics calculated along the chaotic trajectory. Panel (c): same as Panel (b), but for C_2. Panel (d), same as Panel (b), but for C_3. Panel (e): same as Panel (a), but for UPOs with $T \leq 6.4$. The circle indicates the region where major discrepancy is found with respect to Panel (a). Panel (f): Projection of the invariant measure of the system in the (C_1, C_2) space. The rectangle depicts the region that is sparsely covered by UPOs with $T \leq 6.4$; compare with the circle in panel (e).

is robust. Let us focus on the statistics of the distances between the tier 1 UPOs and the chaotic trajectory. The average distance increases by a rather substantial factor ≈ 1.9: excluding the higher period UPOs amounts to losing not only about 80% of all UPOs, but also the geometrically longest and dynamically least unstable ones, so that the ability to cover the attractor is substantially reduced. Additionally, if we restrict our analysis to the $T \leq 6.4$ UPOs, the average distance between the tier 1 UPOs and the chaotic trajectory increases by a factor ≈ 2.5 when considering the rectangular region in Fig. 4(f), whereas it increases by a factor ≈ 1.8 outside of it. This further supports what has been shown in Figs. 4(e) and 4(f) regarding the scarcity of the $T \leq 6.4$ UPOs in the low-energy region of the attractor.

4. Unstable periodic orbits and unstable dimension variability

4.1. Finite-time Lyapunov exponents and unstable dimension variability

In [43–45] it was shown how the UDV could be explained in terms of the presence of fluctuations between positive and negative values

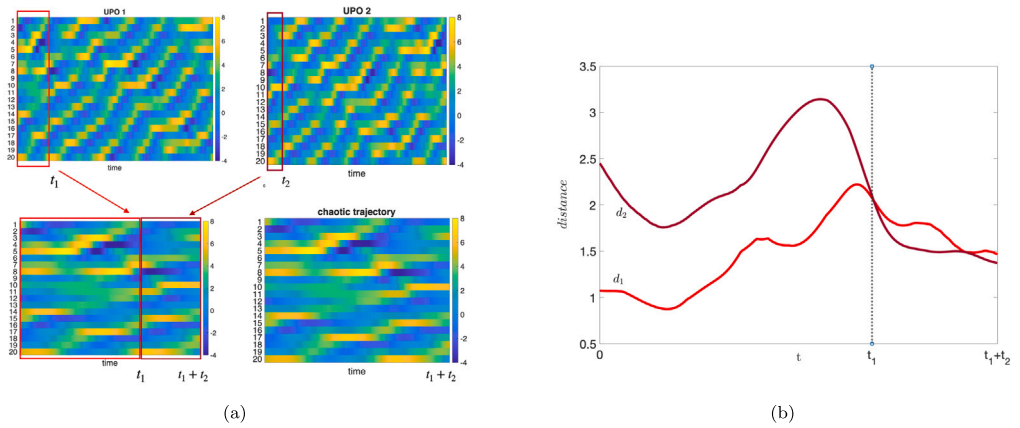

Fig. 5. Panel (a): Subsequent shadowing of a segment of chaotic trajectory performed by two different UPOs (here UPO1 and UPO2). UPO1 (UPO2) has period $T_1 = 10.8748$ ($T_2 = 10.7626$) and possesses 5 (4) positive LEs. In each of the four subpanels the x-axis indicates the time, the y-axis indicates the index j of the X_j, $j = 1, \ldots, 20$ variables of the Lorenz '96 system, and the colorbar indicates the value of each variable. The space–time diagram of UPO1 and UPO2 over their full period is reported in upper row. The chaotic trajectory (lower row, right hand side) is shadowed by UPO1 (bright red) for a time duration $t_1 = 1.78$ and then by UPO2 (dark red) for a time duration $t_2 = 0.87$. The two portions of the UPOs performing the shadowing are reported in the lower row, left hand side. Panel (b): Time evolution of distance d_1 and d_2 between the chaotic trajectory and UPO1 and UPO2, respectively.

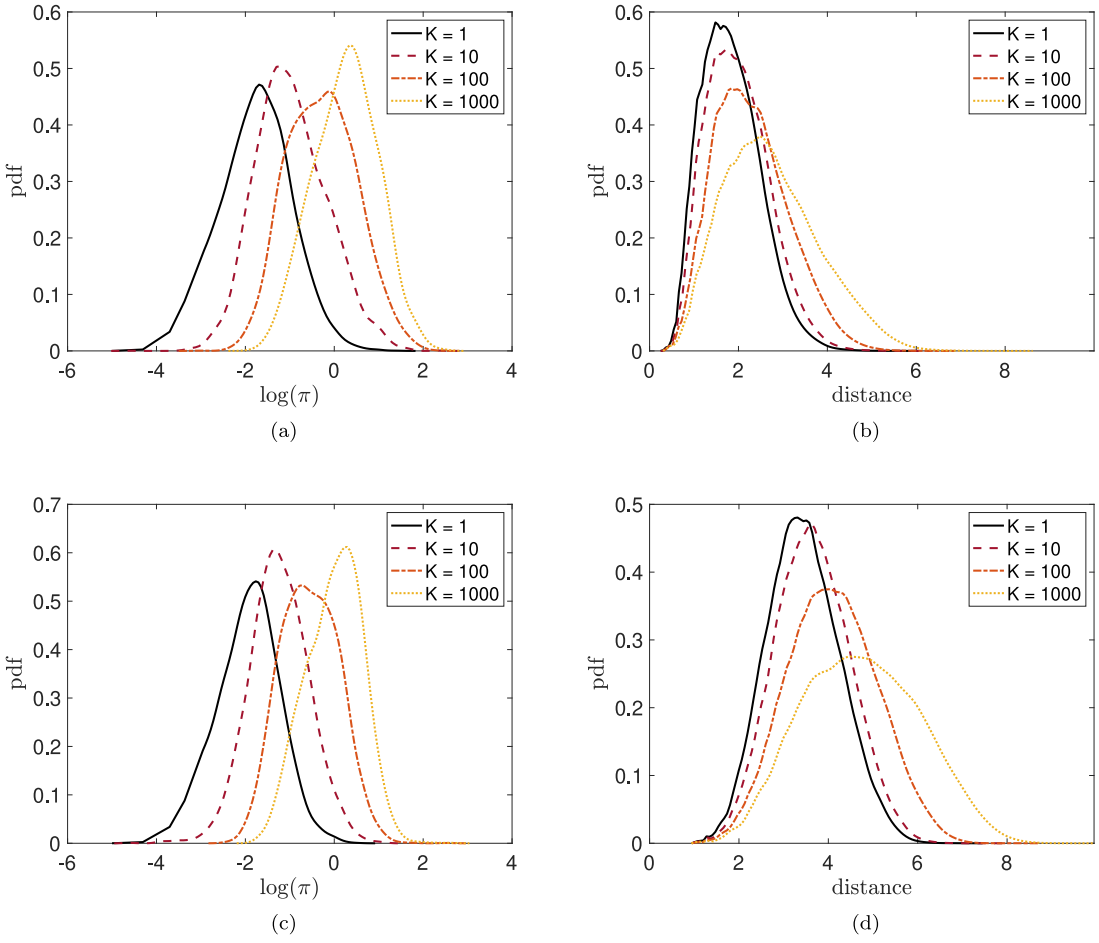

Fig. 6. Panel (a): Probability distribution function of the log of the persistence π of the tier 1 orbits (solid black line; mean persistence = 0.21), and the shadowing orbits with the looser definition with $K = 10$ orbits (dashed red line, mean persistence 0.53), $K = 100$ (dashed and dotted orange line, mean persistence 1.03) and $K = 1000$ (dotted yellow line, mean persistence 1.67). Panel (b): Distribution of the distances from the chaotic trajectory when considering the looser definition of shadowing orbits that allows for fluctuations within the first K=10 tiers (red dashed line, mean distance 1.97), tier 100 orbits (dashed and dotted orange line, mean distance 2.26), tier 1000 orbits (dotted yellow line, mean distance 2.7) and again first tier orbits (black solid line, mean distance 1.81). Panel (c): Same as Panel (a), for $T \leq 6.4$ UPOs. Panel (d): Same as Panel (b), for $T \leq 6.4$ UPOs.

of one specific FTLE (the one corresponding to the LEs with smallest absolute value) computed over a time scale τ also when considering very large values of τ. See definitions and some essential information on FTLEs in Appendix A.2. The presence of changeovers in the sign of the FTLE over such long time scales was proposed as evidence of the trajectory following closely UPOs having different UDs number.

Table 1
We report the distance statistic for both the complete (left hand side of the table) and the reduced database of UPOs having $T \leq 6.4$ (right hand side of the table). The first and second rows indicate the 95th (5th respectively) percentile of the distribution of distances of the shadowing orbits from the chaotic trajectory in the first tier and for $K = 10, 100, 1000$. The third row reports the mean value of the distance for each tier. The fourth row reports the probability of achieving pairwise distance across the attractor smaller than the mean value reported in the third row. We can appreciate how such probability never exceeds 1%, confirming the closeness of the trajectory with the shadowing orbits.

K	1	10	100	1000	1	10	100	1000
95th percentile	2.99	3.21	3.76	4.64	4.82	5.05	5.76	6.90
5th percentile	0.87	0.95	1.05	1.17	2.15	2.29	2.48	2.68
$mean(dist)$	1.81	1.96	2.26	2.70	3.44	3.64	4.09	4.74
$P(d < mean(dist))(\%)$	0.09	0.10	0.12	0.14	0.17	0.18	0.21	0.25

As mentioned earlier, in the system of interest here the UDV entails fluctuations between 2 and 9 of the number of UDs, hence one could expect to find that several FTLEs feature fluctuations between positive and negative values.

This is indeed confirmed by our data. We consider here the first 10 LEs, ordered from the largest to the smallest. Note that for sufficiently large values of τ the distribution of all the FTLEs corresponding to nonzero LEs converges to a Gaussian with variance $\propto 1/\tau$, in agreement with previous studies [56–58]. A different scaling is found for the FTLE corresponding to the vanishing LE (not shown).

Even considering very long averaging times, the support of the pdf of more than one FTLEs includes zero, meaning that one observes fluctuations about zero for the corresponding time series. Putting aside the vanishing LE, for which this property is trivial, this applies to 5 FTLEs for $\tau = 10\tau_1$, which is already much longer than the period of the longest detected UPO, see Fig. 7(a). Clearly, the number of FTLEs fluctuating about zero decreases as one consider larger values for τ. Nonetheless one finds four of such FTLEs when $\tau = 30\tau_1$ (Fig. 7(b)), and still two for the ultralong averaging time $\tau = 100\tau_1$ (Fig. 7(c)). The fluctuations about zero of the 6th FTLE persist even for much long longer averaging times.

Focusing on indicators of instability provided further information on the heterogeneity of the attractor. Panel 7(d) shows that the sum of the FTLEs corresponding to the four largest backward LEs – this provides the finite-time, local estimate of the Kolmogorov–Sinai entropy – have very large fluctuations, and the distribution has support extending to negative values up to averaging times of about $\tau = 3\tau_1$. We also see – Panel 7(e) – that the largest (ordered) backward FTLE can have negative values for averaging times up to $\tau = 3\tau_1$. This implies that the system features (temporary) return of skill. Finally, Panel 7(f) shows the distribution of the number of positive FTLEs. Note that we have removed from the count the direction of the flow, whose corresponding FTLE obviously fluctuates between small positive and small negative values. We find confirmation that the number of positive FTLEs has very large fluctuations even for very long averaging times τ.

4.2. Finite-time Lyapunov exponents and shadowing unstable periodic orbits

We now wish to provide an interpretation of such variability in terms of UPOs. Our intuition is that the local stability properties of the tangent space, measured in terms of the values of the FTLEs, are somehow encoded in neighbouring UPOs populating that same region of the attractor. In what follows, it is important to keep in mind that UPOs are large-scale structures in the phase space of the system and have a very long period compared to the typical shadowing times. With reference to the notation and framework defined in Fig. 4(a), let us suppose that between the times t_k and $t_k + \tau_k$ the chaotic trajectory $x(t)$ is being shadowed by the UPO U_k, before being approximated by another UPO U_h starting at time $t_h = t_k + \tau_k$. We then have a sequence of shadowing orbits U_k, each one associated to a persistence time τ_k. We then compute the spectrum of the FTLEs of the chaotic trajectory between t_k and $t_k + \tau_k$, and investigate the correlation with the corresponding LEs of the shadowing UPO U_k. Please note that each orbit might be considered more than once when looking at the correlations. Note also that the values of the time intervals τ_k change substantially along the trajectory, hence the various considered FTLEs are in general computed for different time horizons.

Table 2
Temporal correlation between the local properties of the chaotic trajectory and relative shadowing UPOs.

K	1	10	100	1000
λ^τ_{max}	0.23	0.27	0.33	0.39
λ^τ_1	0.15	0.17	0.22	0.27
$h^\tau_{KS,+}$	0.34	0.39	0.45	0.53
h^τ_{KS}	0.25	0.30	0.36	0.44

The results are presented in Table 2. There is a weak yet positive linear correlation between the first LE λ_1 of the shadowing UPOs and the corresponding first FTLE λ^τ_1 when considering the orbits shadowing the trajectory in the first tier. The correlation is stronger if, instead, we consider the largest local LE λ^τ_{max} (i.e. the largest one obtained after reordering the local FTLEs). We need to remember that we are comparing two very different objects: a local property with a global structure. The link between local and global properties is clearer when we consider a measure of instability that encompasses the whole unstable manifold. The linear correlation between the sum of the first four FTLE (KS^τ) and the Kolmogorov–Sinai entropy of the shadowing UPOs (first tier) is about 0.25, whilst a higher value (0.34) is found when considering the sum of the local positive FTLEs (KS^τ_+). As we relax our definition of shadowing according to what described in Section 3.2, we obtain that as we consider larger and larger values of K, the correlation between the corresponding local and global properties increases. Such an increase in the correlation derives from the fact that higher values of K result in longer shadowing times for the UPOs (Fig. 6(a)), while maintaining good proximity to the trajectory (Fig. 6(b)).

It is possible to further explore such an aspect by looking in more detail into the properties of the shadowing UPOs chosen according to the strictest $K = 1$ criterion. We then investigate how the correlation between the local instability properties of the chaotic trajectory and those of the shadowing UPOs change as we consider only UPOs that shadow the trajectory for longer and longer time durations. This allows us to restrict our analysis to what we may consider as the better performing UPOs. Fig. 8 shows that the link between the local instability properties of the trajectory and the UPOs steadily grows stronger as we consider UPOs that shadow for longer and longer times. As an example, the linear correlation between the first LE of the UPO λ_1 and the first FTLE λ^τ_1 of the chaotic trajectory increases from 0.15 - full database - to around 0.35 - when considering only UPOs that shadow for a time duration larger than 1. Of course, by definition, the number of UPOs considered for the statistics decreases as we demand more persistence (see Fig. 6(a)), hence the uncertainty in our estimates grows as the UPOs dataset shrinks in size. Analogously, the linear correlations between the Kolmogorov–Sinai entropy of the shadowing UPO and KS^τ increases from 0.25 to around 0.50. While these correlations might appear as relatively weak, we remark again that we are studying the link between local and global properties. Our results support the conjecture that the local stability properties of the flow can indeed be explained in terms of the properties of the shadowing UPOs. We also remark that, as shown in Fig. 2 of the Supplementary Material, the correlations shown in Fig. 8 are substantially degraded if one considers the database of $T \leq 6.4$ UPOs, as a result of the much reduced ability to sample accurately the attractor of the system that we have already discussed earlier in the paper.

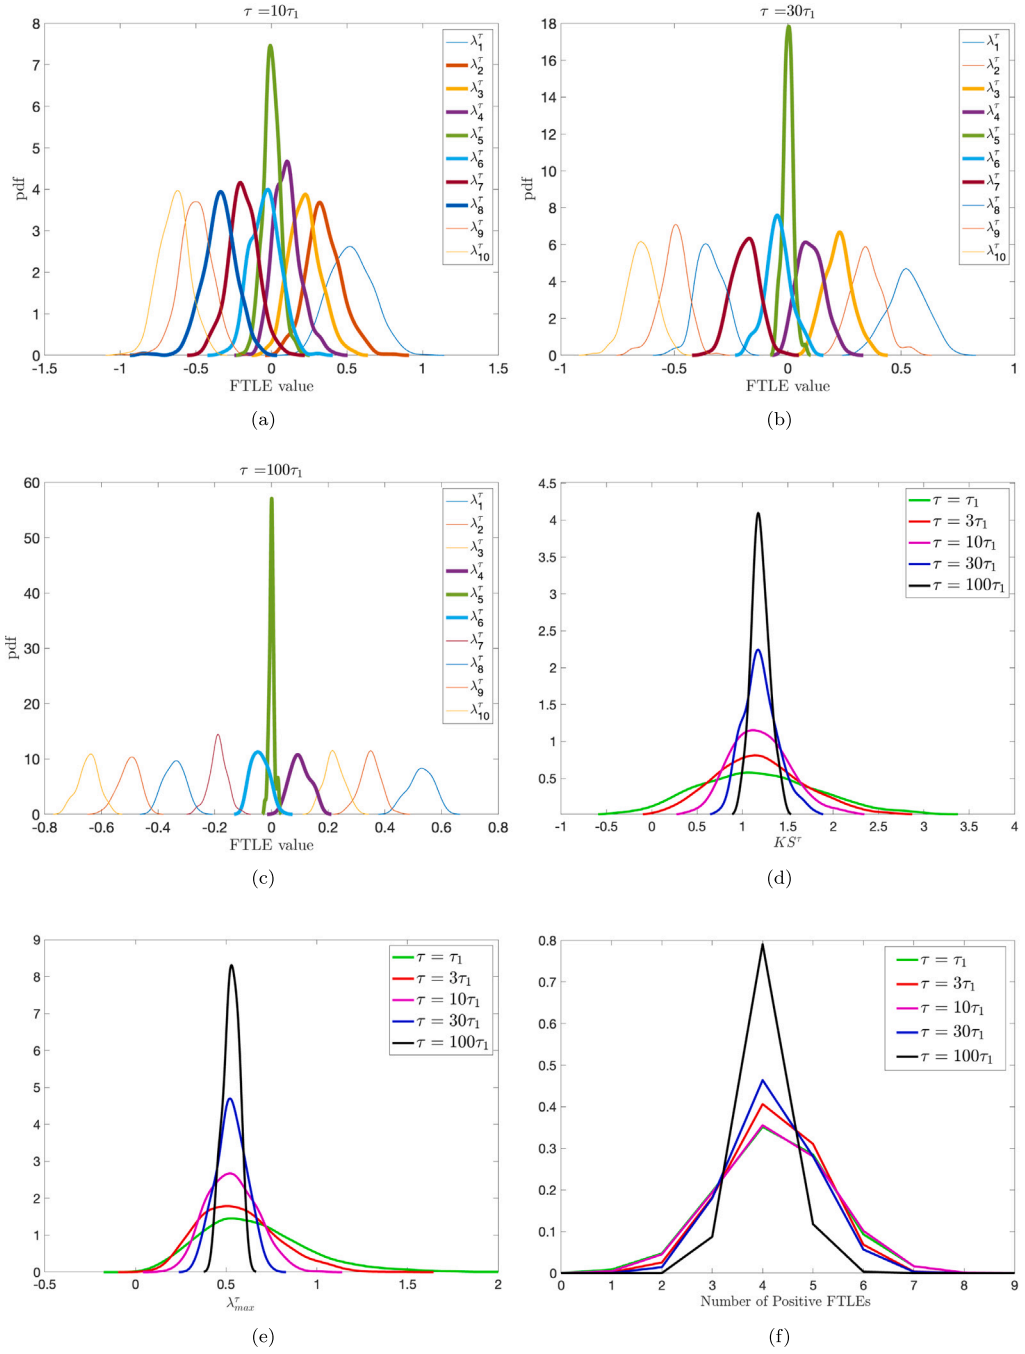

Fig. 7. Evidence of heterogeneity of the tangent space. The 5th LE corresponds to the direction of the flow. Panel (a): Distribution of the first 10 FTLEs with averaging time $\tau = 10\tau_1$. Thick lines correspond to pdfs whose support include zero. Panel (b): Same as (a), with $\tau = 30\tau_1$. Panel (c): Same as (a), with $\tau = 100\tau_1$. Panel (d): Distribution of the sum of the first four backward FTLEs for different averaging times τ. Panel (e): Distribution of the largest FTLE for different averaging times. Panel (f): Distribution of the number of positive FTLEs for different averaging times τ.

4.3. Unstable dimension variability and quality of the shadowing

The UDV is associated with another important phenomenon, namely the so-called shadowing breakdown [97]. This refers to the fact that, whereas for hyperbolic system, the Anosov [8] and Bowen's [98] shadowing lemma guarantees, roughly speaking, that each pseudo-orbit – e.g. the output of a numerical model of the system – stays uniformly close to some true trajectory, this does not necessarily holds true for non-hyperbolic systems. However, even though it is not possible to know in detail whether a computer simulated trajectory is representative of a true trajectory of the system, one can estimate for how long the numerical trajectory remains close to the true one [43,99]. In general, it is expected that the distance between the true trajectory and the pseudo-orbit increases when the trajectory goes through a glitch point, i.e when it performs a transition between regions featuring a different UDs number [97,99].

We wish to test whether a signature of the glitch points emerges as a deterioration of the ranked shadowing properties of the UPOs. We thus look at the statistics of distances between UPOs and the chaotic trajectory at the transitions points, i.e. when a new UPO takes over as best local approximation of the trajectory. We distinguish between the case where the new shadowing UPO has the same number of positive

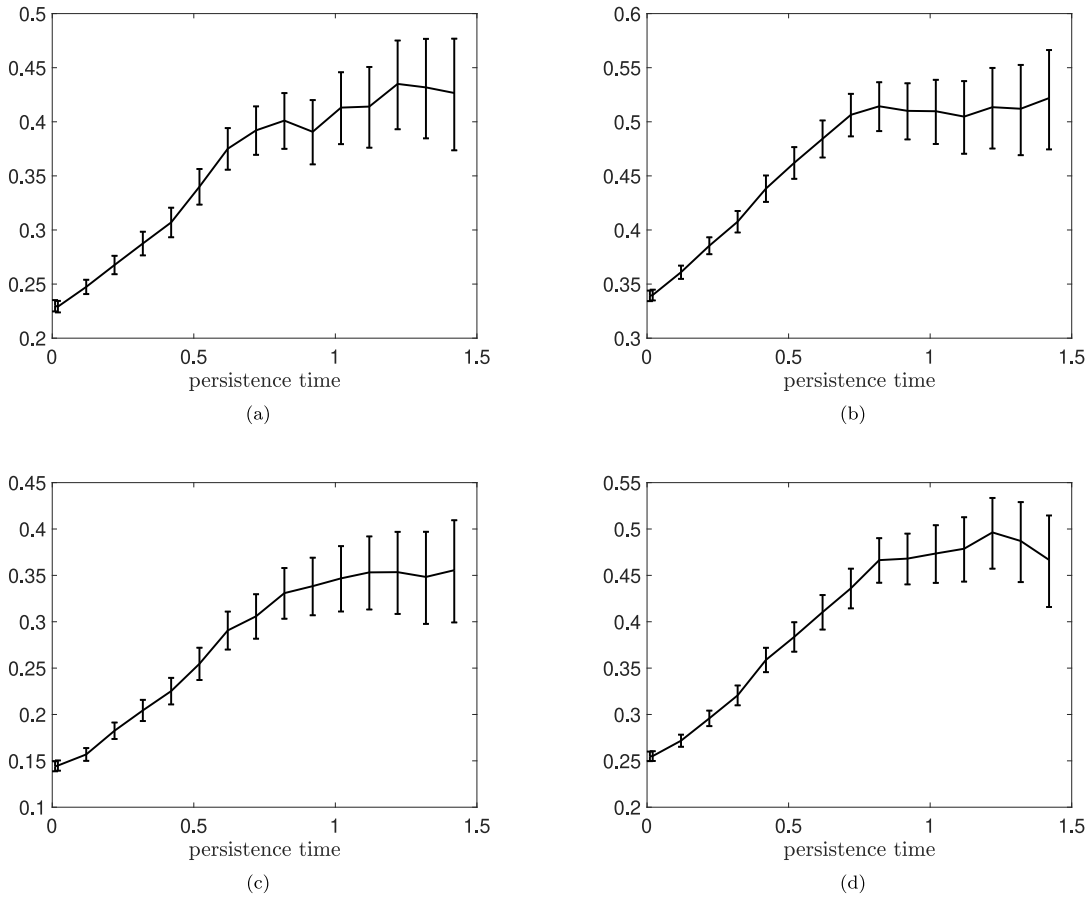

Fig. 8. Panel (a): correlation between the first FTLE λ_1^τ and λ_1 of the corresponding shadowing UPO as a function of the minimum persistence time of the UPO. Panel (b): correlation between h_{KS}^τ (sum of the first four FTLE) and h_{KS} of the corresponding shadowing UPO as a function of the minimum persistence time of the UPO. Panel (c): Same as (a), for the largest FTLE λ_{max}^τ. Panel (d): Same as (b), for the local Kolmogorov–Sinai entropy $h_{KS,+}^\tau$(sum of the positive FTLE). The error bar corresponds to the 95% confidence interval around the correlation value. The leftmost point of the diagrams corresponds to the values reported in Table 2 for $K = 1$.

LEs as the old one or not. If a change in the UD occurs, the trajectory goes through a glitch point, as in the case shown in Fig. 5, where at time t_1 the UDs number of the shadowing UPO decreases from 5 to 4. Indeed, the distribution of distances shown in Fig. 9(a) are statistically different at the 99% confidence level as determined by the Kolmogorov–Smirnov and the average distance between the trajectory and the shadowing UPO is larger than in the case of glitch point. Clearly, since we are focusing on transition points, the quality of the shadowing is in both cases slightly lower than on average.

In order to better characterise the transitions points, we investigate the properties of the UPOs belonging to the first tiers - i.e. the best local approximations to the trajectory - and check how similar they are in terms of number of UD. Fig. 9(b) shows that the variance in the number of UD of the UPOs across the first $K = 100$ tiers is higher (99.9% confidence level) when considering glitch rather than non-glitch transition points. Additionally, the variance of UD is much smaller away from the transitions points.

We propose the following interpretation. During each shadowing window, the chaotic trajectory is surrounded by UPOs that are dynamically similar, so that low variability in the number of UD is found. When the trajectory approaches a transition point where the shadowing UPO changes, the UPOs around the chaotic trajectory become less homogeneous and the variability in the number of UD increases. At a glitch point, the hetereogeneity of the UPOs is higher because the system is performing a transition between two qualitatively different regions of the phase space.

5. Coarse-grained dynamics via finite-state Markov chains

5.1. A statistical analysis of transition points

We take here a different angle for exploring the heterogeneity of the attractor of the system. We propose a coarse grained representation of the dynamics by constructing a finite state Markov chain process [100]. Our approach differs from previous analysis [101,102] because of the way we perform the partition of the phase space. Here, the neighbourhood of all UPOs with same number of UD are considered as a single state, and the mechanism of ranked shadowing outlined in Section 3.2 dictates the sequence of transitions from one state to another other.

We proceed as follows. We consider the $K = 1$ shadowing algorithm, which at each time-step t selects the UPO U_k of the full database that minimises the distance with the chaotic trajectory. We define the states $S = \{2, 3, 4, 5, 6, 7, 8\}$, where each $s_i \in S$ is the UDs number of the shadowing UPO. We say that the system is in state s_i at time t if s_i is the UDs number of the shadowing UPO at time t. Note that we have excluded the state corresponding to the few detected UPOs with 9 positive LEs because it makes little sense to include it in a statistical analysis. The stochastic variable $\varsigma : \{1, \dots, N_{max}\} \subset \mathcal{N} \to S$, describes the discrete Markov chain process as $\varsigma(t) = s_i$, and the stochastic matrix describing the process can then be inferred in a frequentist way as

$$P_{i,j}^{dt} = \frac{\#\{k : (\varsigma(k) = s_j) \wedge (\varsigma(k+1) = s_i)\}}{\#\{k : (\varsigma(k) = s_j)\}} \tag{4}$$

where # denotes the cardinality and the denominator ensures that the matrix is properly normalised.

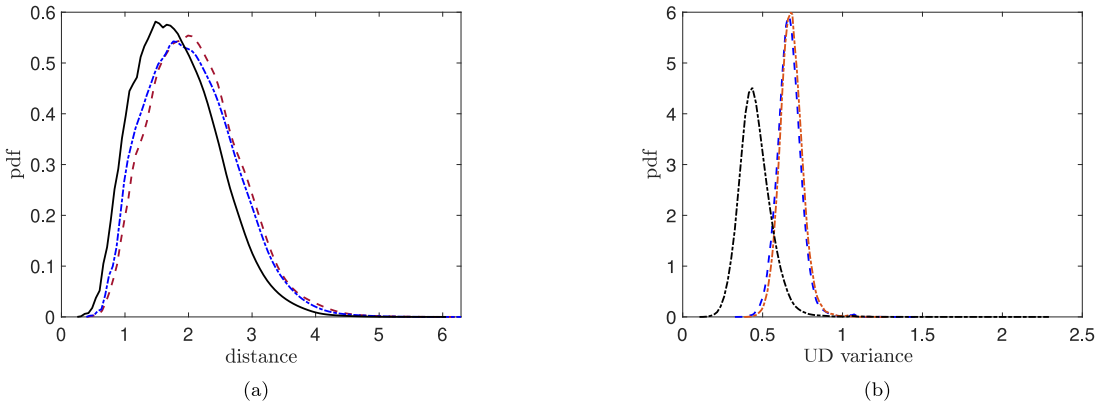

Fig. 9. Panel (a): Probability distribution function for the distance of the first tier orbit limited to transition points: dark red dashed line refers to transitions associated to glitch points (mean distance 2.11), blue dotted line refers to transitions between orbits with the same UDs number (mean distance 2.03). In black probability distribution function for the distance of the first tier orbit along the full chaotic trajectory (mean distance 1.81). Panel (b): Probability distribution function of the variance in the number of UD of the shadowing orbits within the first $K = 100$ tiers. In black variance across the entire chaotic trajectory (mean variance 0.45), in blue variance associated to transition points with same UDs number (non-glitch points) (mean variance 0.66), in dark red variance in the UDs number at glitch points (mean variance 0.67).

Through the study of the spectral properties of the stochastic matrix describing the process it is possible to obtain information on the diffusion properties of the system. In particular, under the assumption of ergodicity, the first eigenvector w_1 corresponding to the unitary eigenvalue $\nu_1 = 1$ determines the unique invariant measure that will be attained exponentially fast and independently of the initial ensemble. Each subdominant eigenvector w_k with corresponding eigenvalue ν_k sums to zero and describes a mode of the anomaly in the measure. The decay occurs on the timescale $\tau_k = -\frac{dt}{\ln \Re\{\nu^{(k)}\}}$, where $\Re\{c\}$ indicates the real part of the complex number c. In particular, the process of relaxation to the invariant measure from a generic initial condition is dominated by the longest timescale $\tau_2 = -\frac{dt}{\ln \Re\{\nu^{(2)}\}}$.

The first four subdominant eigenvalues are real. We have verified that P^{dt} is ergodic and tested the Markovianity of the process by verifying that the eigenvalues of the matrix $P_{i,j}^{dt\times n}$ obtained by sampling the shadowing every n time-steps has eigenvalues that scale with the nth power. Specifically, the relative normalised difference between the 10th power of the eigenvalues of $P_{i,j}$ and the eigenvalues of the stochastic matrix obtained sampling the process every 10 time-steps assumes a value of about 0.5%. We also verified that the eigenvectors are very similar to those of $P_{i,j}^{dt\times n}$ $w_k \cdot w_k^{dt\times n} \approx 1$ for $n = 10$ $\forall k$. We propose that such a coarse-grained representation of the dynamics characterises the statistics of the switching behaviour between clusters of UPOs with the same number of UD.

Fig. 10 shows the first five eigenvectors of P^{dt}, ordered according to the eigenvalue. The first eigenvector (in blue) returns the unique invariant measure of the system. Most often UPOs feature $UD = 4$ and $UD = 5$, in agreement with the distribution presented in Fig. 2(b). On the other hand, the modes corresponding to ν_2 (in orange, $\tau_2 = 0.3982$) and ν_3 (in yellow, $\tau_3 = 0.3134$) are responsible for the diffusion between anomalously stable and anomalously unstable UPOs. Indeed, both eigenvectors are characterised by negative components corresponding to $UD = 3$ and $UD = 4$, with positive values for larger UDs number. This means that if the initial ensemble is anomalously (un)stable, thus having an larger (smaller) number of members with high UDs number than dictated by the invariant measure, the anomaly will be damped mainly through these two modes.

The modes corresponding to ν_4 and ν_5 are responsible for the transitions between typical and atypical UPOs. In particular, the mode corresponding to ν_4 describes the diffusion from UPOs with $UD = 4$ to those with other values for UD. Finally, ν_5 describes the diffusion from UPOs with $UD = 5$ to those with other values for UD. It is also interesting to note that $\tau_4 = 0.2507 > \tau_5 = 0.2138$, in agreement with the fact that the dominant mode of the invariant measure is given by $UD = 4$.

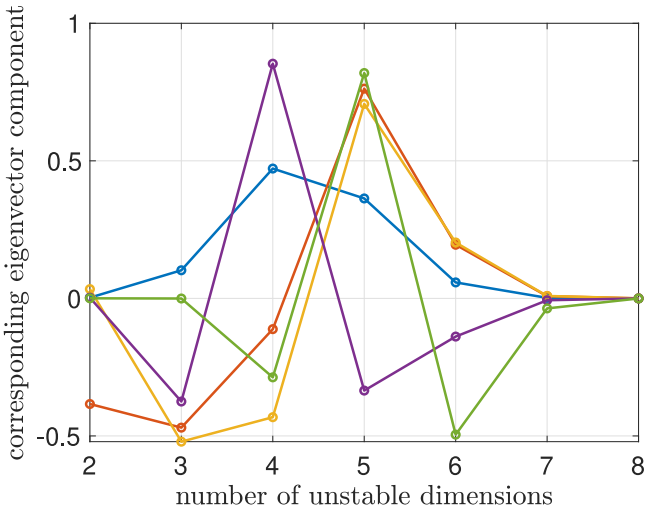

Fig. 10. Eigenvectors of the transition matrix P^{dt}. For each eigenvector (represented by different colours) we represent the value of its different components corresponding to the different states of the system. $\nu_1 = 1$ (in blue) returns the invariant measure of the system and it is in agreement with the distribution presented in Fig. 2(b). The subdominant eigenvalues are $\nu_2 = 0.9752$ in orange, $\nu_3 = 0.9686$ in yellow, $\nu_2 = 0.9608$ in purple and $\nu_2 = 0.9543$ in green.

Using Eq. (13) in [103], it is possible to evaluate the KS entropy h_{KS}^P of the finite state Markov chain described by the stochastic matrix $P_{i,j}^{dt}$. One obtains $h_{KS}^P \approx 0.31$. This figure clearly indicate that we are investigating a system that creates information in the Shannon sense. Nonetheless, as a result of the procedure of coarse graining used for defining the stochastic matrix, the estimate of the KS entropy is lower than the value one can directly derive for the full system $h_{KS} \approx 1.23$, which is given by the sum of the first 4 LEs shown in Fig. 1(a).

5.2. Mixing between large-scale regions of the attractor with different instability

We now wish to investigate the relaxation of an ensemble towards the invariant measure through a different coarse-graining procedure. Following [64], we construct a Markov chain where the states are the neighbourhoods of the shadowing UPOs. We define here the states $S = \{1, \dots, P\}$, where each $i \in S$ is index of the shadowing UPO. The system is in state i at time t if the ith UPO is shadowing the chaotic trajectory at time t. The stochastic variable $\varsigma : \{1, \dots, N_{max}\} \subset \mathcal{N} \to S$, describes

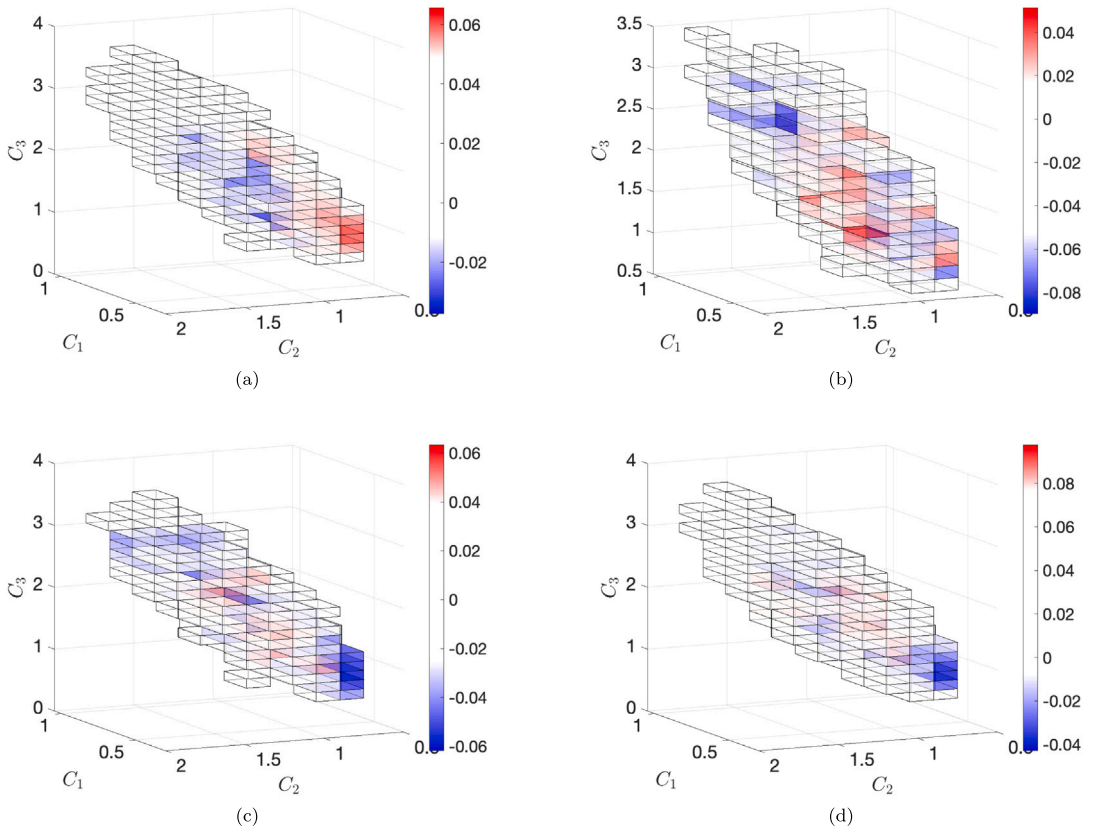

Fig. 11. Projection of the subdominant eigenvectors (panel (a)) w^2, (panel (b)) w^3, (panel (c)) w^4, (panel (d)) w^5. of the stochastic matrix $Q^{dt}_{i,j}$ (see Section 5) in the space defined by the normalised moments (C_1, C_2, C_3) (see Eq. (3)).

the discrete Markov chain process as $\varsigma(t) = i$, and the stochastic matrix describing the process can then be inferred in a frequentist way by performing the $K = 1$ shadowing using P UPOs and evaluating:

$$Q^{dt}_{i,j} = \frac{\#\{k : (\varsigma(k) = j) \wedge (\varsigma(k+1) = i)\}}{\#\{k : (\varsigma(k) = j)\}} \tag{5}$$

Ideally, one like to use the entire dataset of UPOs, so that $P = N_{UPO}$. Yet, given the very large number of UPOs considered in this study and the amount of available data, it is impossible to robustly evaluate the stochastic matrix by considering all UPOs. Hence, we choose as states of the system the neighbourhood of $P = 1000$ UPOs. Such UPOs have been randomly chosen within the whole dataset in such a way to respect the distribution of periods as in the original database. In order to enforce robustness, we also verify that each UPO has been selected in the shadowing for at least the length of its period. We remark that, given the protocol above, the results presented below are weakly dependent on the specific choice of UPOs for $P = 1000$ as well as on the chosen value for P.

We have verified that the process described by $Q^{dt}_{i,j}$ is ergodic, and tested its Markovianity (in particular, similarly to what we described earlier, we estimated an error of 0.4% when considering sampling every 10 time-steps. We also verified that the eigenvectors are very similar to those of $Q^{dt\times n}_{i,j}$ ($w_k \cdot w^{dt\times n}_k \approx 1$ within less than 1% for $n = 10\ \forall k$)). The first eigenvalues are $\nu_1 = 1, \nu_2 = 0.9887, \nu_3 = 0.9848, \nu_4 = 0.9844, \nu_5 = 0.9831$ and the corresponding characteristic decay times are $\tau_2 = 0.8771, \tau_3 = 0.6511, \tau_4 = 0.6348, \tau_5 = 0.5850$. The KS entropy production rate of the stochastic matrix $Q^{dt}_{i,j}$ is $h^Q_{ks} \approx 0.15$: also in this case, one finds clear evidence of chaotic behaviour but underestimates the rate of creation of information of the system.

We want to represent the eigenvectors of the stochastic matrix in the projected space spanned by the normalised moments C_1, C_2, and C_3 introduced in Eq. (3) and used in Fig. 4. We proceed as follows. We consider the compact set in $\mathbb{R}^3$ that contains the projected attractor of the system and take into account a partition in 5288 cubes of side 0.15. Each projected UPO of the system intersects a certain number of cubes, and each cube of the partition might contain contributions from different orbits. We measure the quantity of mass contained in each cube accordingly to the distribution given by the different subdominant eigenvectors. In order to do so, we set a fixed number of points N to be represented in phase space, and we distribute such points to each UPO accordingly to the weight given by the corresponding component of the subdominant eigenvector $w^{(k)}$. For each UPO we represent the selected points equally spaced in time in the phase space and distinguish between negative and positive contributions. We finally quantify the total amount of mass contained in each cube of the partition by computing the algebraic sum of the points contained in it.

Fig. 11 presents the outcome of such a procedure for the eigenvectors corresponding to the first four subdominant eigenvalues. By looking at pattern of positive and negative values, one realises that the eigenvector corresponding to the first subdominant eigenvalue ν_2 - see Fig. 11(a) - is responsible for the transfer of mass between regions characterised by low values of the moments and the core of the attractor, compare with Figs. 4(a) and 4(f). Considering that there is a clear relationship between the value of the moments and instability of the UPOs that preferentially populate the corresponding region of the attractor - see Figs. 4(b) and 4(d) - this is in good agreement with Fig. 10, where it is shown that the first subdominant eigenvector describes the transfer of mass between UPOs with low instability and those with typical instability. The eigenvector corresponding to ν_3 - see Fig. 11(b) describes the transfer of mass between anomalously high and anomalously low values of the moments, also in qualitative agreement with the results shown in Fig. 10. The eigenvectors corresponding to ν_4 and ν_5 - see Figs. 11(c) and 11(d), respectively - describe the transfer of

mass between typical and atypical (anomalously high and anomalously low) values of the moments. Also in this case a good agreement is found with the corresponding eigenvectors in Fig. 10.

We remark that if we perform this analysis taking as starting point the reduced set of $T \leq 6.4$ UPOs, unsurprisingly, the resulting subdominant eigenvector is rather different from the one depicted in Fig. 11(a) because low-period UPOs are unable to sample accurately the low-energy region of the attractor, see Fig. 3 in the Supplementary Material.

6. Conclusions

The use of Lyapunov analysis is well accepted as a key tool for understanding the dynamical and statistical properties of complex systems [22]. In this paper we suggest that this effort should be complemented by taking a global approach based on the study of the skeletal dynamics associated with the UPOs [26]. While their computation is a challenging task [31], UPOs indeed have a great potential for clarifying features and dynamical processes in complex systems [37,40,59,104].

As a step in the direction of better understanding the properties of chaos in high-dimensional non-hyperbolic systems, we have here investigated strong violations of hyperbolicity in a specific version of the Lorenz '96 model [68,69] with 20 degrees of freedom, taken as prototypical example of chaotic system with a nontrivial unstable manifold, i.e. with a dimensionality larger than one. We compute an extremely large set of UPOs immersed in the attractor, covering a large range of periods and geometrical lengths. While such a set is clearly incomplete, it allows for a rather accurate shadowing of the trajectory. We have verified that the system features UDV, where the UDs number ranges (at least) from 2 to 9. Fingerprints of the UDV are also found when performing a Lyapunov analysis of the system: several FTLEs fluctuate around zero also when very long averaging times are considered. Our results extend previous findings [43–45] because in our system we observe a much more pronounced UDV, as the heterogeneity of the UPOs is very high.

We have also found that the local stability properties of the tangent space, measured in terms of their FTLES, are highly correlated to those of the UPOs, measured by their corresponding LEs, populating the same region of the phase space. This is a nontrivial result if one considers the fact that UPOs describe global features on the attractor. The agreement becomes more evident as we focus our statistics on UPOs that shadow the trajectory for a long time.

We can rank the UPOs according to how close they are to the trajectory. While the closest orbit typically keeps its rank for a relatively short time, we find that often the forthcoming rank 1 orbit has the same UDs number as the previous one. The less frequent transitions to the neighbourhood of an UPO with a different UDs number – glitch points – are usually accompanied by a degradation of the quality of the shadowing.

On slower time scales, the systems performs transitions between regions of the attractor associated with anomalously high and low instability, which are preferentially populated by UPOs with large vs low UDs number, respectively. Following [64], we construct coarse-grained versions of the system based on using the neighbourhoods of the UPOs as building blocks and the dynamics as generator of the random process. It is then possible to see such transitions as slow relaxation processes of an arbitrary initial ensemble towards the invariant measure, with the subdominant modes of the transfer operator being associated with slow fluctuations between high instability and low instability states. High (low) energy is typically associated with high (low) instability, in agreement with the thermodynamical understanding of the model. Regions associated with anomalous energy and corresponding anomalous instability can be though of as corresponding, by and large, to quasi-invariant sets of the system [66].

In a (relatively) high-dimensional chaotic systems like the one considered here it is extremely hard if not practically impossible to compute all UPOs up to a given period T. The task of finding long-period UPOs becomes particularly daunting in the case of highly unstable orbits. Hence, our UPOs database is incomplete and biased. One might reasonably ask whether it is better to restrict our analysis to low-period UPOs, which are better sampled. Indeed, our results show the extreme importance of taking into account also the poorly sampled long-period UPOs. In fact, low-period UPOs cover poorly a specific region (the low-energy one) of the attractor. Hence, restricting our analysis to such orbits leads to a serious degradation of the quality of the shadowing and to losing key information on the dynamical and statistical properties of the system, in agreement with previous results [38,105].

While we have focused on a rather specific model, our findings wish to stimulate more general investigations on the link between the heterogeneity of the attractor and the breakdown of hyperbolicity. Such a heterogeneity has important implications in terms of robustness of the system, of the accuracy of the numerical models used to approximate its evolution, and of the efficacy of data assimilation strategies aimed at optimally merging observations and dynamics, as discussed in [59] in the context of geophysical flows.

There is a clear link between hyperbolicity of a system and applicability of linear response theory for computing the change in its statistical properties resulting from a perturbation to its dynamics [11,106]; see also more recent developments in [107,108]. Hyperbolicity is also critical for the definition of rigorous algorithms aimed at implementing Ruelle's response formulas and evaluating separately the two contributions to the overall response coming from the stable and unstable components of the tangent space [109–111]. The serious breakdown of hyperbolicity discussed in this paper could reasonably be seen as a very serious obstacle for the applicability of linear response formulas. Nonetheless, it has been amply shown that the linear response theory applies to a great degree of accuracy for the Lorenz '96 model, albeit in a slightly different configuration from the one considered here [89]. As discussed in [112], linear response theory is very effective in providing accurate climate change projections using climate models of arbitrary level of complexity. There is no reasons to believe that such models are anywhere close to hyperbolicity. In fact the opposite seems to hold when the multiscale effects of atmosphere-ocean coupling is taken into consideration [113].

The clarification of the apparent mismatch between what is rigorously guaranteed by mathematical theorems and what is heuristically shown in multiple applications regarding the applicability of linear response theory in complex systems seems a topic of great scientific relevance both on theoretical grounds and for reasons of practical significance.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The Supplementary Material of this publication, including extensive datasets and MATLAB codes, is available at [114].

Acknowledgements

The authors wish to thank P. Cvitanovic, J. Dorrington, D. Faranda, G. Pavliotis, and J. Yorke for constant encouragement. The authors also acknowledge the constructive criticisms of two anonymous reviewers. VL acknowledges the support received from the EPSRC project EP/T018178/1 and from the EU Horizon 2020 project TiPES (grant no. 820970). CCM has been supported by an EPSRC studentship as part of the Centre for Doctoral Training in Mathematics of Planet Earth (grant number EP/L016613/1). AG was supported by the Moscow Center of Fundamental and Applied Mathematics at INM RAS, RU (Agreement with the Ministry of Education and Science of the Russian Federation No. 075-15-2022-286). YS was supported by JSPS Grant-in-Aid for Scientific Research (B), JP No. 21H01002.

Appendix A. Mathematical framework

A.1. Chaotic dynamical systems

Let us consider a continuous-time autonomous dynamical system $\dot{x} = f(x)$ on a compact manifold $\mathcal{M} \subset \mathbb{R}^n$. We define a state at time $t \in \mathbb{R}_{t>0}$ as $x(t) = S^t x_0$, where $x_0 = x(0)$ is the initial condition and S^t is the evolution operator. We assume that the system is dissipative ($\nabla \cdot f < 0$) with $\Omega \subset M$ compact attracting invariant set. The attractor supports a probability measure ρ, invariant and ergodic with respect to S^t, specified in the following manner:

$$\langle \varphi \rangle = \int \rho(dx)\varphi(x) = \lim_{T\to\infty} \frac{1}{T} \int_0^T \varphi(S^t x_0)dt \tag{6}$$

for any sufficiently regular function (observable) $\varphi : M \to \mathbb{R}$ and for almost all initial conditions x_0 belonging to the basin of attraction of Ω.

A.2. Lyapunov exponents and instability of the flow

We now want to characterise dynamical systems that feature sensitive dependence on initial conditions on the attractor. Following Pesin's theory [13], this is most easily accomplished by using Lyapunov analysis [21,22]. The separation in time of infinitesimally nearby trajectories can be asymptotically quantified through specific global dynamical indicators [22]. We define the matrix

$$\Lambda(x) = \lim_{t\to\infty} (J_t^\top(x) J_t(x))^{\frac{1}{2t}} \tag{7}$$

where $J_t(x) = \nabla_x f(S^t x)$ is the tangent linear matrix of the flow at time t with initial condition $x \in \Omega$ and $\top$ indicates the transpose. It is possible to prove that the matrix $\Lambda(x)$ exists and its eigenvalues Λ_i are real and constant almost everywhere with respect to the measure ρ. We call Lyapunov exponents (LE) of the systems the objects defined as $\lambda_i = \log(\Lambda_i)$. Usually they are ordered by size in descending order $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ and a positive first LE λ_1 indicates that the system is chaotic. LEs are asymptotic quantities and refer to average properties over the attractor. The number of positive LEs is the number of UDs. The full Lyapunov spectrum allows to define the Kaplan–Yorke dimension, which is conjectured to correspond to the order 1 Renyi dimension [115] of the attractor as follows:

$$D_{KY} = m + \frac{\sum_{i=1}^m \lambda_i}{|\lambda_{m+1}|} \tag{8}$$

with m being the highest index for which the sum of the largest m LEs is strictly positive. The quantity D_{KY} can be thought as an approximate value of the number of excited degrees of freedom acting in the system [116]. The degree of chaoticity of a dynamical system can also be quantified via the Kolmogorov–Sinai entropy (approximated via Pesin's theorem in the case the invariant measure is of the Sinai–Ruelle–Bowen type) [16] as:

$$h_{KS} = \sum_{i=1}^n \lambda_i \tag{9}$$

where n is the number of positive LEs, which corresponds to the UDs number. The quantity h_{KS} quantifies the production of information of the system.

The local instability properties of the attractor are described by the Finite Time Lyapunov Exponents (FTLE). These quantify the amount of stretching about the trajectory with initial condition x on the attractor over a finite time interval $[0, t]$. They are local objects since their value depends on x and t. They can be computed as the logarithm of the eigenvalues of the matrix $\Lambda(x,t) = (J_t^T(x)J_t(x))^{1/(2t)}$. One can also define as $h^t_{KS} = \sum_{i=1}^n \lambda_i^t$. The value of the jth FTLE can vary a lot along the attractor and in some cases, for a given t, the support of its probability distribution can include zero, meaning that the local FTLE can have a different sign compared to the corresponding jth LE. Correspondingly, the number of positive FTLEs can in general fluctuate and be different from n. Note also that $\lambda_1^t \leq \lambda_{max}^t$, where the latter is the largest among the λ_j^t's and defines the local largest growth rate. Finally, one can define a local version of the Kolmogorov–Sinai entropy as $h^t_{KS,+} = \sum_{\{\lambda_i^t>0\}} \lambda_i^t \geq h^t_{KS}$.

A.3. Unstable periodic orbits

A periodic orbit of period T is an exact periodic solution of the evolution equation and it is defined as follows:

$$S^T(x) = x. \tag{10}$$

This representation is not unique. In fact, if Eq. (10) is satisfied, $S^{nT}(x) = x$ is verified as well $\forall n \in \mathbb{N}$. In this work we will consider a periodic orbit to be identified by its prime period $T > 0$ (we do not consider equilibria) and an initial condition x_0. The attractor of a chaotic dynamical systems is densely populated by unstable periodic orbits [28]. It is known that a good understanding of the UPOs of the model plays a fundamental role in the characterisation of the system. As an example, it is possible to reconstruct the invariant measure of the system through the use of trace formulas, by considering the following expression for the average of any measurable observable φ:

$$\langle \varphi \rangle = \lim_{t\to\infty} \frac{\sum_{U^p, p\leq t} w^{U^p} \bar{\varphi}^{U^p}}{\sum_{U^p, p\leq t} w^{U^p}} \tag{11}$$

where U^p is a UPO of prime period p, w^{U^p} is its weight and $\bar{\varphi}^{U^p}$ is the average in time of the observable along the orbit. For uniformly hyperbolic dynamical systems this result is exact and the weight can be obtained, to a first approximation, by $w^{U^p} \propto \exp(-p h_{ks}^{U^p})$ [117], with h_{ks} being the Kolmogorov–Sinai entropy of the system; see also discussion in [39].

Appendix B. Supplementary material

The full address is the following: http://dx.doi.org/10.6084/m9.figshare.24242458.v1. The supplementary material contains text and datasets.

References

[1] H. Poincaré, Science et Méthode, Flammarion, Paris, 1908.
[2] D. Ruelle, Henri poincaré's science et méthode, Vol. S88, Publications Mathématiques de l'IHÉS, 1998, pp. 179–181.
[3] E.N. Lorenz, Deterministic nonperiodic flow, J. Atmos. Sci. 20 (2) (1963) 130–141.
[4] D. Ruelle, F. Takens, On the nature of turbulence, Comm. Math. Phys. 20 (3) (1971) 167–192.
[5] T.-Y. Li, J.A. Yorke, Period three implies chaos, Amer. Math. Monthly 82 (10) (1975) 985–992.
[6] A. Hammerlindl, B. Krauskopf, G. Mason, H.M. Osinga, Global manifold structure of a continuous-time heterodimensional cycle, 2019, arXiv preprint arXiv:1906.11438.
[7] S. Smale, et al., Differentiable dynamical systems, Bull. Am. Math. Soc. 73 (6) (1967) 747–817.
[8] D.V. Anosov, Geodesic flows on closed Riemannian manifolds of negative curvature, Tr. Mat. Inst. Imeni VA Steklova 90 (1967) 3–210.
[9] A. Katok, B. Hasselblatt, Introduction to the modern theory of dynamical systems, in: Encyclopedia of Mathematics and Its Applications, Cambridge University Press, 1995.
[10] D. Ruelle, Chaotic Evolution and Strange Attractors: The Statistical Analysis of Time Series for Deterministic Nonlinear Systems, Cambridge University Press, Lezioni Lincee, Cambridge, England, 1989.
[11] D. Ruelle, General linear response formula in statistical mechanics, and the fluctuation–dissipation theorem far from equilibrium, Phys. Lett. A 245 (3–4) (1998) 220–224.
[12] C. Bonatti, Survey towards a global view of dynamical systems, for the c1-topology, Ergodic Theory Dynam. Systems 31 (4) (2011) 959–993.
[13] Y.B. Pesin, Characteristic Lyapunov exponents and smooth ergodic theory, Russian Math. Surveys 32 (1977) 55–114.

[14] G. Benettin, L. Galgani, A. Giorgilli, J.-M. Strelcyn, Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. part 2: Numerical application, Meccanica 15 (1) (1980) 21–30.

[15] L.-S. Young, Mathematical theory of Lyapunov exponents, J. Phys. A 46 (2013) 254001.

[16] J.-P. Eckmann, D. Ruelle, Ergodic theory of chaos and strange attractors, Theory Chaotic Attract. (1985) 273–312.

[17] J.M. Nese, Quantifying local predictability in phase space, Physica D 35 (1) (1989) 237–250.

[18] H.D.I. Abarbanel, R. Brown, M.B. Kennel, Variation of Lyapunov exponents on a strange attractor, J. Nonlinear Sci. 1 (2) (1991) 175–199.

[19] D. Gallez, A. Babloyantz, Lyapunov exponents for nonuniform attractors, Phys. Lett. A 161 (3) (1991) 247–254.

[20] E. Aurell, G. Boffetta, A. Crisanti, G. Paladin, A. Vulpiani, Predictability in the large: an extension of the concept of Lyapunov exponent, J. Phys. A: Math. Gen. 30 (1997) 1–26.

[21] M. Cencini, F. Ginelli, Lyapunov analysis: from dynamical systems theory to applications, J. Phys. A 46 (2013) 250301.

[22] A. Pikovsky, A. Politi, Lyapunov Exponents: A Tool to Explore Complex Dynamics, Cambridge University Press, 2016.

[23] L. Barreira, Y. Pesin, Nonuniform hyperbolicity: dynamics of systems with nonzero Lyapunov exponents, in: Encyclopedia of Mathematics and Its Applications, Cambridge University Press, 2007.

[24] C. Bonatti, L.J. Díaz, M. Viana, Dynamics beyond Uniform Hyperbolicity: A Global Geometric and Probabilistic Perspective, Vol. 3, Springer Science & Business Media, 2004.

[25] W. Zhang, B. Krauskopf, V. Kirk, How to find a codimension-one heteroclinic cycle between two periodic orbits, Discrete Contin. Dyn. Syst. 32 (8) (2012) 2825.

[26] P. Cvitanović, Invariant measurement of strange sets in terms of cycles, Phys. Rev. Lett. 61 (24) (1988) 2729.

[27] P. Cvitanović, Periodic orbits as the skeleton of classical and quantum chaos, Physica D 51 (1–3) (1991) 138–151.

[28] P. Cvitanovic, R. Artuso, R. Mainieri, G. Tanner, G. Vattay, N. Whelan, A. Wirzba, Chaos: Classical and Quantum, Vol. 69, ChaosBook. org (Niels Bohr Institute, Copenhagen 2005, 2005, p. 25.

[29] P. Gaspard, Chaos, scattering and statistical mechanics, Chaos (2005).

[30] E. Kazantsev, Unstable periodic orbits and attractor of the barotropic ocean model, Nonlinear Process. Geophys. 5 (4) (1998) 193–208.

[31] G.J. Chandler, R.R. Kerswell, Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow, J. Fluid Mech. 722 (2013) 554–595.

[32] D. Lucas, R.R. Kerswell, Recurrent flow analysis in spatiotemporally chaotic 2-dimensional Kolmogorov flow, Phys. Fluids 27 (4) (2015) 045106.

[33] L. Van Veen, S. Kida, G. Kawahara, Periodic motion representing isotropic turbulence, Fluid Dyn. Res. 38 (1) (2006) 19.

[34] M.C. Krygier, J.L. Pughe-Sanford, R.O. Grigoriev, Exact coherent structures and shadowing in turbulent Taylor-Couette flow, 2021, arXiv preprint arXiv:2105.02126.

[35] P. Cvitanović, J. Gibson, Geometry of the turbulence in wall-bounded shear flows: periodic orbits, Phys. Scr. 2010 (T142) (2010) 014007.

[36] T. Kreilos, B. Eckhardt, Periodic orbits near onset of chaos in plane couette flow, Chaos 22 (4) (2012) 047505.

[37] B. Suri, L. Kageorge, R.O. Grigoriev, M.F. Schatz, Capturing turbulent dynamics and statistics in experiments with unstable periodic orbits, Phys. Rev. Lett. 125 (6) (2020) 064501.

[38] G. Yalniz, N.B. Budanur, Inferring symbolic dynamics of chaotic flows from persistence, Chaos 30 (3) (2020) 033109.

[39] M. Dhamala, Y.-C. Lai, Unstable periodic orbits and the natural measure of nonhyperbolic chaotic saddles, Phys. Rev. E 60 (1999) 6176–6179.

[40] G. Kawahara, S. Kida, Periodic motion embedded in plane couette turbulence: regeneration cycle and burst, J. Fluid Mech. 449 (2001) 291.

[41] J. Page, P. Norgaard, M.P. Brenner, R.R. Kerswell, Recurrent flow patterns as a basis for turbulence: predicting statistics from structures, 2022, arXiv preprint arXiv:2212.11886.

[42] Y.-C. Lai, Y. Nagai, C. Grebogi, Characterization of the natural measure by unstable periodic orbits in chaotic attractors, Phys. Rev. Lett. 79 (1997) 649–652.

[43] T. Sauer, C. Grebogi, J.A. Yorke, How long do numerical chaotic solutions remain valid? Phys. Rev. Lett. 79 (1997) 59–62.

[44] T.D. Sauer, Shadowing breakdown and large errors in dynamical simulations of physical systems, Phys. Rev. E 65 (2002) 036220.

[45] R.F. Pereira, S.E. de S. Pinto, R.L. Viana, S.R. Lopes, C. Grebogi, Periodic orbit analysis at the onset of the unstable dimension variability and at the blowout bifurcation, Chaos 17 (2) (2007) 023131.

[46] E. Kalnay, Atmospheric Modeling, Data Assimilation and Predictability, Cambridge University Press, Cambridge, 2003.

[47] T.N. Palmer, Predicting uncertainty in forecasts of weather and climate, Rep. Progr. Phys. 63 (2000) 71–116.

[48] J. Slingo, T. Palmer, Uncertainty in weather and climate prediction, Phil. Trans. R. Soc. A 369 (1956) (2011) 4751–4767.

[49] A. Carrassi, M. Bocquet, L. Bertino, G. Evensen, Data assimilation in the geosciences: An overview of methods, issues, and perspectives, WIREs Clim. Change 9 (5) (2018) e535.

[50] Y. Wu, Z. Shen, Y. Tang, A flow-dependent targeted observation method for ensemble Kalman filter assimilation systems, Earth Space Sci. 7 (7) (2020) e2020EA001149, e2020EA001149 2020EA001149.

[51] Y. Chen, A. Carrassi, V. Lucarini, Inferring the instability of a dynamical system from the skill of data assimilation exercises, Nonlinear Process. Geophys. 28 (4) (2021) 633–649.

[52] J.M. Nese, J.A. Dutton, Quantifying predictability variations in a low-order ocean-atmosphere model: A dynamical systems approach, J. Clim. 6 (2) (1993) 185–204.

[53] S. Yoden, M. Nomura, Finite-time Lyapunov stability analysis and its application to atmospheric predictability, J. Atmos. Sci. 50 (11) (1993) 1531–1543.

[54] C. Nicolis, S. Vannitsem, J.-F. Royer, Short-range predictability of the atmosphere: Mechanisms for superexponential error growth, Q. J. R. Meteorol. Soc. 121 (523) (1995) 705–722.

[55] S. Vannitsem, C. Nicolis, Lyapunov vectors and error growth patterns in a t21l3 quasigeostrophic model, J. Atmos. Sci. 54 (2) (1997) 347–361.

[56] L. De Cruz, S. Schubert, J. Demaeyer, V. Lucarini, S. Vannitsem, Exploring the Lyapunov instability properties of high-dimensional atmospheric and climate models, Nonlinear Process. Geophys. 25 (2) (2018) 387–412.

[57] D. Pazó, J.M. López, A. Politi, Universal scaling of Lyapunov-exponent fluctuations in space–time chaos, Phys. Rev. E 87 (2013) 062909.

[58] T. Laffargue, K.-D.N.T. Lam, J. Kurchan, J. Tailleur, Large deviations of Lyapunov exponents, J. Phys. A 46 (2013) 254002.

[59] V. Lucarini, A. Gritsun, A new mathematical framework for atmospheric blocking events, Clim. Dynam. 54 (1) (2020) 575–598.

[60] A. Gritsun, Unstable periodic trajectories of a barotropic model of the atmosphere, Russian J. Numer. Anal. Math. Modelling 23 (4) (2008).

[61] A. Gritsun, Statistical characteristics, circulation regimes and unstable periodic orbits of a barotropic atmospheric model, Phil. Trans. R. Soc. A 371 (1991) (2013) 20120336.

[62] S. Schubert, V. Lucarini, Dynamical analysis of blocking events: spatial and temporal fluctuations of covariant Lyapunov vectors, Q. J. R. Meteorol. Soc. 142 (698) (2016) 2143–2158.

[63] D. Faranda, G. Messori, P. Yiou, Dynamical proxies of north atlantic predictability and extremes, Sci. Rep. 7 (1) (2017) 41278.

[64] C.C. Maiocchi, V. Lucarini, A. Gritsun, Decomposing the dynamics of the Lorenz 1963 model using unstable periodic orbits: Averages, transitions, and quasi-invariant sets, Chaos 32 (3) (2022) 033129.

[65] R. Barrio, A. Dena, W. Tucker, A database of rigorous and high-precision periodic orbits of the Lorenz model, Comput. Phys. Comm. 194 (2015) 76–83.

[66] G. Froyl, K. Padberg-Gehle, Almost-invariant and finite-time coherent sets: directionality, duration, and diffusion, in: Ergodic Theory, Open Dynamics, and Coherent Structures, Springer, 2014, pp. 171–216.

[67] L.A. Smith, C. Ziehmann, K. Fraedrich, Uncertainty dynamics and predictability in chaotic systems, Q. J. R. Meteorol. Soc. 125 (560) (1999) 2855–2886.

[68] E.N. Lorenz, Predictability: A problem partly solved, in: Proc. Seminar on Predictability, Vol. 1, 1996.

[69] E.N. Lorenz, Designing chaotic models, J. Atmos. Sci. 62 (2005) 1574–1587.

[70] J. Miller, J. Yorke, Finding all periodic orbits of maps using newton methods: sizes of basins, Physica D 135 (3–4) (2000) 195–211.

[71] D.L. van Kekem, A.E. Sterk, Travelling waves and their bifurcations in the Lorenz-96 model, Physica D 367 (2018) 38–60.

[72] D.L. van Kekem, A.E. Sterk, Wave propagation in the Lorenz-96 model, Nonlinear Process. Geophys. 25 (2) (2018) 301–314.

[73] D. Wilks, Effects of stochastic parametrizations in the Lorenz '96 system, Q. J. R. Meteorol. Soc. 131 (606) (2005) 389–407.

[74] H.M. Arnold, I.M. Moroz, T.N. Palmer, Stochastic parametrizations and model uncertainty in the Lorenz system, Phil. Trans. R. Soc. A 371 (1991) (2013) 20110479.

[75] G. Vissio, V. Lucarini, A proof of concept for scale-adaptive parametrizations: the case of the Lorenz '96 model, Q. J. R. Meteorol. Soc. 144 (2018) 63–75.

[76] A. Chattopadhyay, P. Hassanzadeh, D. Subramanian, Data-driven predictions of a multiscale Lorenz 96 chaotic system using machine-learning methods: Reservoir computing, artificial neural network, and long short-term memory network, Nonlinear Process. Geophys. 27 (3) (2020) 373–389.

[77] I.I. D. J. Gagne, H.M. Christensen, A.C. Subramanian, A.H. Monahan, Machine learning for stochastic parameterization: Generative adversarial networks in the Lorenz '96 model, J. Adv. Modelling Earth Syst. 12 (3) (2020) http://dx.doi.org/10.1029/2019MS001896, e2019MS001896. e2019MS001896.

[78] M. Gelbrecht, V. Lucarini, N. Boers, J. Kurths, Analysis of a bistable climate toy model with physics-based machine learning methods, Eur. Phys. J. Spec. Top. 230 (14) (2021) 3121–3131.

[79] R. Blender, V. Lucarini, Nambu representation of an extended Lorenz model with viscous heating, Physica D 243 (1) (2013) 86–91.

[80] A.E. Sterk, D.L. van Kekem, Predictability of extreme waves in the Lorenz-96 model near intermittency and quasi-periodicity, Complexity 2017 (2017) 9419024.
[81] G. Hu, T. Bódai, V. Lucarini, Effects of stochastic parametrization on extreme value statistics, Chaos 29 (8) (2019) 083102.
[82] A. Trevisan, F. Uboldi, Assimilation of standard and targeted observations within the unstable subspace of the observation–analysis–forecast cycle system, J. Atmos. Sci. 61 (1) (2004) 103–113.
[83] J. Brajard, A. Carrassi, M. Bocquet, L. Bertino, Combining data assimilation and machine learning to emulate a dynamical model from sparse and noisy observations: A case study with the Lorenz 96 model, J. Comput. Sci. 44 (2020) 101171.
[84] D.S. Wilks, Comparison of ensemble-mos methods in the Lorenz '96 setting, Meteorol. Appl. 13 (3) (2006) 243–256.
[85] W. Duan, Z. Huo, An approach to generating mutually independent initial perturbations for ensemble forecasts: Orthogonal conditional nonlinear optimal perturbations, J. Atmos. Sci. 73 (2016) 997–1014.
[86] S. Hallerberg, D. Pazo, J.M. Lopez, M.A. Rodriguez, Logarithmic bred vectors in spatiotemporal chaos: Structure and growth, Phys. Rev. E 81 (6) (2010) 066204.
[87] M. Carlu, F. Ginelli, V. Lucarini, A. Politi, Lyapunov analysis of multiscale dynamics: the slow bundle of the two-scale Lorenz 96 model, Nonlinear Process. Geophys. 26 (2) (2019) 73–89.
[88] R.V. Abramov, A.J. Majda, New approximations and tests of linear fluctuation-response for chaotic nonlinear forced-dissipative dynamical systems, J. Nonlinear Sci. 18 (3) (2008) 303–341.
[89] V. Lucarini, S. Sarno, A statistical mechanical approach for the computation of the climatic response to general forcings, Nonlinear Process. Geophys. 18 (1) (2011) 7–28.
[90] V. Lucarini, Stochastic perturbations to dynamical systems: A response theory approach, J. Stat. Phys. 146 (4) (2012) 774–786.
[91] G. Gallavotti, V. Lucarini, Equivalence of non-equilibrium ensembles and representation of friction in turbulent flows: the Lorenz 96 model, J. Stat. Phys. 156 (6) (2014) 1027–1065.
[92] G. Vissio, V. Lucarini, Mechanics and thermodynamics of a new minimal model of the atmosphere, Eur. Phys. J. Plus 135 (10) (2020) 807.
[93] P. Grassberger, I. Procaccia, Measuring the strangeness of strange attractors, Physica D 9 (1) (1983) 189–208.
[94] Y. Saiki, Numerical detection of unstable periodic orbits in continuous-time dynamical systems with chaotic behaviors, Nonlinear Process. Geophys. 14 (5) (2007) 615–620.
[95] R. Barrio, A. Dena, W. Tucker, A database of rigorous and high-precision periodic orbits of the Lorenz model, Comput. Phys. Comm. 194 (2015) 76–83.
[96] M.C. Krygier, J.L. Pughe-Sanford, R.O. Grigoriev, Exact coherent structures and shadowing in turbulent Taylor–Couette flow, J. Fluid Mech. 923 (2021) A7.
[97] S. Dawson, C. Grebogi, T. Sauer, J.A. Yorke, Obstructions to shadowing when a Lyapunov exponent fluctuates about zero, Phys. Rev. Lett. 73 (14) (1994) 1927.
[98] R. Bowen, ω-limit sets for axiom a diffeomorphisms, J. Differential Equations 18 (2) (1975) 333–339.
[99] T. Sauer, C. Grebogi, J.A. Yorke, How long do numerical chaotic solutions remain valid? Phys. Rev. Lett. 79 (1) (1997) 59.
[100] R. Livi, P. Politi, Nonequilibrium Statistical Physics: A Modern Perspective, Cambridge University Press, 2017.
[101] G. Froyland, Extracting dynamical behavior via Markov models, in: Nonlinear Dynamics and Statistics, Springer, 2001, pp. 281–321.
[102] G. Froyl, M. Dellnitz, Detecting and locating near-optimal almost-invariant sets and cycles, SIAM J. Sci. Comput. 24 (6) (2003) 1839–1863.
[103] P. Gaspard, Time-reversed dynamical entropy and irreversibility in Markovian random processes, J. Stat. Phys. 117 (3) (2004) 599–615.
[104] P. Cvitanović, Recurrent flows: The clockwork behind turbulence, J. Fluid Mech. 726 (2013) 1–4.
[105] D. Lasagna, Sensitivity of long periodic orbits of chaotic systems, Phys. Rev. E 102 (5) (2020) 052220.
[106] D. Ruelle, A review of linear response theory for general differentiable dynamical systems, Nonlinearity 22 (4) (2009) 855.
[107] V. Baladi, Linear response despite critical points, Nonlinearity 21 (6) (2008) T81.
[108] V. Baladi, Linear response, or else, in: ICM Seoul 2014, Proceedings, Vol. III, 2014, pp. 525–545.
[109] A. Ni, Approximating linear response by nonintrusive shadowing algorithms, SIAM J. Numer. Anal. 59 (6) (2021) 2843–2865.
[110] N. Chandramoorthy, Q. Wang, A computable realization of Ruelle's formula for linear response of statistics in chaotic systems, 2020, arXiv e-prints arXiv: 2002.04117.
[111] A.A. Śliwiak, Q. Wang, A trajectory-driven algorithm for differentiating SRB measures on unstable manifolds, SIAM J. Sci. Comput. 44 (1) (2022) A312–A336.
[112] M. Ghil, V. Lucarini, The physics of climate variability and climate change, Rev. Modern Phys. 92 (2020) 035002.
[113] S. Vannitsem, V. Lucarini, Statistical and dynamical properties of covariant Lyapunov vectors in a coupled atmosphere-ocean model—multiscale effects, geometric degeneracy, and error dynamics, J. Phys. A 49 (22) (2016) 224001.
[114] V. Lucarini, C.C. Maiocchi, A. Gritsun, Y. Sato, Supplementary material for the article heterogeneity of the attractor of the Lorenz '96 model: Bridging the gap between Lyapunov analysis and unstable periodic orbits, 2023, http://dx.doi.org/10.6084/m9.figshare.24242458.v1.
[115] E. Ott, Chaos in Dynamical Systems, Cambridge University Press, 2002.
[116] P. Frederickson, J.L. Kaplan, E.D. Yorke, J.A. Yorke, The Liapunov dimension of strange attractors, J. Differential Equations 49 (2) (1983) 185–207.
[117] C. Grebogi, E. Ott, J.A. Yorke, Unstable periodic orbits and the dimensions of multifractal chaotic attractors, Phys. Rev. A 37 (5) (1988) 1711.