

Sensitivity Analysis of Chaotic Systems using Unstable Periodic Orbits*

Davide Lasagna[†]

Abstract. A well-behaved adjoint sensitivity technique for chaotic dynamical systems is presented. The method arises from the specialisation of established variational techniques to the unstable periodic orbits of the system. On such trajectories, the adjoint problem becomes a time periodic boundary value problem. The adjoint solution remains bounded in time and does not exhibit the typical unbounded exponential growth observed using traditional methods over unstable non-periodic trajectories (Lea et al., Tellus 52 (2000)). This enables the sensitivity of period averaged quantities to be calculated *exactly*, regardless of the orbit period, because the stability of the tangent dynamics is decoupled effectively from the sensitivity calculations. We demonstrate the method on two prototypical systems, the Lorenz equations at standard parameters and the Kuramoto-Sivashinky equation, a one-dimensional partial differential equation with chaotic behaviour. We report a statistical analysis of the sensitivity of these two systems based on databases of unstable periodic orbits of size $\sim 10^5$ and $\sim 4 \times 10^4$, respectively. The empirical observation is that most orbits predict approximately the same sensitivity. The effects of symmetries, bifurcations and intermittency are discussed and future work is outlined in the conclusions.

Key words. adjoint method, sensitivity analysis, chaotic systems, unstable periodic orbits, symmetry, Lorenz, Kuramoto-Sivashinky, boundary value problem

1. Introduction. Sensitivity analysis refers broadly to mathematical techniques that, applied to a given system, provide the gradient of an output quantity of interest with respect to design variables parametrising the system at hand. There is a large range of problems across the physical and engineering sciences where sensitivity analysis techniques are tremendously useful: gradient-based optimisation, inverse problems, control design, model calibration, analysis and reduction or uncertainty quantification.

For time dependent problems, it is common for the output quantity to be the time average of an observable over a finite reference trajectory, generated by a parametrised dynamical system as the solution of an initial value problem from some given initial condition. The interest is typically on the asymptotic state of this trajectory and how parameters affect the time average of the observable. Traditional sensitivity techniques, both ‘tangent’ and ‘adjoint’ methods, attempt to calculate the *sensitivity of the time average* with respect to the parameters, the sensitivity of the asymptotic state, by calculating the *average of the sensitivity*, i.e. by time averaging the solution of a linear initial value problem obtained from linearisation of the dynamics around the reference trajectory [34].

For chaotic dynamical systems, however, the initial value problem determining the evolution of a trajectory from a given initial condition is inherently ill-conditioned [39]. Small perturbations to the system parameters have exponentially growing effects on the *future* evolution of unstable trajectories. In such settings, the instability of the tangent dynamics is inextricably intertwined to the sensitivity calculations. Linearised equations, both tangent

*In TSFP10 conference proceedings at http://tsfp10.org/TSFP10_program/2/156.pdf

[†]Faculty of Engineering and Environment, University of Southampton, University Road, SO17 1BJ, Southampton, United Kingdom (davide.lasagna@soton.ac.uk).

and adjoint, do not reproduce the inevitable saturation of exponentially growing modes that will occur later on at some point in time. The resulting *average of the sensitivity* leads to unphysically large gradients and does not coincide in the limit with the required *sensitivity of the time average* [25, 11, 34, 39].

In this paper, we aim to show that these well known limitations can be completely overcome if time periodic trajectories, i.e. unstable periodic orbits (UPOs), are used as reference trajectories for sensitivity analysis, rather than using chaotic non-periodic trajectories. For a time periodic trajectory, it is natural to drop the concept of initial conditions and the causality associated with the forward direction of time. As a result, it is more proficuous to think of UPOs as solutions of a *nonlinear time periodic boundary value problem*, i.e. as trajectories satisfying at each point the governing equations *plus* a periodicity constraint, rather than a simple initial condition. Crucially, this boundary value problem is well-conditioned: excluding pathological cases, i.e. bifurcations, small parameter perturbations have a bounded effect on the periodic solutions, *regardless of their stability characteristics and length*. In other words, small parameter perturbations do not result in exponentially diverging trajectories, but in a bounded state space displacement/distortion of an orbit that remains periodic in this process. The linearisation of the dynamics around a periodic trajectory leads to a *linear time periodic boundary problem* that does not admit exponentially growing modes but only a periodic solution. Stability and sensitivity are effectively decoupled and the *average of the sensitivity* coincides with the *sensitivity of the average*, the sensitivity of the asymptotic (periodic) state. This approach shares similar perspectives, discussed in the paper, to those underpinning the Least Squares Shadowing (LSS) method [40], recently introduced to circumvent the above mentioned limitations of traditional sensitivity approaches on open trajectories.

There are also other deeper system theoretical reasons that make UPOs attracting. For instance, Periodic Orbit Theory introduces appropriate weighted sums over UPOs to express ergodic statistics of hyperbolic systems [3], although, it appears that there is no established theory that expresses the sensitivity of the asymptotic state from the sensitivity of UPOs. In this paper, however, we do not aim to deal with fundamental aspects of Periodic Orbit Theory. Rather, our aim is to show that a well-behaved sensitivity technique can be formulated on unstable time periodic trajectories. The methods that we outline in this paper are analogous to variational methods for time periodic systems previously reported in the literature, for the tangent approach [38, 42] and for the adjoint approach [12, 19, 43]. However, in the vast majority of cases these methods have been discussed, for a variety of reasons, in the context of problems where the periodic trajectory is *stable*. These formalisms, though, apply regardless of the stability of the orbit, although more care must be taken in the numerical solution.

The main contributions of this paper are thus: the discussion of tangent and adjoint sensitivity techniques time periodic trajectories and associated numerical algorithms, the discussion of connections to the above-mentioned limitation of traditional initial value approaches and the discussion of potential pitfalls and caveats of the approach. To this end, we demonstrate the method on UPOs of the Lorenz equations at the standard parameters, in [section 4](#), and then consider a model problem relevant to applications, i.e. the sensitivity with respect to feedback control for UPOs of the Kuramoto-Sivashinsky equation, in [section 5](#). Their modest size makes them convenient playgrounds to test ideas and develop algorithms and enables a statistical analysis of UPO sensitivity, across inventories of UPOs of size $\sim 10^5$ and 4×10^4 ,

respectively.

The sensitivity methods reported in this paper only apply to systems with discrete symmetries. For instance, for the KS equation, we restrict the evolution on the flow-invariant subspace of odd solutions, resulting in a dynamical system with a discrete symmetry, [24, 23]. Methods for sensitivity analysis time-periodic invariant solutions of systems with continuous symmetries are left to future work.

Tangent and adjoint sensitivity techniques for open (non-periodic) trajectories are first presented, in section 2. These methods are well known but are reported here to highlight the differences with analogous sensitivity techniques for time periodic trajectories, discussed in section 3. For convenience, we use a finite-dimensional setting and assume the dynamical system is defined by a set of ordinary differential equations (ODEs). Similar techniques can be also be developed in the continuous setting for partial differential equations.

2. Tangent and adjoint sensitivity of non-periodic unstable trajectories. Let

$$(1a) \quad \begin{cases} \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t); \mathbf{p}), & t \in [0, T] \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases}$$

be an initial value problem (IVP), where $\mathbf{x}(t) \in \mathbb{R}^{N_{\mathbf{x}}}$ denotes the reference state space trajectory originating at some initial condition \mathbf{x}_0 and t is time. The field $\mathbf{f} : \mathbb{R}^{N_{\mathbf{x}}} \times \mathbb{R}^{N_{\mathbf{p}}} \rightarrow \mathbb{R}^{N_{\mathbf{x}}}$ is a smooth vector function of \mathbf{x} and depends, additionally, on a vector of variables $\mathbf{p} \in \mathbb{R}^{N_{\mathbf{p}}}$, the design variables of the problem. The variables \mathbf{p} parametrise implicitly, in some unknown fashion, the trajectories of (1). Let now

$$(2) \quad K(t) = K(\mathbf{x}(t), \mathbf{p}) : \mathbb{R}^{N_{\mathbf{x}}} \times \mathbb{R}^{N_{\mathbf{p}}} \rightarrow \mathbb{R}$$

denote an observable of the system, also referred to as cost in the following. Its average over the reference trajectory,

$$(3) \quad \bar{K} = \frac{1}{T} \int_0^T K(\mathbf{x}(t), \mathbf{p}) dt,$$

will depend implicitly on the parameters too.

The quantity of interest is the *sensitivity of the average* to parameter perturbations, expressed by the gradient $(\bar{K})_{\mathbf{d}\mathbf{p}}$. In practice, sensitivity techniques only provide the *average of the sensitivity*, obtained by time averaging the result of a linear problem (see discussion in [34]). In the tangent approach, the average of the sensitivity is obtained by differentiating (3) with respect to the parameters, leading to

$$(4) \quad \overline{(K_{\mathbf{d}\mathbf{p}})} = \frac{1}{T} \int_0^T [K_{\partial\mathbf{x}}(t) \cdot \mathbf{x}_{\partial\mathbf{p}}(t) + K_{\partial\mathbf{p}}(t)] dt,$$

where subscripts ‘d’ and ‘∂’ denote total and partial differentiation. We follow the notation that gradients are row vectors.

The quantity $\mathbf{x}_{\partial\mathbf{p}}(t) \in \mathbb{R}^{N_{\mathbf{x}} \times N_{\mathbf{p}}}$ is the state sensitivity. The linear problem for its evolution is obtained by differentiating (1) with respect to the parameters, leading to the tangent

equations, the IVP

$$(5a) \quad \begin{cases} \dot{\mathbf{x}}_{\partial \mathbf{p}}(t) = \mathbf{f}_{\partial \mathbf{x}}(t) \cdot \mathbf{x}_{\partial \mathbf{p}}(t) + \mathbf{f}_{\partial \mathbf{p}}(t), & t \in [0, T] \\ \mathbf{x}_{\partial \mathbf{p}}(0) = \mathbf{0}. \end{cases}$$

where $\mathbf{f}_{\partial \mathbf{x}}(t) \in \mathbb{R}^{N_{\mathbf{x}} \times N_{\mathbf{x}}}$ and $\mathbf{f}_{\partial \mathbf{p}}(t) \in \mathbb{R}^{N_{\mathbf{x}} \times N_{\mathbf{p}}}$, the Jacobians of the vector field, contain the partial derivatives of \mathbf{f} with respect to the state and parameters, respectively. The state sensitivity describes the effects that parameter perturbations have on the *future* evolution of state space trajectories originating from the same initial condition \mathbf{x}_0 , hence $\mathbf{x}_{\partial \mathbf{p}}(0) = \mathbf{0}$. For an unstable reference trajectory, the state sensitivity grows exponentially at an average rate determined by the leading Lyapounov exponent of the system. This is a reflection of the ill-conditioning of the IVP (1), whereby small parameter perturbations lead to large changes in the solution [41]. When the length T of the trajectory is increased, with the idea of obtaining the sensitivity of converged averages, the integral (4) does not converge to a finite value but grows exponentially and yields unphysical gradients [25, 34, 39].

The same limitation applies when an adjoint approach is used to obtain the gradient (4), even if the state sensitivity $\mathbf{x}_{\partial \mathbf{p}}(t)$ is never computed explicitly. The governing equation (1) is adjointed to the time average (3), via the vector of adjoint variables $\boldsymbol{\lambda}(t) \in \mathbb{R}^{N_{\mathbf{x}}}$, forming the Lagrangian

$$(6) \quad \mathcal{L} = \frac{1}{T} \int_0^T K(t) + \boldsymbol{\lambda}^\top(t) \cdot [\dot{\mathbf{x}}(t) - \mathbf{f}(t)] dt.$$

Integrating by parts with respect to time the term $\boldsymbol{\lambda}^\top \cdot \dot{\mathbf{x}}$ leads to

$$(7) \quad \mathcal{L} = \frac{1}{T} \int_0^T K(t) - \dot{\boldsymbol{\lambda}}^\top(t) \cdot \mathbf{x}(t) - \boldsymbol{\lambda}^\top(t) \cdot \mathbf{f}(t) dt + [\boldsymbol{\lambda}^\top(t) \cdot \mathbf{x}(t)]_0^T,$$

and differentiation with respect to the parameters yields

$$(8) \quad \mathcal{L}_{\text{dp}} = \overline{(K_{\text{dp}})} = \frac{1}{T} \int_0^T \underbrace{[K_{\partial \mathbf{x}}(t) - \dot{\boldsymbol{\lambda}}^\top(t) - \boldsymbol{\lambda}^\top(t) \cdot \mathbf{f}_{\partial \mathbf{x}}(t)]}_{\mathbf{a}} \cdot \mathbf{x}_{\partial \mathbf{p}}(t) - \boldsymbol{\lambda}^\top(t) \cdot \mathbf{f}_{\partial \mathbf{p}}(t) + K_{\partial \mathbf{p}}(t) dt + \underbrace{[\boldsymbol{\lambda}^\top(t) \cdot \mathbf{x}_{\partial \mathbf{p}}(t)]_0^T}_{\mathbf{b}},$$

where the first equality holds because it is assumed that when parameters are perturbed the governing equation remains satisfied, for any choice of $\boldsymbol{\lambda}(t)$. The adjoint variables are selected such that \mathbf{a} and \mathbf{b} vanish, circumventing the (expensive) calculation of the state sensitivity. Requiring $\mathbf{a} = \mathbf{0}$ yields the adjoint equation (10a), while its initial conditions are obtained by requiring $\mathbf{b} = \mathbf{0}$, i.e.

$$(9) \quad \mathbf{0} = \mathbf{b} = [\boldsymbol{\lambda}^\top(t) \cdot \mathbf{x}_{\partial \mathbf{p}}(t)]_0^T = \boldsymbol{\lambda}^\top(T) \cdot \mathbf{x}_{\partial \mathbf{p}}(T) - \boldsymbol{\lambda}^\top(0) \cdot \mathbf{x}_{\partial \mathbf{p}}(0)$$

For any $\boldsymbol{\lambda}(0)$, the second term at $t = 0$ is zero because the initial conditions are assumed not to change, while for any $\mathbf{x}_{\partial \mathbf{p}}(T)$, the first term at $t = T$ vanishes if $\boldsymbol{\lambda}(T) = \mathbf{0}$. This results in

the adjoint problem

$$\begin{aligned} (10a) \quad & \dot{\boldsymbol{\lambda}}(t) = -\mathbf{f}_{\partial \mathbf{x}}^{\top}(t) \cdot \boldsymbol{\lambda}(t) + K_{\partial \mathbf{x}}^{\top}(t) \\ (10b) \quad & \boldsymbol{\lambda}(T) = \mathbf{0}, \end{aligned}$$

a linear initial value problem with homogeneous terminal conditions. Because of such conditions, problem (10) needs to be solved *backward* in time. This adds a number of well-known complications since the time varying coefficients in (10) are defined on a trajectory that proceeds *forward* in time.

Similarly to the tangent sensitivity problem, the solution of (10) grows at an average exponential rate when integrated in time over an unstable trajectory. Hence, the integral

$$(11) \quad \mathcal{L}_{\mathbf{d}\mathbf{p}} = \overline{(K_{\mathbf{d}\mathbf{p}})} = \frac{1}{T} \int_0^T -\boldsymbol{\lambda}^{\top}(t) \cdot \mathbf{f}_{\partial \mathbf{p}}(t) + K_{\partial \mathbf{p}}(t) dt.$$

does not converge as the length of the trajectory is increased, but the gradient gets exponentially larger, and unphysical, for increasing T [25, 11, 34, 39].

3. Tangent and adjoint sensitivity of periodic unstable trajectories. For time periodic problems, it is practical to first introduce the so-called Linsteadt-Poincaré transformation

$$(12) \quad s = 2\pi t/T = \omega t,$$

with ω being the fundamental frequency of a particular reference periodic trajectory $\mathbf{x}^o(s)$, to uniformly transform time $t \in [0, T]$ into the loop parameter $s \in [0, 2\pi]$. This transformation makes the natural frequency appear explicitly in the evolution equation

$$\begin{aligned} (13a) \quad & \omega \dot{\mathbf{x}}^o(s) = \mathbf{f}(\mathbf{x}^o(s); \mathbf{p}), \quad s \in [0, 2\pi] \\ (13b) \quad & \mathbf{x}^o(0) = \mathbf{x}^o(2\pi) \end{aligned}$$

and removes the explicit dependence of the period T from the limits of integration in the calculation of the period averages, denoted by brackets, i.e.

$$(14) \quad \langle K \rangle = \frac{1}{2\pi} \int_0^{2\pi} K(\mathbf{x}^o(s), \mathbf{p}) ds.$$

Problem (13) is a nonlinear time periodic boundary value problem (BVP). Its solutions, infinitely many for a chaotic system, are pairs $\{\mathbf{x}^o(s), \omega\}$, pairs of unstable periodic orbits and the associated frequencies. Numerical methods to solve BVPs like (13) are well known [2]. In this paper, the modest size of the systems we consider enables us to use an approach known as *quasi-linearisation* in the literature, to reduce the nonlinear problem to a sequence of linear BVPs solved with standard finite difference technique for temporal discretisation (see Appendices A and B). For high-dimensional systems, shooting strategies might be more practical [2, 43].

Similarly to solutions of the IVP (1), the solutions of the BVP (both the trajectory and the frequency) depend implicitly on the parameters \mathbf{p} , in some unknown fashion, because the

vector field where UPOs reside is parametrised. Small perturbations of the BVP (13), e.g. a small parameter perturbation, produces, in general, small changes to its solutions, i.e. a small displacement/distortion of the trajectory and a small change in the frequency. In other words, the condition number of problem (13) is only ‘moderate’ and does not increase exponentially with the length of the orbit, as for the IVP approach [41].

For tangent sensitivity analysis, the gradient of the period average is obtained by differentiating (14), leading to

$$(15) \quad \langle K_{\text{d}\mathbf{p}} \rangle = \langle K_{\partial\mathbf{x}}(s) \cdot \mathbf{x}_{\partial\mathbf{p}}(s) + K_{\partial\mathbf{p}}(s) \rangle.$$

The evolution equation for the state sensitivity $\mathbf{x}_{\partial\mathbf{p}}(s)$ is obtained by differentiating (13) with respect to the parameters, to obtain

$$(16a) \quad \begin{cases} \omega \dot{\mathbf{x}}_{\partial\mathbf{p}}(s) = \mathbf{f}_{\partial\mathbf{x}}(s) \cdot \mathbf{x}_{\partial\mathbf{p}}(s) + \mathbf{f}_{\partial\mathbf{p}}(s) - \omega_{\partial\mathbf{p}} \dot{\mathbf{x}}^o(s), & s \in [0, 2\pi] \\ \mathbf{x}_{\partial\mathbf{p}}(0) = \mathbf{x}_{\partial\mathbf{p}}(2\pi) \\ \mathbf{0} = \mathbf{x}_{\partial\mathbf{p}}(0)^\top \cdot \mathbf{f}(0). \end{cases}$$

a linear time periodic BVP. Such boundary conditions reflect the fact that the arbitrary orbit origin $\mathbf{x}^o(0) = \mathbf{x}^o(2\pi)$ is a single point that moves in a single direction when parameters are varied, so that the perturbed orbit remains periodic. Unlike in (5b), $\mathbf{x}_{\partial\mathbf{p}}(0) \neq \mathbf{0}$ and every point on the orbit is allowed to move in state space when parameters are varied. The gradient $\omega_{\partial\mathbf{p}}$ is identically zero for forced oscillations, but it is non-zero and unknown for limit-cycle-type oscillations and its value can be obtained as part of the solution of (16). The phase locking constraint (16c) ensures that (16) has a unique solution. Without this condition, problem (16) has a one-parameter family of solutions, corresponding to arbitrary shifts along the orbit: if $\mathbf{x}_{\partial\mathbf{p}}(s)$ is a solution, $\mathbf{x}_{\partial\mathbf{p}}(s + \Delta)$ is a solution too, for any real Δ . Other equivalent forms of this constraint can be also introduced [23, 42].

Solutions of (16) are time periodic and hence bounded, regardless of the stability characteristics and length of the underlying reference trajectory. The exponential growth of solutions of (5) over unstable non-periodic trajectories is avoided in a natural fashion by the choice of the boundary conditions. This implies that the linearisation underpinning (15) and (16) remains valid regardless of the length of the orbit, and thus

$$(17) \quad \langle K \rangle_{\text{d}\mathbf{p}} = \langle K_{\text{d}\mathbf{p}} \rangle.$$

In other words, the sensitivity of the asymptotic state (now periodic), i.e. the *sensitivity of the average*, coincides with the *average of the sensitivity* obtained from the linear problem.

As anticipated in the introduction, the proposed sensitivity method for periodic trajectories shares several similarities with Least-Squares Shadowing (LSS) methods for sensitivity analysis of chaotic trajectories [5]. The common ground is that the initial condition of the perturbed trajectory is permitted to change, to avoid the exponential divergence. In LSS, the perturbation of the initial condition $\mathbf{x}_{\partial\mathbf{p}}(0)$ and the entire sensitivity solution $\mathbf{x}_{\partial\mathbf{p}}(t)$ are obtained from the solution of a constrained optimisation problem, whereby the average norm of $\mathbf{x}_{\partial\mathbf{p}}(t)$ is minimised subject to the constraint of the sensitivity equation (Equation 13 in [41]). In stark contrast, for a periodic trajectory, the sensitivity solution is obtained, more

naturally, by requiring temporal periodicity. It is important to point out, though, that for a periodic trajectory, the computed sensitivity is exact, and the method does not display the “shadowing error” of LSS [40] at the trajectory end points.

For the adjoint method, we proceed as in section 2, by formulating a Lagrangian and differentiating it with respect to the parameters, to obtain

$$(18) \quad \mathcal{L}_{\text{dp}} = \langle K \rangle_{\text{dp}} = \left\langle \underbrace{[K_{\partial \mathbf{x}}(s) - \omega \dot{\boldsymbol{\lambda}}^\top(s) - \boldsymbol{\lambda}^\top(s) \cdot \mathbf{f}_{\partial \mathbf{x}}(s)] \cdot \mathbf{x}_{\partial \mathbf{p}}(s) - \boldsymbol{\lambda}^\top(s) \cdot \mathbf{f}_{\partial \mathbf{p}}(s)}_{\mathbf{a}} + K_{\partial \mathbf{p}}(s) + \omega_{\partial \mathbf{p}} \boldsymbol{\lambda}^\top(s) \cdot \dot{\mathbf{x}}^o(s) \right\rangle + \underbrace{[\boldsymbol{\lambda}^\top(s) \cdot \mathbf{x}_{\partial \mathbf{p}}(s)]_0^{2\pi}}_{\mathbf{b}}$$

which differs from (8) only by the last term in the brackets arising from the time transformation. Requiring \mathbf{a} to vanish leads to the adjoint equation (20a), the same as (10a). Requiring \mathbf{b} to vanish, i.e.

$$(19) \quad \mathbf{0} = \mathbf{b} = [\boldsymbol{\lambda}^\top(s) \cdot \mathbf{x}_{\partial \mathbf{p}}(s)]_0^{2\pi} = \boldsymbol{\lambda}^\top(2\pi) \cdot \mathbf{x}_{\partial \mathbf{p}}(2\pi) - \boldsymbol{\lambda}^\top(0) \cdot \mathbf{x}_{\partial \mathbf{p}}(0) = [\boldsymbol{\lambda}(2\pi) - \boldsymbol{\lambda}(0)]^\top \cdot \mathbf{c},$$

leads to time periodic boundary conditions (20b), because $\mathbf{x}_{\partial \mathbf{p}}(0) = \mathbf{x}_{\partial \mathbf{p}}(2\pi) = \mathbf{c}$ for some unknown \mathbf{c} . This results in the adjoint problem

$$(20a) \quad \begin{cases} \omega \dot{\boldsymbol{\lambda}}(s) = -\mathbf{f}_{\partial \mathbf{x}}^\top(s) \cdot \boldsymbol{\lambda}(s) + K_{\partial \mathbf{x}}^\top(s) \\ (20b) \quad \boldsymbol{\lambda}(0) = \boldsymbol{\lambda}(2\pi) \\ (20c) \quad 0 = \boldsymbol{\lambda}^\top(0) \cdot \mathbf{f}(0) \end{cases}$$

which should be compared to (10). The constraint (20c) is required to break the translational invariance; the solution is not unique but there is a one-parameter family of adjoint solutions.

Similarly to the tangent problem, the adjoint problem is a linear time periodic BVP, the adjoint solution is periodic and bounded and cannot grow exponentially in time. An appropriate numerical method must be used to enforce the boundary conditions. When shooting techniques are used for the solution [2, 43], the adjoint equation *does not need to be marched backwards* but the “forward” and adjoint equations can be marched forward simultaneously, alleviating well-known numerical issues associated to (10). Numerical methods required to solve (20) and (16) are fundamentally similar to those used to search periodic trajectories in the first place, enabling reuse of algorithms and computational tools (see Appendix B).

Upon solution, the gradient of the period average can then be calculated as

$$(21) \quad \langle K \rangle_{\text{dp}} = \left\langle K_{\partial \mathbf{p}}(s) - \boldsymbol{\lambda}^\top(s) \cdot \mathbf{f}_{\partial \mathbf{p}}(s) - \omega_{\partial \mathbf{p}} \boldsymbol{\lambda}^\top(s) \cdot \dot{\mathbf{x}}^o(s) \right\rangle.$$

The unknown frequency gradient still appears in (21). Rather than from the (expensive) solution of the tangent sensitivity problem, this quantity can be obtained at little additional cost using Fredholm’s alternative ([17], pg. 146). Fredholm’s alternative expresses a compatibility constraint for the solution of linear time periodic BVPs. Specifically, problem (16) has time-periodic solutions if and only if

$$(22) \quad \int_0^{2\pi} \mathbf{w}^\top(s) \cdot [\mathbf{f}_{\partial \mathbf{p}}(s) - \omega_{\partial \mathbf{p}} \dot{\mathbf{x}}^o(s)] ds = 0$$

for all 2π -periodic solutions of the homogeneous adjoint problem

$$\begin{aligned} (23a) \quad & \omega \dot{\mathbf{w}}(s) = -\mathbf{f}_{\partial \mathbf{x}}^\top(s) \cdot \mathbf{w}(s) \\ (23b) \quad & \mathbf{w}(0) = \mathbf{w}(2\pi) \end{aligned}$$

A non-trivial solution can be obtained by augmenting (23) with a inhomogeneous phase locking condition $\mathbf{w}(0)^\top \cdot \mathbf{f}(0) = 1$, and part of the computational cost can be amortised by solving it along with (20). The gradient of the frequency is then

$$(24) \quad \omega_{\partial \mathbf{p}} = \frac{\int_0^{2\pi} \mathbf{w}^\top(s) \cdot \mathbf{f}_{\partial \mathbf{p}}(s) \, ds}{\int_0^{2\pi} \mathbf{w}^\top(s) \cdot \dot{\mathbf{x}}^o(s) \, ds}$$

3.1. Non-uniqueness of the state sensitivity. The gradient $\langle K \rangle_{\text{dP}}$ quantifies the global, average effects of perturbations but does not convey information about how these are distributed during the UPO period. Knowledge of such a distribution could be useful, for instance, to pinpoint instants where parameters have the largest influence on the cost.

Using the tangent sensitivity formulation, and assuming a single scalar parameter p , one could recall (15) and use the quantity $K_{\partial \mathbf{x}}(s) \cdot \mathbf{x}_{\partial p}(s) + K_{\partial p}(s)$ for this purpose. This would be incorrect, however, because the sensitivity $\mathbf{x}_{\partial p}(s)$ fails to convey uniquely the desired insight into the orbit deformation, since the orbit origin, and thus the phase locking condition to break the time-translational invariance, is arbitrary. Similar considerations apply for the adjoint variables. In other words, the global change in state space structure of an orbit is unique, but there is a one-parameter family of time-preserving sensitivity trajectories that express this change locally. These can be generated as

$$(25) \quad \mathbf{x}_{\partial p}^\tau(s) = \mathbf{x}_{\partial p}(s) + \tau \mathbf{f}(s),$$

by a shift along the orbit. Note that for any τ ,

$$(26) \quad \langle K_{\partial \mathbf{x}}(s) \cdot \mathbf{x}_{\partial p}(s) \rangle = \langle K_{\partial \mathbf{x}}(s) \cdot \mathbf{x}_{\partial p}^\tau(s) \rangle$$

holds, as $\langle K_{\partial \mathbf{x}}(s) \cdot \mathbf{f}(s) \rangle$ is the closed path integral of the conservative vector field $K_{\partial \mathbf{x}}$ and it is identically zero. Here we wish to develop a method to express the local change uniquely.

The phase locking condition (16c) selects one sensitivity trajectory by enforcing the sensitivity to be orthogonal to the flow at the arbitrary origin. However, on the rest of the orbit the orthogonality is not guaranteed, i.e., $\mathbf{x}_{\partial p}(s) \cdot \mathbf{f}(s) \neq 0$ in general. The idea, inspired by LSS [41], is then to calculate *a posteriori* the shift τ^* that minimises the average squared norm of the sensitivity, the quantity $\tau^* = \arg \min_{\tau} \langle \|\mathbf{x}_{\partial p}^\tau(s)\|^2 \rangle$. A simple variational calculation shows that this shift is given by

$$(27) \quad \tau^* = -\langle \mathbf{x}_{\partial p}^\top(s) \cdot \mathbf{f}(s) \rangle / \langle \mathbf{f}^\top(s) \cdot \mathbf{f}(s) \rangle.$$

This shift is unique for a given periodic trajectory and can be used to define and illustrate the local changes to the orbit structure induced by parameter perturbations. An alternative,




Figure 1. Panel (a): shortest UPO for selected values of ρ . Panel (b): variation of the period average $\langle z \rangle$ with ρ (dashed line, left axis) and its gradient $\langle z \rangle_{d\rho}$ (right axis), computed using the adjoint algorithm (ADJ, solid line) and approximated via finite differences (FD, circles).

similarly inspired by LSS and which we do not pursue here, is to remove the component of the sensitivity tangent to the direction of the flow $\mathbf{x}_{\partial\rho}^{\parallel}(s)$ and retain the component orthogonal to it, $\mathbf{x}_{\partial\rho}^{\perp}(s)$, to obtain a solution that has least possible squared average. This is achieved by a straightforward orthogonalisation,

$$(28) \quad \mathbf{x}_{\partial\rho}^{\perp}(s) = \mathbf{x}_{\partial\rho}(s) - \mathbf{x}_{\partial\rho}^{\parallel}(s) = \mathbf{x}_{\partial\rho}(s) - \mathbf{n}(s)(\mathbf{x}_{\partial\rho}(s) \cdot \mathbf{n}(s))$$

so that $\mathbf{x}_{\partial\rho}^{\perp}(s) \cdot \mathbf{f}(s) = 0, \forall s \in [0, 2\pi]$ holds, where $\mathbf{n}(s) = \mathbf{f}(s)/\|\mathbf{f}(s)\|$ is the unit-norm loop tangent vector. This operation, however, changes the flow of time and (26) ceases to hold.

4. Application to the Lorenz Equations. The adjoint sensitivity algorithm discussed in section 3 is demonstrated on UPOs of the celebrated Lorenz equations [26]. These are

$$(29) \quad \begin{cases} \dot{x}(t) = \sigma(x(t) - y(t)) \\ \dot{y}(t) = \rho x(t) - y(t) - x(t)z(t) \\ \dot{z}(t) = x(t)y(t) - \beta z(t) \end{cases}$$

where the standard parameters $\sigma = 10$, $\beta = 8/3$ and $\rho = 28$ are used throughout. Let z be the observable of the system, ρ the sensitivity parameter and let us calculate the sensitivity $\langle z \rangle_{d\rho}$. An approximation of the sensitivity $(\bar{z})_{d\rho}$ around $\rho = 28$ obtained using, e.g., finite difference quotients of long-time averages or similar techniques, is about unity.

We first validate the approach by considering the shortest UPO, with cycle length $n = 2$, with n the number of intersections with the standard Poincaré section $z = \rho - 1$, $\dot{z}(s) > 0$. This UPO spans both lobes of the attractor, has symbol sequence AB (in the notation of [37]) and is invariant under the symmetry transformation $S : [x, y, z] \rightarrow [-x, -y, z]$. We find numerically this orbit at $\rho = 25$ and perform continuation in ρ up to $\rho = 45$, with intervals $\Delta\rho = 0.2$. UPOs for selected values of ρ are shown in Figure 1-(a). The period average $\langle z \rangle$, shown in panel (b) as a function of ρ (dashed line, left axis), follows approximately a linear behaviour. The slope

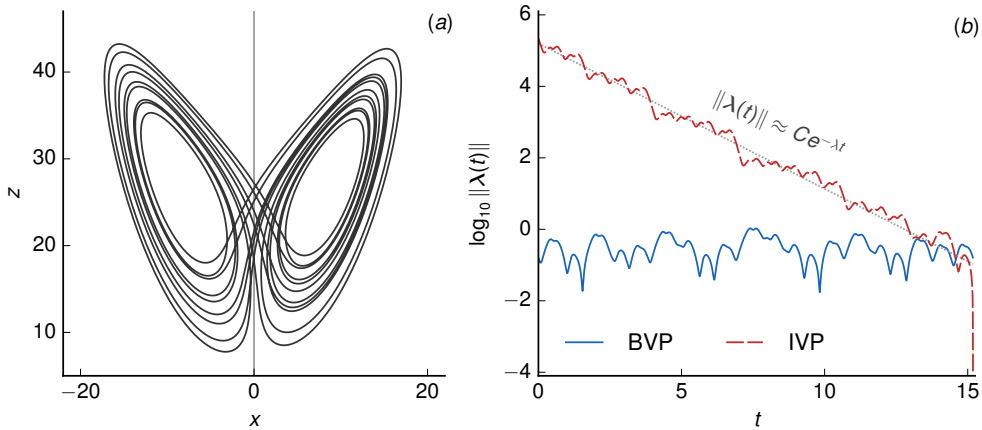


Figure 2. Panel (a): the $n = 20$ UPO ($T \approx 15.1827$) used as reference trajectory to solve the adjoint BVP (20) and IVP (10). Panel (b): the norm of the adjoint variables from solution of these two problems.

of this curve, obtained using a second order accurate finite difference approximation for a few values of ρ (blue circles, right axis, denoted by FD), spans approximately the range $[0.99, 1.04]$. The sensitivity $\langle z \rangle_{d\rho}$ calculated using the adjoint method discussed previously (solid line) matches with the finite difference approximation up to discretisation errors involved in the solution of the relevant BVPs and the truncation/cancellation errors in the finite difference quotient.

We now calculate the sensitivity for a long UPO ($n = 20$, Figure 2-(a)) using the BVP approach (20), and the IVP approach (10), using the backward-in-time integration. For the latter, the UPO is treated as an open trajectory, and we select an arbitrary point on it as the initial/final condition. In Figure 2-(b), the time histories of the norm of the adjoint variables from the solution of two problems are reported. When the traditional IVP approach is used (dashed red line), the norm of the adjoint variables starts from zero, rapidly jumps to a finite value and then grows exponentially at an average rate equal to the positive Lyapounov exponent of the system at these parameters, $\lambda \approx 0.906$ [33]. The sensitivity of the time average calculated from (11) is unphysical, $(\bar{z})_{d\rho} \approx 3017.6$. In stark contrast, the norm of the adjoint variables from the solution of the BVP (20) (solid blue line) is a periodic function. The sensitivity of the period average (21) is physically meaningful, $\langle z \rangle_{d\rho} \approx 1.01847$. This applies regardless of the length of the UPOs (see Appendix D for sensitivity analysis of an UPO with $T \approx 1080.5$).

4.1. Statistics of UPO gradients. The purpose of this section is to analyse the statistical distribution of the sensitivity across UPOs of the Lorenz equations. We will sometimes make use of the expectation $E[\cdot]$, defined as the equally weighted average over all UPOs. This average has no formal justification, but we use it here and in the following sections primarily to give a quantitative measure of the distribution, for instance, to quantify the standard deviation of the sensitivity across UPOs of same cycle length.

At the standard parameters, we find a set of 107406 unique orbits with cycle length $n \leq 20$, about 3% less than the number of orbits (111011) known to compose the symbolic dynamics

for $n \leq 20$ [37]. This set includes symmetry-invariant UPOs and pairs of symmetry-equivalent orbits, where one orbit is generated from the other via action of the symmetry. Uniqueness is verified using the graph-based approach described in Appendix C. UPOs are refined iteratively until the estimated error on the orbit baricenter $\epsilon_{\mathbf{h}}$ in subsequent iterations is smaller than 10^{-6} (see Appendix C for details), with a typical time step of $5 \cdot 10^{-4}$. The gradient $\langle z \rangle_{d\rho}$ is also typically accurate to this precision.




Figure 3. Panel (a): normalised histogram of $\langle z \rangle_{d\rho}$. The horizontal bar denotes the mean of the distribution. Panel (b): scatter plot of period and gradient of UPOs with even cycle number. Circles identify UPOs with symbol sequence $A^{n-1}B$, for all n . The grey band defines the range displayed in panel (c). Panel (c): expectation of gradient and period. The vertical bars span the standard deviation of the distributions for each n .

Figure 3-(a) shows the normalised histogram of $\langle z \rangle_{d\rho}$ constructed by binning the range between 0.995 and 1.02 with 500 bins. No data point falls outside this range. The distribution of the gradients is remarkably narrow; the right tail falls abruptly while the left tail decays more gently, displaying several humps. The gradient expectation (horizontal line in the figure) is $E[\langle z \rangle_{d\rho}] = 1.0177$, with standard deviation 0.00117. In Figure 3-(b) the scatter of the orbits' period T and gradient $\langle z \rangle_{d\rho}$ is reported for UPOs of even cycle number n .

The distribution of gradients is structured. Data points are organised in clusters, causing the humps in the left tail of the histogram of Figure 3-(a). The subset of UPOs with symbol sequence pattern $A^{n-1}B$, open circles, shown for $2 \leq n \leq 20$, displays a wavy pattern and produce extreme values of the gradient, up to $n = 19$, after which the orbit $A^{18}B^2$ has the lowest gradient. Tracing UPOs with such pattern for $n > 20$ and examining the asymptotics

Table 1

Characteristics of the first nineteen UPOs with symbol sequence $A^{n-1}B$. All reported digits are converged. Data points in the third and last columns correspond to the open circles in [Figure 3-\(b\)](#).

n	sequence	T	$T_{d\rho}$	$\langle z \rangle$	$\langle z \rangle_{d\rho}$
2	AB	1.5586522	-0.0392528	23.0421039	1.0181392
3	A^2B	2.3059072	-0.0573500	23.2322548	1.0190202
4	A^3B	3.0235837	-0.0740579	23.4646321	1.0197131
5	A^4B	3.7256417	-0.0899305	23.6658224	1.0194636
6	A^5B	4.4177664	-0.1051958	23.8338832	1.0184538
7	A^6B	5.1030384	-0.1199824	23.9742312	1.0168975
8	A^7B	5.7834108	-0.1343760	24.0920415	1.0149585
9	A^8B	6.4602592	-0.1484423	24.1913462	1.0127595
10	A^9B	7.1346327	-0.1622367	24.2751574	1.0103966
11	$A^{10}B$	7.8073881	-0.1758120	24.3456991	1.0079511
12	$A^{11}B$	8.4792710	-0.1892228	24.4045969	1.0054978
13	$A^{12}B$	9.1509718	-0.2025301	24.4530089	1.0031120
14	$A^{13}B$	9.8231719	-0.2158068	24.4917076	1.0008768
15	$A^{14}B$	10.4965892	-0.2291449	24.5211223	0.9988917
16	$A^{15}B$	11.1720295	-0.2426673	24.5413439	0.9972860
17	$A^{16}B$	11.8504579	-0.2565465	24.5520879	0.9962404
18	$A^{17}B$	12.5331044	-0.2710385	24.5525992	0.9960262
19	$A^{18}B$	13.2216410	-0.2865488	24.5414594	0.9970803
20	$A^{19}B$	13.9185025	-0.3037700	24.5162058	1.0001659

of their sensitivity would be interesting and is left to future analysis. [Table 1](#) reports the period T , the period average of z and the gradients of these two quantities for such a subset.

The long orbits can occasionally predict gradients significantly lower than average, causing the gentle decay of the distribution of in [Figure 3-\(a\)](#). Nevertheless, the standard deviation of the gradient across UPOs of same cycle length seems to decrease with n . This is illustrated in panel (c), reporting the expected gradient $E[\langle z \rangle_{d\rho}]$ and the expected standard deviation (vertical bars) as a function of the expected period $E[T]$ across orbits of same cycle length n .

We have used linear regression of long-time averages obtained from integration of the equations for twenty one values of ρ at equally sampled points in the interval across $\rho = 28 \pm 0.25$, to approximate the gradient $(\bar{z})_{d\rho}$. Accurate averages of z were obtained by augmenting the equations with the quadrature equation $\dot{\xi}(t) = (z(t) - \xi(t))/t$ (the variable $\xi(\tau)$ is the cumulative average $\xi(\tau) = 1/\tau \int_0^\tau z(t) dt = \bar{z}^\tau$) and integrating the augmented system using a linear multi-step, variable order integrator with relative and absolute tolerance of 10^{-7} from the Sundials suite of integrators [18].

We repeated this approach over 30 trials, using random initial conditions and discarding initial transients. At $\tau = 10^6$, the average value of $(\bar{z}^\tau)_{d\rho}$ across the trials was equal to 1.0027, with a standard deviation of 0.0012. Different values of the gradient $(\bar{z})_{d\rho}$ have been reported by other authors, e.g., 0.96 in [25] and 1.01 ± 0.04 in [39]. Overall, this approximation lies at the very end of the left tail of [Figure 3-\(a\)](#). This is quite far from where the vast majority of UPOs are located, in particular the short orbits. It is speculated here that the mismatch between the numerical approximation of the gradient and the distribution of [Figure 3-\(a\)](#) by the fact that

the numerical calculation of statistical quantities over long trajectories in chaotic dynamical systems is a delicate subject [29]. In absence of hyperbolicity, as for the Lorenz equations at standard parameters [32], “there is no fundamental reason for computer simulated long-time statistics to be even approximately correct” [28]. In stark contrast, periodic trajectories have desirable convergence properties. Computations in finite precision arithmetic on periodic orbits are *globally* insensitive since the catastrophic temporal propagation of the effects of floating-point errors is constrained by the need to satisfy the periodic boundary conditions of the problem, i.e. robustness is built in the numerical algorithm. Hence, periodic trajectories identified from numerical computation converge when the time step is refined and so do period averages and their gradients, regardless of the orbit length and its stability properties.

5. Feedback control in Kuramoto-Sivashinsky equation. In this section, we consider a sensitivity problem for UPOs of the Kuramoto-Sivashinsky equation (KSEq). The KSEq is a well known one-dimensional partial differential equation exhibiting complex spatio-temporal behaviour. It arises in a variety of physically relevant context, e.g. reaction diffusion systems [22] and the dynamics of liquid films falling along a wall [31]. A fairly large number of numerical and theoretical investigations on bifurcation, chaos and control [1, 14] have used the KSEq as a model problem. Of particular interest in the current context are early [8] and more recent [24, 9] investigations on UPOs of the KSEq.

The equation considered here is

$$(30) \quad v_{\partial t}(x, t) = (v^2)_{\partial x}(x, t) - v_{\partial^2 x}(x, t) - \kappa v_{\partial^4 x}(x, t) + h(x, t)$$

where $v(x, t)$ is the scalar field variable, x denotes space, t is time and the subscript $(\cdot)_{\partial^n}$ denotes the n -th partial derivative with respect to time or space. The additional source term $h(x, t)$ in (30) represents the distributed volume force actuation that is used to control the system, via a to-be-determined feedback controller \mathcal{G} such that

$$(31) \quad h(x, t) = \mathcal{G}(v(x, t)).$$

In such settings, (30) is still an autonomous system, parametrised by a set of controller variables. The equation is parametrised by the diffusivity constant κ , while the domain length is fixed to 2π . This set up is equivalent to fixing the viscosity, $\kappa = 1$, and varying the length of the domain, where the equivalence is $L = 2\pi/\sqrt{\kappa}$, [24]. We consider 2π -periodic, odd solutions, satisfying $v(x + 2\pi) = v(x)$ and $v(-x) = -v(x)$. Considering odd solutions reduces the continuous translational symmetry

$$(32) \quad \tau_{\Delta}^{\dagger} : v(x, t) \rightarrow v(x + \Delta, t), \quad \Delta \in \mathbb{R}$$

to a discrete shift-by- π symmetry

$$(33) \quad S^{\dagger} : v(x, t) \rightarrow v(x + n\pi, t), \quad n \text{ odd.}$$

The energy density

$$(34) \quad K^{\dagger}(t) = K^{\dagger}(v(x, t)) = \frac{1}{4\pi} \int_0^{2\pi} v(x, t)^2 dx$$

is the target observable for control. We wish to calculate the gradient $\langle K^\dagger \rangle_{\text{d}\mathcal{G}}$, i.e., the sensitivity of the period averaged cost with respect to the controller \mathcal{G} driving the actuation. For a given UPO, the controller $-\lambda \langle K^\dagger \rangle_{\text{d}\mathcal{G}}$, for a scalar $\lambda > 0$, has the effect of displacing/distorting the UPO in state space in a way that is optimal for the reduction of $\langle K^\dagger \rangle$. No other choice of \mathcal{G} results in a larger variation of the period average, at least in a linear sense. The idea is different from previous chaos control studies whereby an UPO with desirable period averaged characteristics is stabilised by feedback [30]. Here, UPOs serve as special trajectories where the sensitivity of the nonlinear chaotic state with respect to the feedback can be extracted exactly. When such information is available for many UPOs, insight can be gained into how the attractor and its statistical moments change under control. Empirically, if most UPOs display similar sensitivity, varying the controller in the direction of the expectation might result in an overall favourable displacement of the attractor. The dynamics will still be chaotic but would be mitigated by actuation, resulting in a reduction of the average cost.

A spectral Fourier discretisation is used, where

$$(35) \quad v(x, t) = \sum_{k=-N}^N i v_k(t) e^{ikx},$$

$i = \sqrt{-1}$, k is the integer wave number and the symmetry $v_{-k} = -v_k$ applies, as Fourier coefficients are imaginary, $v_k \in \mathbb{R}$, $v_0 = 0$, for odd functions. Galerkin projection results in

$$(36) \quad \dot{v}_k(t) = (k^2 - \kappa k^4) v_k(t) - k \sum_{m=-N}^N v_m(t) v_{k-m}(t) - h_k(t), \quad k = 1, \dots, N,$$

where the control variables $h_k(t) \in \mathbb{R}$ arise from a discretisation of $h(x, t)$ analogous to (35). Numerically, we solve (36) by evaluating the convolution sum directly and using the variable order, multi-step stiff integrator from the Sundials suite [18] for time marching, typically setting relative error tolerances to 10^{-6} .

With $\mathbf{v} = (v_1, v_2, \dots, v_N)^\top \in \mathbb{R}^N$ denoting the state variables vector, and similarly for the control variables \mathbf{h} , the finite-dimensional equivalents of the energy density (34) and the symmetry transformation (33) are

$$(37) \quad K(t) = K(\mathbf{v}(t)) = \mathbf{v}^\top(t) \cdot \mathbf{v}(t)$$

and

$$(38) \quad S : v_k \rightarrow \begin{cases} v_k & \text{if } k \text{ is even} \\ -v_k & \text{if } k \text{ is odd} \end{cases}$$

respectively, where K is invariant under S . Assuming linearity of \mathcal{G} , spatial discretisation yields

$$(39) \quad \mathbf{h}(t) = \mathbf{G}\mathbf{v}(t)$$

where the entries of $\mathbf{G} \in \mathbb{R}^{N \times N}$, the control gains, parametrise the dynamics and the UPOs of (36). Sensitivity results will be reported in matrix form, with the gradient $\langle K \rangle_{\text{d}\mathbf{G}} \in \mathbb{R}^{N \times N}$




Figure 4. Panel (a): Spatio-temporal evolution of a chaotic realisation of the uncontrolled KSEq at $\kappa = (2\pi/38.5)^2$. Panel (b): (v_1, v_2) projection of a long trajectory.

and individual entries denoted as $\langle K \rangle_{dG_{km}}$. However, notation for operations such as inner products, norms, etc. involving the gradient $\langle K \rangle_{d\mathbf{G}}$ will be used as if it was a row vector. Equation (36) can be rewritten as

$$(40) \quad \dot{\mathbf{v}}(t) = \mathbf{f}(\mathbf{v}(t); \mathbf{G}),$$

and it is thus formally equivalent to (1). We use notation such as $\mathbf{v}^o(\mathbf{G})$, with s implicit, to denote a periodic solution of (40) for a particular \mathbf{G} , or $\mathbf{v}^o(G_{km})$ for a particular parameter. Similarly, $\langle K(\mathbf{v}^o) \rangle$ and $\langle K(\mathbf{G}) \rangle$ are short hands for $\langle K(\mathbf{v}^o(\mathbf{G})) \rangle$ used to denote the period average of an UPO for a particular \mathbf{G} , with implicit \mathbf{G} and \mathbf{v}^o , respectively, depending on the context.

We investigate control in a “weakly turbulent” regime at $\kappa = (2\pi/39)^2$. This value corresponds to a domain size slightly larger than what studied in [24], where $L = 38.5$. At $L = 38.5$, the post-transient dynamics lands onto an attractor composed of a “center” region (S_C) and two side regions ($S_{L/R}$). A short spatio-temporal sequence of the solution at $L = 38.5$ for the uncontrolled system is reported in Figure 4-(a), while panel (b) shows the (v_1, v_2) projections of a longer trajectory, with labels indicating the three regions of the attractor. In this regime, the dynamics is quite intermittent and trajectories can jump and remain in the side regions for quite some time, before returning to the center region.

In the center region, visited in this realisation during $22 \lesssim t \lesssim 29$, the solution displays two steady boundary layers, whilst most of the spatio-temporally complex activity, a ‘left-right’ wavy motion, takes place in the domain interior. This activity is the result of multiple concurrent instabilities, yet the dynamics is still relatively low dimensional. Linearisation of (36) around the trivial state $v_k = 0, k = 1 \dots N$, leads to a diagonal system with entries, the eigenvalues, $\lambda_k = k^2 - \kappa k^4$. At $\kappa = (2\pi/38.5)^2$, there are six unstable modes for $k < \sqrt{1/\kappa} = 38.5/2\pi \approx 6.127 \dots$, with $k = 4$ the most unstable. This wave number defines the scale of the coherent structures featuring prominently in the center region. In the side regions, $17 \lesssim t \lesssim 22$ and $29 \lesssim t \lesssim 32$, the character of the solution changes drastically and describes the oscillatory growth of an instability that saturates back to the center region. At $\kappa = (2\pi/39)^2$, considered

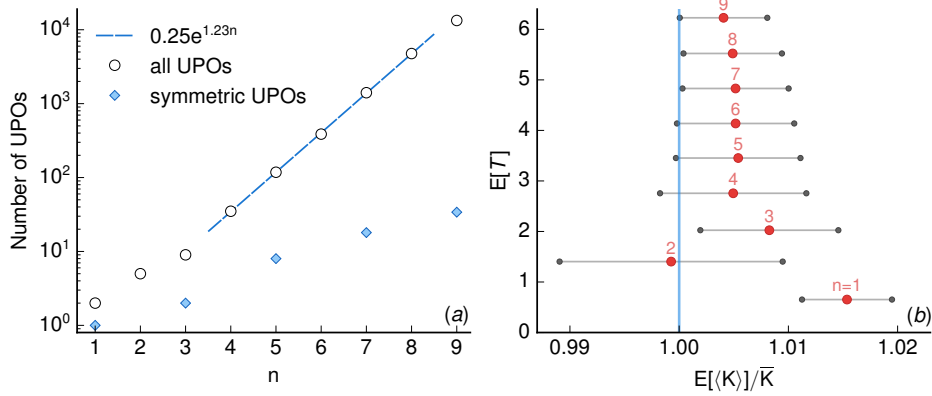


Figure 5. Panel (a): number of orbits as a function of n , open circles, and exponential fit for $4 \leq n \leq 8$. The blue diamonds denote the number of symmetry invariant UPOs. Panel (b): expectation of the normalised period averaged energy density, as a function of the expectation of the period. The horizontal bars span the standard deviation of the period average, for each n .

in what follows, the side regions of the attractor collapse and long trajectories span the center region only. The evolution is thus considerably simpler and not intermittent, facilitating the numerical search of UPOs and the interpretation of sensitivity and control results.

Truncation of (35) at $N = 32$ (twice the resolution used in [24]) results in an overall dynamic range of about six orders of magnitude. Robustness of statistical quantities to perturbations in the truncation is a subtle issue [9]. Here, the attractor structure remains intact as the resolution is increased/decreased to $N = 64/16$. However, the system appears to be structurally unstable with respect to perturbations in the control gains, as discussed in subsection 5.4.

5.1. Search results. A database of 20073 unique UPOs has been generated. The large number of orbits is used to explore the statistical distribution of key quantities across UPOs, most importantly the sensitivity with respect to control. Like for the Lorenz system, UPOs are either symmetry-invariant, or come into symmetry-equivalent pairs. Only one of the two orbits is stored in the database. We count orbits using the cycle length $n \geq 1$, the number of ‘up’ intersections with the Poincaré section $v_1 = 0, v_1(s) > 0$. For practical purposes, we only search and store UPOs with $n \leq 9$.

Figure 5-(a) shows the number of UPOs found as a function of n (open circles). An exponential fit for orbits with $4 \leq n \leq 8$ provides an estimate of the topological entropy equal to 1.23 [3]. The fact that the data follows closely the exponential fit, at least up to $n = 8$, suggests that the vast majority of UPOs of length in this range has been found. Symmetry-invariant UPOs (blue diamonds, for odd n only) are the vast minority (for $n = 9$ we found 34 over 13345).

In panel (b), the expectation of the period averaged energy density across UPOs of same n , normalised by the long-time average, is reported as a function of the period expectation (red circles). The horizontal bars show the span of the standard deviation among UPOs, for each n . Even short UPOs give remarkably accurate predictions of the average energy density.

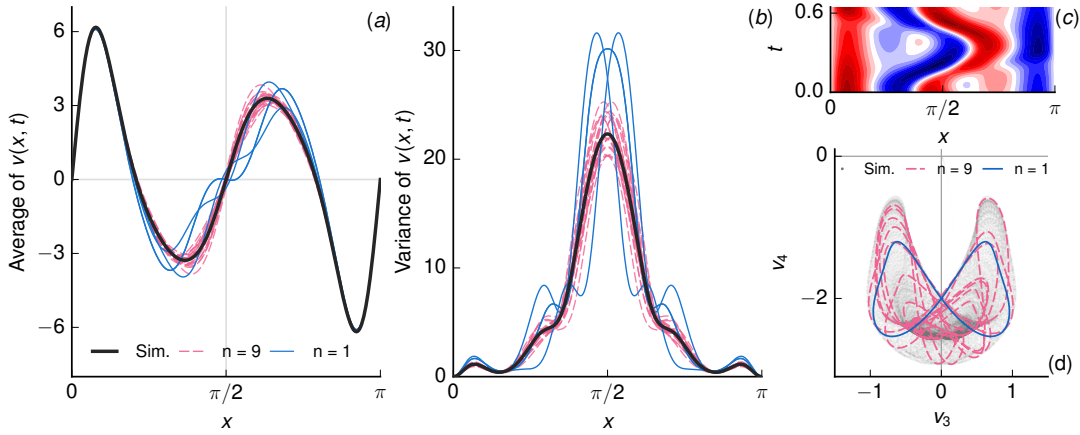


Figure 6. Panel (a): long-time and period averages of $v(x, t)$ for a long chaotic realisation and for selected UPOs of varying length. Panel (b): long-time and period variances. Panel (c): the shortest UPO (same colors as in Figure 4). Panel (d): (v_3, v_4) projections of UPOs and long chaotic realisation.

Longer orbits yield on average more accurate predictions and the standard deviation amongst orbits of same length decreases with n .

The average spatial structure of the chaotic realisation of Figure 4 is well reproduced by UPOs. This is illustrated in Figure 6, where statistics for all short ($n = 1$, solid blue lines) and a subset of ten long ($n = 9$, dashed pink lines) UPOs are compared to long-time statistics of a chaotic realisation ($T = 1000$ thick black line). Panel (a) shows the long-time and period averaged solutions, \bar{v} and $\langle v \rangle$, whilst panel (b) shows the the long-time and period variances $\overline{(v - \bar{v})^2}$ and $\langle (v - \langle v \rangle)^2 \rangle$. The UPOs reproduce particularly well the boundary layers, where the chaotic activity is weaker. These contribute in large part to the overall energy density. The second order statistics in the domain interior, are somewhat less well reproduced, with the $n = 1$ UPOs over-predicting the variance at $\pi/2$ by about 40%. As in Figure 5, more accurate predictions are obtained for longer UPOs, as these span a larger volume of the attractor and underpin longer segments of chaotic trajectories. This is illustrated in panel (d), where the projections on the subspace (v_3, v_4) of a short and a long ($n = 1/9$, blue solid/pink dashed line) UPOs are reported, together with a long chaotic trajectory (grey dots) sketching the attractor (this corresponds to the center part S_C).

The instantaneous spatio-temporal features of chaotic solutions are also well explained by UPOs. For instance, the shortest UPO ($n=1$), in Figure 6-(c) describes the “minimal” spatio-temporal “building block” of the chaotic dynamics and its time scale [8], the left-and-right motion of the wave in the domain center, while longer UPOs reproduce variations and combinations of this pattern (see e.g. the $n = 5$ UPO in Figure 14-(a)).

5.2. Sensitivity analysis: role of symmetries. We begin the sensitivity analysis results with a discussion on the role of the symmetry (38) on the sensitivity calculations. System (40) is equivariant under (38), if the vector field \mathbf{f} and the transformation S commute, i.e. if

$$(41) \quad S(\mathbf{f}(\mathbf{v}; \mathbf{G})) = \mathbf{f}(S(\mathbf{v}); \mathbf{G})$$

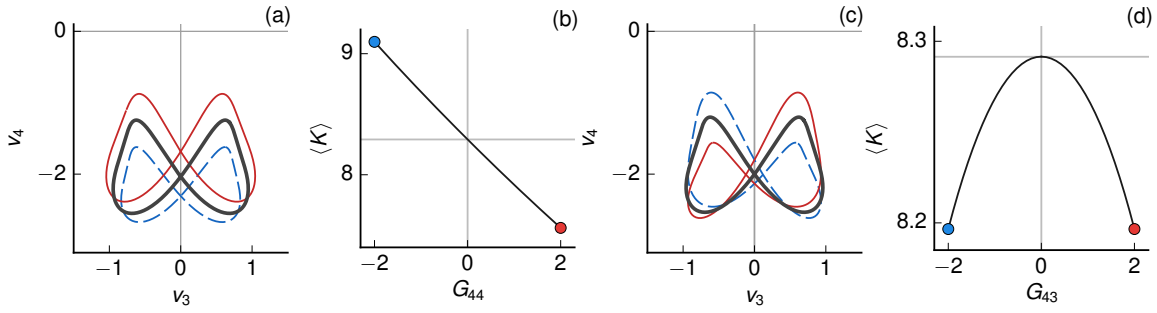


Figure 7. Panels (a) and (c): (v_3, v_4) projections of the original (thick black line) and perturbed UPOs, for perturbations of $G_{44} = \pm 2$ and $G_{43} = \pm 2$, respectively. Negative perturbations correspond to the dashed orbits. Panels (c) and (d): period averaged energy density as a function of the perturbed parameter. Circles denote the value of the parameter at which projections are shown in the corresponding panels. Note that the slope of the bifurcation curves in (b) and (d) at $G_{44}/G_{43} = 0$ equal exactly the sensitivity from the adjoint calculation.

holds. When $\mathbf{G} = 0$, (41) holds identically.

Algebraic manipulations omitted here show that (41) also holds for perturbations of the control gains G_{km} , when $k + m$ is even. When these control gains are varied, symmetry-invariant UPOs move in state space whilst preserving the symmetry. An instance is shown in Figure 7-(a), for the shortest ($n = 1$) UPO. The original UPO (thick black orbit) and two orbits obtained from continuation along G_{44} ($G_{44} = +2/-2$, solid red/dashed blue orbits, respectively) are shown, projected on the (v_3, v_4) subspace. When this parameter is varied, the UPO moves towards/away from the origin for positive/negative perturbations, whilst remaining symmetric. The period averaged energy density varies accordingly, Figure 7-(b). In practice, the displacement of UPOs in state space results in the displacement of the entire attractor and thus time averages from long chaotic trajectory vary similarly.

On the other hand, perturbations of control gains G_{km} , when $k + m$ is odd, break (41). Depending on the sign of the perturbation, the symmetry can be broken in one direction or the other. This is illustrated in Figure 7-(c), for parameter G_{43} . For any perturbation G_{43} , the perturbed orbit satisfies globally $\mathbf{v}^o(s, G_{43}) = S(\mathbf{v}^o(s, -G_{43}))$, $\forall s$. Because the energy density is invariant under the symmetry, i.e., $K(\mathbf{v}) = K(S(\mathbf{v}))$, it follows that $\langle K(G_{43}) \rangle = \langle K(-G_{43}) \rangle$ for all G_{43} . Hence, the period averaged energy density is an even function of G_{43} and the gradient $\langle K \rangle_{dG_{43}}$ must vanish identically at the origin.

The effect of symmetries on UPOs that are not symmetry-invariant is illustrated in Figure 8. Panels (a) and (c) show projections on the (v_3, v_4) subspace of the three $n = 1$ UPOs, the symmetry-invariant UPO, in (a), and the symmetry-equivalent pair, in (c). The dashed orbit in the latter panel is obtained by application of (38). Panels (b) and (d) display the gradient $\langle K \rangle_{d\mathbf{G}}$. The entries of the gradient decay exponentially with m and k , hence only the relevant part is shown. As anticipated, for the symmetric orbit the gradient with respect to gains where $k + m$ is odd, indicated by black dots in Figure 8-(b), is identically zero and the gradient display a checker board pattern. The orbits in Figure 8-(c) do not have this property and all entries are generally non zero (the figure shows the gradient of the orbit denoted by




Figure 8. Panels (b)-(d): the gradient $\langle K \rangle_{dG}$ for the two shortest UPOs ($n = 1$). Panels (a)-(c): their projections on the (v_3, v_4) subspace.

the solid line). However, the gradient obtained for the symmetric counterpart satisfies

$$(42) \quad \begin{cases} \langle K(\mathbf{v}^o) \rangle_{dG_{km}} = \langle K(S(\mathbf{v}^o)) \rangle_{dG_{km}} & \text{for } k + m \text{ even} \\ \langle K(\mathbf{v}^o) \rangle_{dG_{km}} = -\langle K(S(\mathbf{v}^o)) \rangle_{dG_{km}} & \text{for } k + m \text{ odd,} \end{cases}$$

Hence, the joint contribution of the orbits of a symmetry-equivalent pair displays the same pattern as symmetry-invariant orbits.

This result has significant implications for sensitivity analysis and optimisation of chaotic dynamics using, but not limited to, the methods discussed here. The primary reason is that many such systems possess symmetries that would lead to the same behaviour illustrated here. For instance, many fluid flow configurations of practical interest, e.g. shear flows in channels or pipes, possess continuous symmetries along the homogeneous spatial directions. Hence, a first order sensitivity analysis might not be sufficient to reveal the true influence of system parameters. A second order analysis might be necessary for optimisation over finite amplitude perturbations. The development of second order sensitivity techniques is thus an important avenue for future work.

5.3. Sensitivity analysis: statistics of gradients. Figure 9-(a) shows the expectation of the sensitivity $\langle K \rangle_{dG}$. The expectation displays the same checker board pattern observed for symmetry-invariant orbits, because opposing contributions from pairs of symmetric-equivalent orbits are included in the average.

In panels (b) and (c), normalised histograms of two components with large magnitude, $\langle K \rangle_{dG_{44}}$ and $\langle K \rangle_{dG_{66}}$, are reported. Other components display similar behaviour (distributions for gradients components for odd $k + m$ have zero odd moments). Panel (d) shows the scatter of these two components and three contours of the joint probability density function obtained using a kernel density estimation technique with bandwidth 0.005. Analysis of these and other distributions shows that the vast majority of UPOs predicts approximately the same sensitivity, i.e. gradients are tightly aligned around the expectation of panel (a), denoted by lines in the other panels. We attribute this behaviour to the fact that the attractor has a rather simple structure, the system displays a limited range of dynamical behaviour and all




Figure 9. Panel (a): gradient expectation. Panel (d): scatter of components $\langle K \rangle_{dG_{44}}$ and $\langle K \rangle_{dG_{66}}$ (black dots). In red, contour lines of the joint probability density function of the two components, for density 1, 10, and 100. Panels (b-c): normalised histograms of $\langle K \rangle_{dG_{44}}$ and $\langle K \rangle_{dG_{66}}$, respectively.




Figure 10. Scatter $\langle K \rangle_{dG_{44}}$ and $\langle K \rangle_{dG_{66}}$ across UPOs of different cycle length n . The orbits denoted by larger symbols ($n=1$ and 6) are those for which continuation analysis is performed in [Figure 11](#).

UPOs describe variations of the same spatio-temporal pattern. It is argued that in systems with larger number of active degrees of freedom, with higher-dimensional and more convoluted attractors, the scatter of sensitivity among UPOs will be larger.

[Figure 10](#) shows the scatter of the same two gradient components for orbits of cycle lengths $1 \leq n \leq 9$. UPOs with $n \leq 5$ appear to be mostly concentrated around the expectation (blue lines), although for $n \geq 6$ a handful of orbits displays much larger or smaller gradient than average, producing the two long tails of the density distribution of [Figure 8-\(d\)](#).




Figure 11. Panel (a): the two $n = 6$ UPOs with largest/smallest sensitivity $\langle K \rangle_{dG_{44}}$. Continuation analysis along G_{44} for these two orbits is shown in panel (b). Here, the two lines tangent to the bifurcation curves have slope equal to what illustrated in the relevant panel in figure (Figure 10), with numerical value of the gradient denoted by the label. Panel(c): bifurcation curve for the shortest UPO ($n = 1$), displayed in the inset in the same panel.

We perform continuation along G_{44} for the two $n = 6$ UPOs with extreme value of the sensitivity $\langle K \rangle_{dG_{44}}$ (denoted in the relevant panel of Figure 10 with larger symbols). One would expect that UPOs with such large difference in sensitivity have significantly different characteristics. However, the projections of these two UPOs on the (v_3, v_4) subspace, reported in Figure 11-(a), show that they are actually similar to each other. The continuation analysis, panel (b), reveals that the two orbits actually sit on the upper and lower branches of the same bifurcation curve, originating in a saddle-node bifurcation just before $G_{44} = 0$. The slope of the lines tangent to the bifurcation curve is the sensitivity predicted by the adjoint analysis at $G_{44} = 0$. These two UPOs yield larger/smaller than average sensitivity simply because they sit close to the bifurcation. They survive the continuation process for larger positive values of G_{44} than shown in the figure. Asymptotically, the branches tend to have slopes that agree well with the expectation and with UPOs that are not about to disappear. For instance, for the shortest $n = 1$ UPO, panel (c), bifurcation occurs for large absolute values of G_{44} , and the gradient is physically relevant. It is argued that other pairs of UPO with larger/smaller than average gradient for $n \geq 6$ display the same behaviour. This has not been checked, as tracing bifurcation curves for all UPOs would be too expensive. This is aggravated by the fact that there are many bifurcation parameters in the present problem and saddle-node bifurcations might lurk only in some directions in parameter space, biasing the sensitivity analysis result only in those directions.

These results suggests that care should be taken in interpreting the sensitivity of UPOs. Near a saddle-node bifurcation, the nonlinear BVP (13) is ill-conditioned: small changes in the parameters result in ‘large’ changes in the solution, distorting the sensitivity results. On the positive side, however, the ill-conditioning is unrelated to the length of the trajectory, e.g. increasing exponentially with time as for the IVP (1), [41]. In addition, such bifurcations can be identified with relative ease using standard bifurcation techniques, or with second order sensitivity analysis, as second derivatives would be large close to bifurcations.




Figure 12. Geometry of gradients and projections.

Further gradient statistics relevant for control are now reported. We analyse the geometry and distribution of gradient vectors in parameter space, and quantify their alignment with the gradient expectation $E[\langle K \rangle_{dG}]$, the equally weighted average of UPO gradients. The aim is to obtain a description of the scatter of UPO gradient vectors in parameter space. For the i -th UPO, the gradient $\langle K \rangle_{dG}^i$ is a vector in parameter space pointing in the direction of largest increase of the period averaged cost. Considering Figure 12, let

$$(43) \quad \mathbf{d}^E = E[\langle K \rangle_{dG}]^\top / \|E[\langle K \rangle_{dG}]\| \quad \text{and} \quad \mathbf{d}^i = \langle K \rangle_{dG}^{i\top} / \|\langle K \rangle_{dG}^i\|$$

be the unit vectors pointing along the gradient expectation and the gradient of the i -th orbit, respectively. The projections

$$(44) \quad \langle K \rangle_{dG}^i \cdot \mathbf{d}^E \quad \text{and} \quad \langle K \rangle_{dG}^i \cdot \mathbf{d}^i = \|\langle K \rangle_{dG}^i\|$$

define the linearised change in the period averaged energy density that would be obtained if control gains were set to \mathbf{d}^E and \mathbf{d}^i , respectively. In other words, these projections define the slope $\langle K \rangle_{d\lambda}$ of an hypothetical bifurcation curve $\langle K(\lambda) \rangle$ when control parameters are varied as $-\lambda \mathbf{d}^E$ and $-\lambda \mathbf{d}^i$, respectively, at $\lambda = 0$. The angle $\theta^i = \arccos(\mathbf{d}^i \cdot \mathbf{d}^E)$ quantifies the misalignment angle between the gradient of a particular UPO and the expectation.

Normalised histograms for the two projections (44) and the angle θ are reported in panels (a), (c) and (d) of Figure 13, respectively. A scatter plot of the projections is reported in panel (b). The two marginal histograms in (a-c) are quite narrow, reflecting the compactness of the histograms of individual gradient components in Figure 9.

These two distributions have weak tails that are better seen in the scatter plot as two tails that extend to large values of the gradient norm $\langle K \rangle_{dG}^i \cdot \mathbf{d}^i$. The two closed symbols in (b) correspond to the two $n = 6$ orbits described in Figure 11, suggesting that orbits that are close to bifurcation have typically large gradient norm and can point in direction in parameter space quite away from the expected direction, especially orbits on the lower tail of the scatter plot. The distribution of angles in (d) shows that the mode occurs below 10 degrees. More than 90% of the gradients are within 30 degrees from the expectations and gradients with larger angles correspond to UPOs close to bifurcation. That such a tight distribution is found is remarkable and a clear indication that most UPOs predict a similar sensitivity with respect to actuation. If gradient had random directions, they would be mostly orthogonal in such a high dimensional parameter space ($N^2 = 1024$).

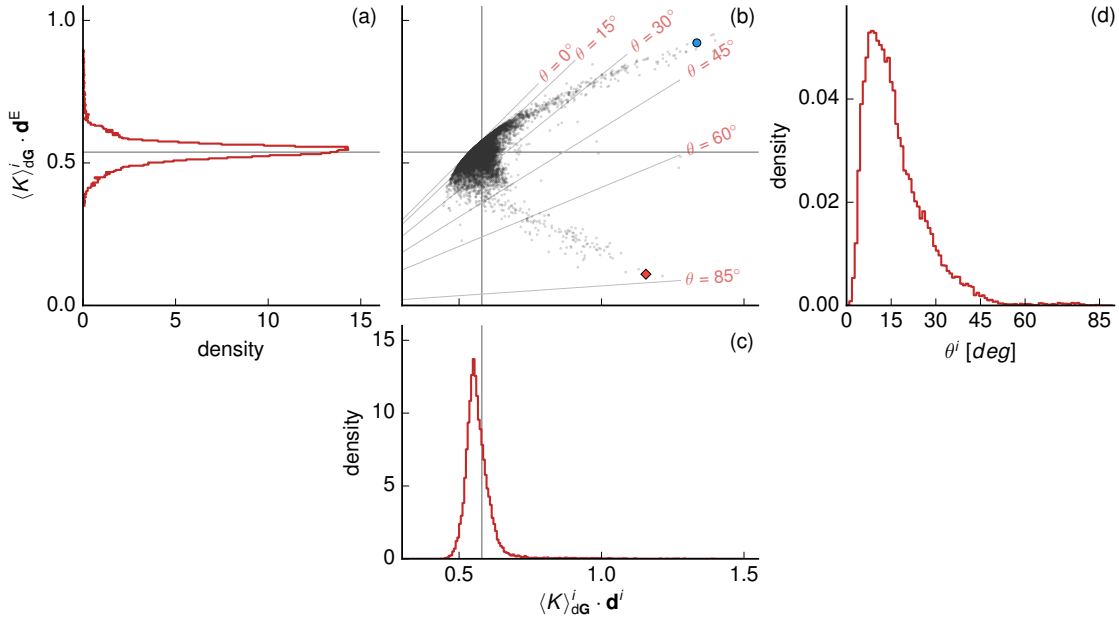


Figure 13. Panels (a), (c) and (d): normalised histograms of the projections (44) and angle θ , respectively. Panel (b): scatter of the two projections. Note that the quantity of panel (c) is also the norm of the gradient.

5.4. Feedback control. We now illustrate relevant mechanisms of control identified by the sensitivity analysis and discuss how these affect the dynamics of controlled chaotic realisations. Figure 14 present such results illustrated for an example UPO, one of the eight symmetric $n = 5$ UPOs, with period $T \approx 3.33$. The space-time evolution of this orbit is shown in panel (a). We solve the adjoint problem for this orbit, obtain the gradient $\langle K \rangle_{dG}$ and then construct a feedback controller $-\lambda \langle K \rangle_{dG} / \|\langle K \rangle_{dG}\|$. Panel (b) shows the actuation spatial distribution that would result upon application of the controller with $\lambda = 1$ to the solution $v(x, t)$ in (a).

In the region adjacent to the domain boundaries, actuation is almost stationary and has opposite sign to the solution of panel (a). Actuation in this region reduces the amplitude of the boundary layer and the associated peak of the energy. In the domain center, where most of the chaotic activity occurs, actuation has a lower intensity and counters the left-right wavy motion in this region. Here the solution changes more significantly under actuation. Panel (c) shows the spatio-temporal distribution of the sensitivity

$$(45) \quad v_{\partial\lambda}^*(x, t) = \sum_{k=-N}^N (v_{\partial\lambda}^*)_k(t) e^{ikx}$$

i.e. the linearised change in the solution $v(x, t)$ under actuation. The Fourier coefficients of the sensitivity $(v_{\partial\lambda})_k$ are obtained from a finite difference approximation along the descent direction, and the least squares technique discussed in subsection 3.1 is then applied. The structure of $v_{\partial\lambda}^*(x, t)$ shows that actuation induces significant changes in the dynamics of the wavy motion in the domain center, and less at the domain boundaries, in spite of the strength

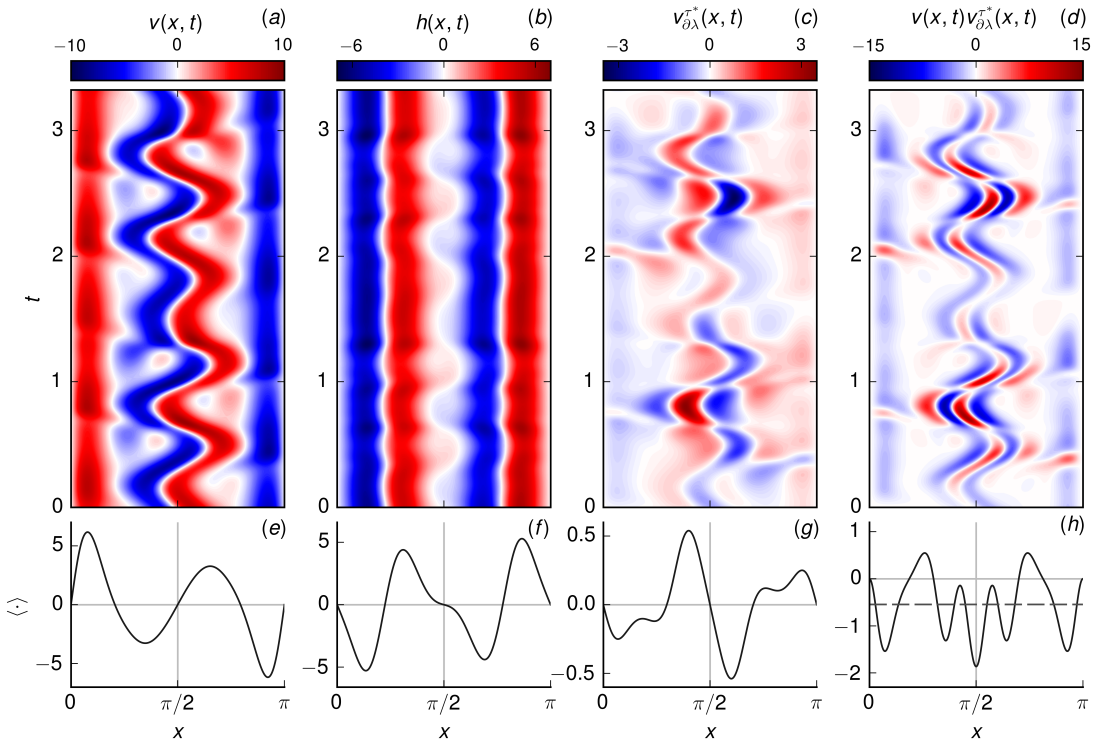


Figure 14. Control mechanisms identified by the sensitivity analysis, illustrated on a symmetric $n = 5$ UPO, panel (a). Panel (b): spatio-temporal control distribution. Panel (c) least-squares sensitivity of the solution. Panel (d): linearised change in the energy $v(x,t)^2/2$. The bottom panels show the time-averaged, space-dependent quantity reported in the corresponding upper panel.

of the control in the boundary region. Note that the UPO period increases under actuation, i.e. control slows down the chaotic activity and increases the time scale of the dynamics (this seems to be a general behaviour as $\omega_{\text{dG}} \cdot \langle K \rangle_{\text{dG}}^{\top} > 0$ for the vast majority of UPOs).

The space-time distribution of

$$(46) \quad \frac{\partial(v^2(x,t)/2)}{\partial\lambda} = v(x,t)v_{\partial\lambda}^{\tau^*}(x,t),$$

is shown in Figure 14-(d). This quantity is the linearised change in the integrand of the cost function (34). The space and period average of this quantity coincides with the norm of the gradient $\langle K \rangle_{\text{dG}}$, indicated in panel (h) with a dashed line. One can notice that the cost does not strictly reduce everywhere in time and space, but there are regions/instants where the energy increases under control. The period average of this change, panel (h), shows that the reduction is consistent near the domain boundaries, but also in the domain center, whereas there are two regions, near $x = \pi/2 \pm \pi/4$ where on average the energy increases.

We now discuss the effects of actuation on chaotic realisations. We construct the controller $-\lambda \mathbf{d}^{\text{E}}$, i.e. we vary the control gains in the direction where, on average, the UPOs are most sensitive to. This is justified empirically by the fact that most UPOs are tightly aligned about this direction. To illustrate the response of the system, we vary λ in the range $[-2, 2]$,

although it is very likely that many orbits do not survive such a large perturbation. We then run numerical simulations of the controlled system and monitor the long time ($T = 5000$ time units) averaged energy density as a function of the control strength λ . The results are reported in Figure 15, where the vertical axis has been normalised by the long-time average of the uncontrolled system. The dashed lines have slope equal to $\|E[\langle K \rangle_{dG}]\|$, the expectation of the reduction of the period averaged cost across UPOs, per unit strength of control λ varying the gains in the direction \mathbf{d}^E . This quantity coincides with mean of the distribution of Figure 13-(a) and is approximately equal to 0.538.




Figure 15. Variation of the normalised long-time average as a function of λ , the strength of the control, from numerical simulation of the controlled PDE (black dots). The dashed line has slope equal to the norm of the gradient expectation, obtained from UPOs. The dotted line near the origin defines the actual slope of the data points. The two right panels show close-ups of the results in the regions indicated in panel (a).

The slope associated to the expectation of the UPOs gradients gives a good overall representation of the linear variation of the long-term dynamics of chaotic solutions induced by the feedback control, especially for negative λ . Closer analysis near $\lambda = 0$, however, shows that the data points are distributed around a straight line (the dotted line) with larger slope, about 0.75, i.e. the local response $(\overline{K})_{\partial\lambda}$ (from a data fit) is about 40% larger than that predicted on average by the UPOs. That the match is not exact is to be expected because the equally weighted average in the expectation has no theoretical basis, but it would be justified, empirically, by the fact that gradients are tightly aligned in parameter space around the expectation. Still the variation of the long-time average with λ is much greater than what predicted by the vast majority of UPOs, much more than for the Lorenz system. In particular, the value 0.75 sits on the right tail of the distribution of Figure 13-(a). Only a few UPOs reach such large value.

Our speculation on this difference is that the system undergoes significant structural changes around $\lambda = 0$. These become particularly evident around $\lambda = 0.1$. For $\lambda < -0.5$ the dynamics do not vary significantly and the response of the UPOs appear to give a very good overall prediction of the response of the controlled chaotic dynamics. For $\lambda < 0$, the aver-

age can lie on two separate “branches”. On the lower “branch”, the dynamics can land and remain, depending on the particular value of λ and the initial condition, on the side regions $S_{L/R}$ of an attractor with similar structure to that of $\kappa = (2\pi/38.5)^2$ in Figure 4-(b), arising $\kappa = (2\pi/39)^2$ as a result of control. Long simulations (thousand of time units) indicate that these are disconnected from the “center” part S_C (associated to the upper “branch”), which justifies the fact that the two “branches” in Figure 15 are distinct. The points on the lower branch are the vast minority, though.

For $\lambda > 0$, with the exception of a narrow range around $\lambda \approx 1.2$, the dynamics become highly intermittent. The probability of trajectories landing on the “side parts” of the attractor along the duration of the integration increases, and the distribution of the data points for the average bends towards to lower “branch”. The much larger slope of the data points for the average in the region $0.1 < \lambda < 0.3$ is thus mostly attributed to the increasing intermittent character of the dynamics under stronger actuation and the increasing preference for the system to land of the “side” parts, where the energy is lower, rather than a genuinely larger sensitivity of the dynamics on the “center” part of the attractor with respect to actuation.

6. Conclusions. In this paper we have discussed a robust adjoint sensitivity method for chaotic dynamical systems. The method is based on the idea of formulating the adjoint problem on time periodic trajectories, rather than on open trajectories. Both types of trajectories are unstable, but in the periodic case the adjoint problem has periodic boundary conditions in time, i.e. it becomes a two-point boundary value problem. The adjoint solution remains bounded at all times and does not exhibit the typical unbounded exponential growth observed in traditional adjoint methods on open unstable trajectories. This special structure enables the exact sensitivity of period averaged statistics with respect to problem parameters to be obtained, *regardless of the length and stability characteristics of the orbit*. Sensitivity information from these trajectories can then reveal important physical features of the parametrised system that would be otherwise hardly accessible.

The solution of the adjoint problem on a periodic orbit requires a strategy different to that used in traditional adjoint analysis for unsteady problems. The backward in time integration is not required and the governing and adjoint equations can in principle be marched forward simultaneously in time. However, methods that have been traditionally used to solve two-point boundary value problems are needed to enforce the time periodic boundary conditions. On the practical side, however, algorithms and software used for the search of periodic orbits can be largely reused.

We have demonstrated the method using two well known low-dimensional chaotic systems: the Lorenz equations at the standard parameters and the Kuramoto-Sivashinsky equation, in a weakly chaotic regime. For the latter problem we have discussed a case that is relevant to applications, e.g. the sensitivity of averages with respect to feedback control. For both problems, we have provided a statistical description of the sensitivity over a large number of UPOs, i.e. over a database of $\sim 10^5$ and $\sim 4 \times 10^4$ UPOs, respectively. In both cases, we have found that the scatter of sensitivity across UPOs is rather limited, i.e. most orbits predicted similar response to parameter perturbations. This might be attributed to the relatively low-dimensional nature of these systems. We hypothesise that an increasingly larger scatter will be observed in systems with increasingly larger number of active degrees of freedom, e.g.

turbulent fluid systems at large Reynolds number.

We have attempted to estimate the gradient using linear regression of long time averages from chaotic realisations of the perturbed equations and compared the value with the distribution of period averaged gradients from UPOs. The agreement is only fair, not in absolute terms (2% error on the Lorenz problem), but relatively. The estimated gradient falls far from where the vast majority of UPOs (in particular the short cycles) points to, although it still falls within the tails of the distribution. This might be speculatively justified by subtle nonlinear dynamics issues, e.g. intermittency, and to the objective difficulty in obtaining properly converged statistics from long simulations. Further analysis is required. In any case, the scatter of UPO gradients is lower than that resulting from ensemble average techniques [11] or the Least Squares Shadowing method [39, 41].

There are some potential caveats and open questions that need to be highlighted. Firstly, one needs to find a periodic trajectory before its sensitivity can be computed. This has been a relatively straightforward step for the two problems discussed here but it can be an objective challenge for high-dimensional systems. Our plan is utilise these techniques to tackle outstanding control/design problems for three-dimensional Navier-Stokes turbulence, in light of the recent advances in understanding the role and dynamics of exact coherent structures [27, 20, 13, 21, 7, 6]. Although many such structures have now been found, the scalability of the approach to real-world turbulent flows remains an open question. In practice, the search of periodic trajectories become increasingly more challenging as the Reynolds number and domain size are increased [7]. Sensitivity techniques that do not require temporal periodicity, such as Least-Squares Shadowing [41], might be better suited in such a scenario. In addition, a sensitivity method for systems with continuous symmetries, typical of shear flows, remains to be developed.

Secondly, a first order sensitivity analysis might not be sufficient, or not informative enough, for problems with symmetries. In such cases, the sensitivity of period averages for symmetry-invariant orbits, or the joint contribution of groups of orbits where the group is invariant, is identically zero. From an optimisation view point, the averages exhibits saddles along these directions in parameter space and thus second order information becomes necessary for effective optimisation over finite parameter perturbations.

A third caveat is related to the fact that periodic orbits (like any other trajectory) can undergo bifurcations in parameter space. The main issue is not the obvious one regarding the differentiability of period averages *at* a bifurcation point, but rather the fact that the sensitivity for an orbit that is *close* to bifurcating might be much larger/smaller than expected, or physically sound. This might lead to erroneous conclusions about the behaviour of the overall system. For instance, an orbit that is about to disappear/appear in, e.g. a saddle-node-type bifurcation, as a certain parameter is varied will yield a sensitivity with respect to that parameter tending to positive or negative infinity. Clearly, the availability of a large number of periodic orbits, a bifurcation study or a second order sensitivity analysis might alleviate this issue.

Fourthly, there are subtle “pathological” features of non-hyperbolic systems, the vast majority of systems in the applied sciences, that suggest care in interpreting the sensitivity results. The primary feature is that statistics might not be smooth functions of the parameters, hence not differentiable with respect to parameters, due to the lack of structural stability

[16, 10]. What is then the interpretation of the sensitivity of an UPO, if the attractor structure in which this orbit resides can suddenly collapse and statistics change abruptly for infinitesimal parameter perturbations? Another issue is that related to intermittency, as illustrated here for the Kuramoto-Sivashinsky equation. If the degree of intermittency varies as parameters are varied, the statistics of the system might change in a way unrelated to what predicted by UPOs.

We wish to address these issues in future work.

Acknowledgements. The author wishes to thank Dr Ati Sharma, Mr Arslan Ahmed and Dr Yongyun Hwang for the fruitful discussions. The Institute of Mathematics and its Applications (ima.org.uk) is also gratefully acknowledged for supporting attendance to the Tenth International Symposium on Turbulence and Shear Flow Phenomena (July 6-9th, 2017, Chicago, IL, USA) where part of the present results have been presented.

Appendix A. UPO search algorithm. We used a Newton-Raphson-type technique traditionally known as *quasi-linearisation* in the literature [2] to solve the nonlinear periodic boundary value problem (BVP), equation (13), and find UPOs numerically. The method is analogous to that reported elsewhere, e.g. [36, 23]. It is reported here primarily to highlight the similarities between the linear time periodic BVP arising in the Newton-Raphson iterations and the linear time periodic BVPs for tangent and adjoint sensitivity (16), (20) and (23). This similarity enables most of the numerical techniques (and software) to be reused for both search and sensitivity analysis. The numerical approach used to solve these linear time periodic BVPs is described in Appendix B.

The method starts with a trial solution $\{\mathbf{x}_k^o(s), \omega_k\}$, $k = 0$, consisting of a frequency guess ω_k and a state space loop $\mathbf{x}_k^o(s)$, a smooth candidate periodic trajectory that does not satisfy the governing equations, that is, where the residual

$$(47) \quad \mathbf{r}_k(s) = \omega_k \dot{\mathbf{x}}_k^o(s) - \mathbf{f}(\mathbf{x}_k^o(s))$$

is generally non zero, although $\mathbf{x}_k^o(0) = \mathbf{x}_k^o(2\pi)$. At each iteration $k = 0, 1, \dots$ the trial solution is adjusted by calculating a correction $\{\mathbf{y}_k(s), \eta_k\}$ such that the new trial solution $\{\mathbf{x}_{k+1}^o(s), \omega_{k+1}\} = \beta_k \{\mathbf{x}_k^o(s) + \mathbf{y}_k(s), \omega_k + \eta_k\}$ solves the nonlinear periodic BVP exactly, i.e.

$$(48) \quad \mathbf{r}_{k+1}(s) = (\omega_k + \eta_k)(\dot{\mathbf{x}}_k^o(s) + \dot{\mathbf{y}}_k(s)) - \mathbf{f}(\mathbf{x}_k^o(s) + \mathbf{y}_k(s)) = 0$$

where $0 < \beta_k \leq 1$ is a damping factor discussed below. Assuming that the correction is small in some sense, linearising the vector field around $\mathbf{x}_k^o(s)$, and neglecting second order terms leads to the linear time periodic BVP

$$(49a) \quad \begin{cases} \omega_k \dot{\mathbf{y}}_k(s) = \mathbf{f}_{\partial \mathbf{x}}(s) \cdot \mathbf{y}_k(s) - \mathbf{r}_k(s) - \eta_k \dot{\mathbf{x}}_k^o(s) \\ (49b) \quad \mathbf{y}_k(0) = \mathbf{y}_k(2\pi) \\ (49c) \quad 0 = \mathbf{y}_k(0)^\top \cdot \mathbf{f}(\mathbf{x}_k^o(0)) \end{cases}$$

where $\mathbf{f}_{\partial \mathbf{x}}(s) = \mathbf{f}_{\partial \mathbf{x}}(\mathbf{x}_k(s))$ is the Jacobian operator evaluated on the current trial solution and the quantities $\{\mathbf{y}_k(s), \eta_k\}$ are the unknowns. The last row defines a phase locking condition to break the translational invariance along the orbit and ensures that the temporal discretisation

of (49) does not result in a singular problem. The iterations converge quadratically, but convergence is only local. In this work we used a backtracking technique to improve robustness with respect to poor initial guesses, obtained from scanning long simulation for near recurrences. This consisted in recursively halving the damping factor β_k at each iteration from $\beta_k = 1$ until $\langle |\mathbf{r}_{k+1}(s)|^2 \rangle < \langle |\mathbf{r}_k(s)|^2 \rangle$. Typically, we observed slow convergence in the initial iterations, $\beta_k \ll 1$, and then fast quadratic convergence near the solution. We defined convergence when $\langle |\mathbf{y}_k(s)|^2 \rangle < 10^{-\alpha}$ and $\langle |\mathbf{r}_k(s)|^2 \rangle < 10^{-\alpha}$ and $|\eta^k| < 10^{-\alpha}$, with $\alpha = 13$, typically. Stagnation of the average residual norm and decreasing small damping factors β_k were used to detect and stop prematurely iterations from initial guesses that would not converge.

Appendix B. Numerical solution of linear time periodic BVPs. The Newton-Raphson iteration problem (49), the tangent sensitivity problem (16), and the adjoint problems (20) and (23) can be all recast into the same standard form for linear time periodic BVPs

$$\begin{aligned} (50a) \quad & \omega \dot{\mathbf{z}}(s) = \mathbf{A}(s) \cdot \mathbf{z}(s) + \mathbf{b}(s) + \kappa \mathbf{c}(s), \quad s \in [0, 2\pi] \\ (50b) \quad & \mathbf{z}(0) = \mathbf{z}(2\pi) \\ (50c) \quad & \alpha = \mathbf{z}(0)^\top \cdot \mathbf{g} \end{aligned}$$

where $\mathbf{z}, \mathbf{b}, \mathbf{c}$ and $\mathbf{g} \in \mathbb{R}^N$, $\mathbf{A} \in \mathbb{R}^{N \times N}$ and $\omega, \kappa, \alpha \in \mathbb{R}$, and where the unknowns are the solution $\mathbf{z}(s)$ and the parameter κ . We use a finite difference temporal discretisation of (50) over the M -point uniform mesh $s_i = i\Delta s, i = 0, \dots, M-1$, with $\Delta s = 2\pi/M$, and seek the solution $\hat{\mathbf{z}} = (\mathbf{z}(s_0), \mathbf{z}(s_1), \dots, \mathbf{z}(s_{M-1}))^\top \in \mathbb{R}^{NM}$ at these points. The fourth-order accurate centered stencil

$$(51) \quad \dot{\mathbf{z}}(s_i) \approx \frac{1}{12\Delta s} \left(\mathbf{z}(s_{i-2}) - 8\mathbf{z}(s_{i-1}) + 8\mathbf{z}(s_{i+1}) - \mathbf{z}(s_{i+2}) \right), \quad i = 0, \dots, M-1.$$

is used to approximate the loop derivatives, using the periodicity $\mathbf{z}(s_{i+M}) = \mathbf{z}(s_i)$ as appropriate. This is analogous to the approach used in [23]. Discretization leads to the set of linear algebraic equations

$$(52) \quad \chi(\mathbf{z}(s_{i-2}) - 8\mathbf{z}(s_{i-1}) + 8\mathbf{z}(s_{i+1}) - \mathbf{z}(s_{i+2})) - \mathbf{A}(s_i) \cdot \mathbf{z}(s_i) - \kappa \mathbf{c}(s_i) = \mathbf{b}(s_i)$$

for $i = 0, \dots, M-1$ and with $\chi = \omega/(12\Delta s)$. Addition of the phase locking constraint (50c) results in a bordered square system of $NM + 1$ linear equations

$$(53) \quad \begin{pmatrix} \mathbf{D} & \hat{\mathbf{c}} \\ \hat{\mathbf{g}}^\top & 0 \end{pmatrix} \begin{pmatrix} \hat{\mathbf{z}} \\ \kappa \end{pmatrix} = \begin{pmatrix} \hat{\mathbf{b}} \\ \alpha \end{pmatrix}$$

In (53), $\hat{\mathbf{g}} = (\mathbf{g}, \mathbf{0}, \dots, \mathbf{0})^\top \in \mathbb{R}^{NM}$ and the matrix $\mathbf{D} \in \mathbb{R}^{NM \times NM}$, defined as

$$(54) \quad \mathbf{D} = \begin{bmatrix} -\mathbf{A}_0 & 8\chi\mathbf{I}_N & -\chi\mathbf{I}_N & & & \chi\mathbf{I}_N & -8\chi\mathbf{I}_N \\ -8\chi\mathbf{I}_N & -\mathbf{A}_1 & 8\chi\mathbf{I}_N & -\chi\mathbf{I}_N & & & \chi\mathbf{I}_N \\ \chi\mathbf{I}_N & -8\chi\mathbf{I}_N & -\mathbf{A}_2 & 8\chi\mathbf{I}_N & -\chi\mathbf{I}_N & & \\ & & & \ddots & & & \\ & & \chi\mathbf{I}_N & -8\chi\mathbf{I}_N & -\mathbf{A}_3 & 8\chi\mathbf{I}_N & -\chi\mathbf{I}_N \\ \chi\mathbf{I}_N & & \chi\mathbf{I}_N & -8\chi\mathbf{I}_N & -\mathbf{A}_{M-2} & 8\chi\mathbf{I}_N & \\ 8\chi\mathbf{I}_N & \chi\mathbf{I}_N & & \chi\mathbf{I}_N & -8\chi\mathbf{I}_N & -\mathbf{A}_{M-1} & \end{bmatrix}$$

with $\mathbf{A}_i = \mathbf{A}(s_i) \in \mathbb{R}^{N \times N}$ and \mathbf{I}_N the identity matrix of appropriate size, arises from the temporal discretisation and has a banded structure with $2N \times 2N$ blocks on the corners arising from the periodic boundary conditions (50b). The system (53) is well conditioned by the additional row arising from the phase locking constraint, even when \mathbf{D} is singular, i.e. on a periodic orbit. For the adjoint problems (20) and (23), where technically the term $\kappa \mathbf{c}(s)$ in (50) is zero, we use the same approach and check that κ is zero to machine accuracy.

We solve the bordered system (53) using specialised dense linear algebra techniques. The mixed block elimination method BEM is used [15]. The algorithm is backward stable regardless of the condition number of \mathbf{D} , if a backward stable solver is available for \mathbf{D} and its transpose. We exploit the cyclic block tri-diagonal structure of \mathbf{D} (actually penta-diagonal, but blocks are lumped together in groups of four for simplicity) and used the algorithm described in Reference [4] for solving the linear systems arising in the BEM algorithm. This uses at its core the Sherman-Morrison-Woodbury formula to handle the corner blocks. LU factorisation and solution of banded linear systems (the banded body of \mathbf{D}) with bandwidth $4N + 1$ and M rows is performed using LAPACK implementations of `gbtrf` and `gbtrs`, respectively. The vast majority of the computational time is spent in calls to these two functions, hence the overall method benefits from using efficient implementations. Note that the factorisation of the banded body of \mathbf{D} can be reused for the solution of the homogeneous and inhomogeneous adjoint problems (20) and (23).

Solving (50) with the proposed finite difference approach becomes impractical for high-dimensional systems, such as, for instance, discretisations of Navier-Stokes fluid problems in three dimensional domains. In particular, the storage of the “space-time” trajectory $\hat{\mathbf{z}}$ can be problematic, but more importantly the construction and solution of problem (53) can be a formidable challenge. A better approach in this scenario, would be to solve the relevant BVPs using shooting techniques [2, 43], greatly reducing the storage burden.

Appendix C. Graph-based algorithm to construct a set of unique orbits. Different initial guesses might converge to solutions that define essentially the same orbit. In some instances, solutions might also describe more than one loop around the same orbit. From a numerical viewpoint, distinguishing these duplicates is not immediate. Here, we describe an efficient graph-based algorithm, scalable to hundreds of thousands orbits, to identify and remove such duplicates.

Step 1: Calculation of orbit invariant features. A vector of invariant features $\mathbf{h} \in R^{N_h}$ is first calculated for each orbit. The features can be thought of as quantities that characterise uniquely the location and structure of a periodic trajectory in state space and are invariant under time translation. In this work, we used the orbit barycenter $\mathbf{h} = \langle \mathbf{x}^o(s) \rangle$ over the first six state variables for the Kuramoto-Sivashinsky problem and the three state variables for the Lorenz equations. If needed, the invariance can be extended to other symmetries, to avoid storing more than once the same orbit (this was done for the KSEq by taking the absolute value of the orbit barycenter). This procedure leads to a set of M points $\mathcal{H} = \{\mathbf{h}^i\}_{i=1}^M$ in the “feature space”, where M is the number of converged solutions, Figure 16-(a). The error $\epsilon_{\mathbf{h}}$ on these invariants should be estimated. Here, we refined the temporal discretisation and used a Richardson extrapolation technique to estimate the error. Integrals were computed using a composite trapezoidal rule, which converges exponentially for periodic functions [35]. This

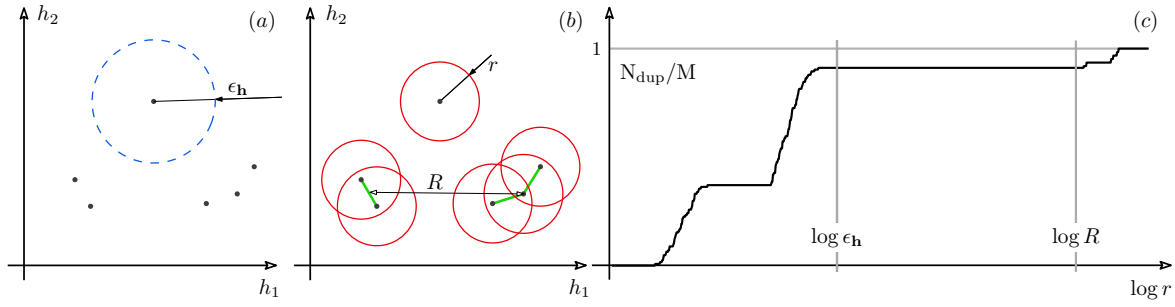


Figure 16. Panel (a): invariant features are calculated for each orbit up to an error estimate ϵ_h , resulting in a collection of points in a N^h dimensional space. Panel (b): a connectivity graph is built (green segments) by connecting points closer than a given distance r . The connected components of this graph are extracted, and the orbit with the shortest period is saved whilst duplicates are removed. Panel (c): typical effect of ball radius r on the number of duplicates identified.

ensures that the error in the calculation of period averages was due to the truncation error of the nonlinear BVP, rather than to the numerical quadrature.

Step 2: Construction of connectivity graph using tree-based nearest neighbour search.

A connectivity graph is constructed by identifying the adjacency sets

$$(55) \quad \mathcal{N}^i = \{\mathbf{h}^j \in \mathcal{H} : \text{dist}(\mathbf{h}^j, \mathbf{h}^i) < r, i \neq j\},$$

the sets of feature vectors that are contained in a ball of radius r around \mathbf{h}^i , where $\text{dist}(\cdot, \cdot)$ is an appropriate distance function, Figure 16-(b). Using an adjacency set representation allows for efficient implementations of graph traversal algorithms. To build such sets we used an efficient space-partitioning data structure, a KD-tree, to perform the nearest neighbours searches. A KD-tree is a tree-based data structure that recursively partitions with a hyper-plane the (feature) space into two sub-partitions, each containing half the points (feature vectors), until all points are exhausted. Building a KD-tree out of M points has $\mathcal{O}(M \log(M))$ complexity, whereas searching for known points in a given range around any point in the feature space has $\mathcal{O}(N_h \cdot M^{1-1/N_h})$ complexity. Hence, this structure enables significantly more efficient searches than a round-robin approach when $M \gg 2^{N_h}$.

Step 3: Identification of connected components. We use a depth-first-search traversal algorithm to identify the connected components from the connectivity graph. Connected components identify clusters of solutions that have ‘similar’ invariant features and thus describe the same UPO. This operation has, in total, $\mathcal{O}(M)$ complexity. For each connected component in the graph, the orbit with the shortest period is saved while the all others are tagged as duplicates and removed. This results in a set of unique orbits.

Step 4: Effect of ball radius r . Because the distribution of solutions in a high-dimensional feature space might be hard to visualise, we study the effect of the ball radius r on the number of duplicates N_{dup} found. Figure 16-(c) shows the typical behaviour of the algorithm, applied to a large dataset of UPOs found by converging initial guesses from a long chaotic

Table 2

Convergence of UPO characteristics and their sensitivity at $\rho = 28$. Converged digits are underlined.

Δt	M	T	$T_{d\rho}$	$\langle z \rangle$	$\langle z \rangle_{d\rho}$
0.020000	54026	<u>1080.4469534509</u>	-26.1239369819	23.5498909445	<u>1.0166218037</u>
0.010000	108052	<u>1080.5131786996</u>	-26.1247514912	23.5517275703	<u>1.0165702500</u>
0.005000	216104	<u>1080.5173671241</u>	-26.1248030601	23.5518442606	<u>1.0165669625</u>
0.002500	432208	<u>1080.5176296848</u>	-26.1248062847	23.5518515967	<u>1.0165667556</u>
0.001250	864416	<u>1080.5176461071</u>	-26.1248064862	23.5518520559	<u>1.0165667426</u>

realisation of the KSEq. For $r > R$, roughly the minimum distance between the barycenters of the connected components, a large number of UPOs are identified incorrectly as duplicated, because individual clusters that represent different UPOs are connected together. For $r < \epsilon_h$, nearby solutions approximating the same orbit form distinct clusters and are thus not tagged as duplicates. Any choice of r in the interval $\epsilon_h \ll r \ll R$ will identify duplicates correctly. This shows that a sensible choice of ϵ_h is possible an order of magnitude lower than R .

Further remarks. In practice, the KD-tree based nearest neighbours search is significantly more efficient when $N_h < \log_2(M)$. This limits the number of invariant features that can be used. In spatially extended systems with many degrees of freedom, good invariant features could be chosen by projecting UPOs onto a low-dimensional linear subspace obtained, for instance, by the span of the first leading Proper Orthogonal Decomposition vectors. The data transformation will ensure maximal efficiency, because the splitting hyper-planes in the construction of the KD-tree will be orthogonal to axis of inertia of the data points.

Appendix D. Sensitivity of long UPOs. We show, empirically, that the adjoint sensitivity calculations using the methods discussed in [Appendix A](#) and [Appendix B](#) are well behaved regardless of the stability characteristics and length of the UPO. To this end we show convergence results for a rather long UPOs at $\rho = 28$, with period $T \approx 1080$ and symbol sequence length 1439, too long to be reported here. An initial guess for the search is obtained by taking a point on the attractor, after transients have decayed, integrating the Lorenz equations for a thousand time units and then finding the nearest recurrence within the next hundred time units. Such long UPOs can be easily found using the approach described in [Appendix A](#), because the instability associated to classical shooting methods is completely avoided by the particular discretisation adopted [2].

The rows of [Table 2](#) show how the same quantities reported in [Table 1](#) vary as the temporal discretisation is refined. Refinement is performed by doubling the number of grid points (M) and using a linear prolongation algorithm to interpolate the solution on the finer grid. Newton-Raphson iterations are then run to convergence with less tight tolerances due to the length of the orbit ($\alpha = 11$ as discussed in [Appendix B](#)). At each refinement iteration approximately one digit of accuracy is gained, due to the fourth order accurate scheme (52). This UPO is shown in [Figure 17](#). What is remarkable and intriguing is that the temporal dynamics of this orbit is indistinguishable from a chaotic trajectory, as they are its statistics (e.g. the density of z , in [Figure 17-\(c\)](#)). The difference, however, is that unlike for similarly long chaotic trajectories, its sensitivity with respect to parameter perturbations can be calculated exactly.




Figure 17. A long UPO of the Lorenz equations. Panel (a), projection on the (x, z) plane; panel (b), time history of $z(t)$; panel (c), normalised histogram of the z variable over 100 bins from 0 to 50, for the UPO and for a long chaotic trajectory (10000 time units).

REFERENCES

- [1] A. ARMAOU AND D. P. CHRISTOFIDES, *Feedback control of the Kuramoto–Sivashinsky equation*, *Physica D: Nonlinear Phenomena*, 137 (2000), pp. 49–61.
- [2] U. M. ASCHER, R. MATTHEI, AND R. D. RUSSELL, *Numerical solution of boundary value problems for ordinary differential equations*. Society for Industrial and Applied Mathematics, 1995.
- [3] D. AUERBACH, P. CVITANOVIĆ, J. P. ECKMANN, G. GUNARATNE, AND I. PROCACCIA, *Exploring Chaotic Motion Through Periodic-Orbits*, *Phys. Rev. Letters*, 58 (1987), pp. 2387–2389.
- [4] M. BATISTA AND A. R. A. I. KARAWIA, *The use of the Sherman–Morrison–Woodbury formula to solve cyclic block tri-diagonal and cyclic block penta-diagonal linear systems of equations*, *Applied Mathematics and Computation*, 210 (2009), pp. 558–563.
- [5] P. J. BLONIGAN AND Q. WANG, *Least squares shadowing sensitivity analysis of a modified Kuramoto–Sivashinsky equation*, *Chaos, Solitons and Fractals*, 64 (2014), pp. 16–25.
- [6] N. B. BUDANUR, K. Y. SHORT, M. FARAZMAND, A. P. WILLIS, AND P. CVITANOVIĆ, *Relative periodic orbits form the backbone of turbulent pipe flow*, *Journal of Fluid Mechanics*, 833 (2017), pp. 274–301.
- [7] G. J. CHANDLER AND R. R. KERSWELL, *Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow*, *Journal of Fluid Mechanics*, 722 (2013), pp. 554–595.
- [8] F. CHRISTIANSEN, P. CVITANOVIĆ, AND V. PUTKARADZE, *Spatiotemporal chaos in terms of unstable recurrent patterns*, *Nonlinearity*, 10 (1997), pp. 55–70.
- [9] P. CVITANOVIĆ, R. L. DAVIDCHACK, AND E. SIMINOS, *On the state space geometry of the Kuramoto–Sivashinsky flow in a periodic domain*, *Siam Journal on Applied Dynamical Systems*, 9 (2010), pp. 1–33.
- [10] S. V. ERSHOV, *Is a perturbation theory for dynamical chaos possible?*, *Physics Letters A*, 177 (1993), pp. 180–185.
- [11] G. L. EYINK, T. HAINE, AND D. J. LEA, *Ruelle’s linear response formula, ensemble adjoint schemes and Lévy flights*, *Nonlinearity*, 17 (2004), pp. 1867–1889.
- [12] F. GIANNETTI, S. CAMARRI, AND P. LUCHINI, *Structural sensitivity of the secondary instability in the wake of a circular cylinder*, *Journal of Fluid Mechanics*, 651 (2010), pp. 319–337.
- [13] J. F. GIBSON, J. HALCROW, AND P. CVITANOVIĆ, *Equilibrium and travelling-wave solutions of plane Couette flow*, *Journal of Fluid Mechanics*, 638 (2009), pp. 243–24.
- [14] S. N. GOMES, M. PRADAS, S. KALLIADASIS, D. T. PAPAGEORGIOU, AND G. A. PAVLIOTIS, *Controlling spatiotemporal chaos in active dissipative-dispersive nonlinear systems*, *Phys. Rev. E*, 92 (2015), pp. 022912–5.
- [15] W. GOVAERTS, *Solution of bordered singular systems in numerical continuation and bifurcation*, *Journal of Computational and Applied Mathematics*, 50 (1994), pp. 339–347.

- [16] J. GUCKENHEIMER AND P. HOLMES, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Volume 42 of Applied Mathematical Sciences, Springer-Verlag.
- [17] J. K. HALE, *Ordinary Differential Equations*, vol. 21, Wiley-Interscience, 1969.
- [18] A. C. HINDMARSH, P. N. BROWN, K. E. GRANT, S. L. LEE, R. SERBAN, D. E. SHUMAKER, AND C. S. WOODWARD, *SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers*, ACM Transactions on Mathematical Software (TOMS), 31 (2005), pp. 363–396.
- [19] Y. HWANG, *Structural sensitivities of soft and steep nonlinear global modes in spatially developing media*, European Journal of Mechanics B/Fluids, 49 (2014), pp. 322–334.
- [20] G. KAWAHARA AND S. KIDA, *Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst*, Journal of Fluid Mechanics, 449 (2001), pp. 291–300.
- [21] G. KAWAHARA, M. UHLMANN, AND L. VAN VEEN, *The significance of simple invariant solutions in turbulent flows*, Annu. Rev. Fluid Mech., 44 (2011), pp. 203–225.
- [22] Y. KURAMOTO, *Diffusion-induced chaos in reaction systems*, Progress of Theoretical Physics Supplement, (1978).
- [23] Y. LAN AND P. CVITANOVIĆ, *Variational method for finding periodic orbits in a general flow*, Phys. Rev. E, 69 (2004), pp. 016217–10.
- [24] Y. LAN AND P. CVITANOVIĆ, *Unstable recurrent patterns in Kuramoto-Sivashinsky dynamics*, Phys. Rev. E, 78 (2008), pp. 026208–12.
- [25] D. J. LEA, M. R. ALLEN, AND T. HAINE, *Sensitivity analysis of the climate of a chaotic system*, Tellus Series a-Dynamic Meteorology and Oceanography, 52 (2000), pp. 523–532.
- [26] E. N. LORENZ, *Deterministic nonperiodic flow*, J. Atmos. Sci., (1963).
- [27] M. NAGATA, *Three-dimensional finite-amplitude solutions in plane couette flow: Bifurcation from infinity*, Journal of Fluid Mechanics, 217 (1990), pp. 519–527.
- [28] T. SAUER, *Shadowing breakdown and large errors in dynamical simulations of physical systems*, Phys. Rev. E, 65 (2002), pp. 036220–5.
- [29] T. SAUER, C. GREBOGI, AND J. A. YORKE, *How Long Do Numerical Chaotic Solutions Remain Valid?*, Phys. Rev. Letters, 79 (1997), pp. 59–62.
- [30] T. SHINBROT, C. GREBOGI, E. OTT, AND J. A. YORKE, *Using small perturbations to control chaos*, Nature, 363 (1993), pp. 411–417.
- [31] G. I. SIVASHINSKY AND D. M. MICHELSON, *On Irregular Wavy Flow of a Liquid Film Down a Vertical Plane*, Progress of theoretical physics, 63 (1980), pp. 2112–2114.
- [32] C. SPARROW, *The Lorenz Equations*, vol. 41 of Bifurcations, Chaos, and Strange Attractors, Springer Science & Business Media, New York, NY, Dec. 2012.
- [33] J. C. SPROTT, *Chaos and Time-series Analysis*, Oxford University Press, Jan. 2003.
- [34] J. THUBURN, *Climate sensitivities via a Fokker–Planck adjoint approach*, Q. J. R. Meteorol. Soc., 131 (2005), pp. 73–92.
- [35] L. N. TREFETHEN AND J. A. C. WEIDEMAN, *The Exponentially Convergent Trapezoidal Rule*, SIAM Rev., 56 (2014), pp. 385–458.
- [36] D. VISWANATH, *The Lindstedt–Poincaré Technique as an Algorithm for Computing Periodic Orbits*, SIAM Rev., 43 (2001), pp. 478–495.
- [37] D. VISWANATH, *Symbolic dynamics and periodic orbits of the Lorenz attractor*, Nonlinearity, 16 (2003), pp. 1035–1056.
- [38] I. VYTYAZ, D. C. LEE, P. K. HANUMOLU, U. K. MOON, AND K. MAYARAM, *Sensitivity analysis for oscillators*, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 27 (2008), pp. 1521–1534.
- [39] Q. WANG, *Forward and adjoint sensitivity computation of chaotic dynamical systems*, Journal of Computational Physics, 235 (2013), pp. 1–13.
- [40] Q. WANG, *Convergence of the Least Squares Shadowing Method for Computing Derivative of Ergodic Averages*, SIAM J. Numer. Anal., 52 (2014), pp. 156–170.
- [41] Q. WANG, R. HU, AND P. BLONIGAN, *Least Squares Shadowing sensitivity analysis of chaotic limit cycle oscillations*, Journal of Computational Physics, 267 (2014), pp. 210–224.
- [42] A. K. WILKINS, B. TIDOR, J. WHITE, AND P. I. BARTON, *Sensitivity Analysis for Oscillating Dynamical Systems*, SIAM J. Sci. Comput., 31 (2009), pp. 2706–2732.
- [43] M. J. ZAHR, P. O. PERSSON, AND J. WILKENING, *A fully discrete adjoint method for optimization of*

flow problems on deforming domains with time-periodicity constraints, Computers and Fluids, 139 (2016), pp. 1–18.