

Constructing periodic orbits of high-dimensional chaotic systems by an adjoint-based variational method

Sajjad Azimi, Omid Ashtari, and Tobias M. Schneider*
*Emergent Complexity in Physical Systems Laboratory (ECPS),
École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland*
(Dated: October 14, 2021)

Chaotic dynamics in systems ranging from low-dimensional nonlinear differential equations to high-dimensional spatio-temporal systems including fluid turbulence is supported by non-chaotic, exactly recurring time-periodic solutions of the governing equations. These unstable periodic orbits capture key features of the turbulent dynamics and sufficiently large sets of orbits promise a framework to predict the statistics of the chaotic flow. Computing periodic orbits for high-dimensional spatio-temporally chaotic systems remains challenging as known methods either show poor convergence properties because they are based on time-marching of a chaotic system causing exponential error amplification; or they require constructing Jacobian matrices which is prohibitively expensive. We propose a new matrix-free method that is unaffected by exponential error amplification, is globally convergent and can be applied to high-dimensional systems. The adjoint-based variational method constructs an initial value problem in the space of closed loops such that periodic orbits are attracting fixed points for the loop-dynamics. We introduce the method for general autonomous systems. An implementation for the one-dimensional Kuramoto-Sivashinsky equation demonstrates the robust convergence of periodic orbits underlying spatio-temporal chaos. Convergence does not require accurate initial guesses and is independent of the period of the respective orbit.

Keywords: spatio-temporal chaos, unstable periodic orbits, adjoint methods, variational methods, matrix-free numerical methods, Kuramoto-Sivashinsky, dynamical systems approach to turbulence

I. INTRODUCTION

Ideas from low-dimensional chaotic dynamical systems have recently led to new insights into high-dimensional spatio-temporally chaotic systems including fluid turbulence. The idea for a dynamical description of turbulence has a long history [1–3] and stems from the observation that turbulent flows often show recognizable transient coherent patterns that recur over time and space [4]. Only in the last 15 years, however, has concrete progress allowed dynamical systems to be truly established as a new paradigm to study turbulence [5–7]. This progress is based on the discovery of unstable non-chaotic steady and time-periodic solutions of the fully nonlinear Navier-Stokes equations which leads to a description of turbulence as a walk through a connected forest of these dynamically connected invariant (‘exact’) solutions in the infinite-dimensional state space of the flow equations [8–11].

Of special importance are time-periodic exactly recurring flows. These so-called unstable periodic orbits capture the evolving dynamics of the flow [12] and form the elementary building blocks of the chaotic dynamics. Periodic orbits have been recognized as being key for understanding chaos since the 1880s [13–15]. Provided results from low-dimensional hyperbolic dissipative systems carry over to high-dimensional spatio-temporally chaotic systems, periodic orbits lie dense in the chaotic set supporting turbulence. The turbulent trajectory thus

almost always shadows a periodic orbit. As a consequence, periodic orbit theory allows to express ergodic ensemble averages of the turbulent flow as weighted sums over periodic orbits. In these ‘cycle expansions’, the statistical weight of an individual orbit is controlled by its stability features [16–21]. Sufficiently complete sets of periodic orbits for three-dimensional fluid flows may thus eventually allow to quantitatively describe statistical properties of turbulence in terms of exact invariant solutions of the underlying flow equations [22]. Even if a full description of turbulence in terms of periodic orbits remains beyond our reach, individual periodic orbits are of significant importance as they capture key physical processes underlying the turbulent dynamics and may inform control strategies [23]. Consequently, robust tools for computing periodic orbits of high-dimensional spatio-temporally chaotic systems including three-dimensional fluid flows are needed.

High-dimensional spatio-temporal systems, including spectrally discretized three-dimensional fluid flow problems, are often characterized by more than $N = 10^6$ highly coupled degrees of freedom. Computing periodic orbits of such high-dimensional strongly coupled systems remains computationally challenging. The commonly used shooting method considers an initial value problem yielding trajectories satisfying the evolution equations and varies the initial condition until the solution closes on itself. To find the initial condition u_0 and the period T , Newton iteration is used to numerically solve the nonlinear equation $g(u_0, T) = f^T(u_0) - u_0$, where f^T is the evolution of the state u_0 over time T . To solve this system of nonlinear coupled equations, a standard Newton method would require constructing the full Jacobian matrix with

* tobias.schneider@epfl.ch

$\mathcal{O}(N^2)$ elements. This is practically impossible for high-dimensional strongly coupled systems with large N . Key for computing periodic orbits of high-dimensional systems are thus *matrix-free* Newton methods that do not construct the Jacobian matrix but only require successive evaluations of the function g , implying time-stepping of the evolution equations. Commonly used algorithms are Krylov subspace methods [24, 25] including the Newton-GMRES-hook-step method by Viswanath [9, 26, 27] as well as slight variations with alternative trust-region optimizations [28, 29].

The matrix-free Newton approach is well suited for computing fixed points, where the ‘period’ T can be chosen arbitrarily, but the Newton approach poses fundamental challenges for periodic orbits. The defining property of a chaotic system is an exponential-in-time separation of trajectories which leads to a sensitive dependence on initial conditions. Very small changes in the initial condition u_0 are thus exponentially amplified by the required time-integration. Finding zeros of g thus becomes an ill-conditioned problem. Consequently, an extremely good initial guess is required for the Newton method to converge. Generating sufficiently accurate initial guesses is very challenging and often impossible. Owing to the finite numerical precision of double-precision arithmetic long and unstable orbits are even entirely impossible to converge. Examples demonstrating the difficulty in finding periodic orbits of high-dimensional systems using shooting methods include the seminal work by Chandler and Kerswell [22], who computed approximately 100 orbits for a two-dimensional model flow and describe the time-consuming and tedious manual work to find initial guesses and trying to converge them. Likewise van Veen et al. [30] recently computed a single periodic orbit for box turbulence with only moderate resolution of 64^3 grid points. The authors reach a moderately small residual of $1.8 \cdot 10^{-4}$ and thus many orders of magnitude larger than machine precision only after “several months of computing on modern GPU cards, due to the poor conditioning of the linear problems associated with Newton’s method”. Consequently, more robust methods with larger radii of convergence than those of shooting methods are needed to compute periodic orbits of high-dimensional spatio-temporally chaotic systems.

For low-dimensional systems more robust methods for finding periodic orbits have been devised. In two- and four-dimensional systems periodic orbits have been constructed using alternatives to shooting methods including variational approaches with Fourier representation in time [31] and methods based on topological degree theory [32, 33], which have not been extended to higher dimensional problems. In the context of nonlinear partial differential equations (PDEs) with one spatial dimension, Lan and Cvitanović [34] followed a different strategy to overcome challenges due to the exponential error amplification associated with the shooting method: Instead of starting from trajectories satisfying the evolution equations and varying the initial condition until the solution

closes on itself, they suggested a variational method that reverses the approach: It starts from a closed loop in state space that does not satisfy the evolution equations and then adapts the loop until it solves the equations and a periodic orbit is found. To adapt the closed loop, the problem is recast as a minimization problem in the space of all closed loops. The loop is driven towards a periodic orbit by minimizing a cost function that measures the deviation of the loop from an integral curve of the vector field induced by the governing equations. No time-marching along the orbit is required and the loop is adapted locally. Consequently, the variational method does not suffer from exponential error amplification and has a large radius of convergence. The robustness of the method has been demonstrated in the one-dimensional Kuramoto-Sivashinsky system [35] for which Lasagna [23] recently found more than 20 000 periodic orbits using $N=32$ Fourier modes to discretize the problem.

Unfortunately, the robust variational method of Lan and Cvitanović cannot be scaled to high-dimensional problems such as fluid turbulence. The method is not matrix-free but requires the explicit construction of Jacobian matrices and their inversion. Moreover, accurate computations of tangents to the loop by finite differences require the loop to be represented by a sufficiently large number of instantaneous fields closely-spaced in temporal direction along the loop. The size of the Jacobian matrix to be inverted scales with the number of instantaneous fields M and the spatial degrees of freedom N as $\mathcal{O}(M^2 N^2)$. This scaling reflects the prohibitively large memory requirements for high-dimensional systems. The only attempt to apply the method to a higher-dimensional system we are aware of is Fazendairo et al. [36, 37] who study forced box-turbulence in a triple-periodic box using Lattice-Boltzmann computations. They provide evidence for the convergence of two periodic orbits but reaching a modestly small residual of $\mathcal{O}(10^{-5})$ on a relatively small 64^3 spatial lattice requires tens of thousands of CPU cores. As stated by Fazendairo et. al., even finding the shortest orbits of 3D flows using the method by Lan and Cvitanović requires petascale computing resources. Despite its robustness, the variational method by Lan and Cvitanović is thus too computationally expensive to be realistically used for high-dimensional spatio-temporally chaotic systems.

Here we propose a novel matrix-free method that provides the same favorable convergence properties of the variational method by Lan and Cvitanović [34, 35] but can be applied to high-dimensional systems. The method combines a variational approach similar to Lan and Cvitanović with an adjoint-based minimization technique inspired by recent work of Farazmand [38] on computing steady state solutions. Combining the variational approach with adjoints allows us to construct an initial value problem in the space of closed loops such that unstable periodic orbits become attracting fixed points of the dynamics in loop-space. Converging to a periodic orbit thus only requires evolving an initial guess under

the dynamics in loop-space. We develop the matrix-free adjoint-based variational method for general autonomous dynamical systems. As a proof-of-concept, the introduced method is applied to the one-dimensional Kuramoto-Sivashinsky equation (KSE) [39, 40]. The KSE is a model system showing spatio-temporal chaos that has commonly been used as a sandbox model to develop algorithms that are eventually applied to three-dimensional fluid flows. We demonstrate the robust convergence of multiple periodic orbits of varying complexity and periods. The implementation utilizes a spectral Fourier discretization in the temporal direction to significantly reduce the prohibitively large memory requirements of the method by Lan and Cvitanović.

The structure of the paper is as follows: First, the proposed method for computing periodic orbits is introduced for a general autonomous system. Section II describes the setup of the variational problem and Sec. III discusses the adjoint-based minimization technique. In Sec. IV, we apply the adjoint-based variational method to the KSE and demonstrate the convergence of periodic orbits in this spatio-temporally chaotic system. Section V summarizes the manuscript and discusses future applications to three-dimensional fluid turbulence.

II. VARIATIONAL METHOD FOR FINDING PERIODIC ORBITS

We consider a general dynamical system for an n -dimensional real field \vec{u} defined over a spatial domain $\Omega \subset \mathbb{R}^d$ and varying in time t ,

$$\begin{aligned} \vec{u} : \Omega \times \mathbb{R} &\rightarrow \mathbb{R}^n, \\ (\vec{x}, t) &\mapsto \vec{u}(\vec{x}, t). \end{aligned}$$

The evolution of the field \vec{u} is first-order in time and governed by an autonomous partial differential equation (PDE) of the form

$$\frac{\partial \vec{u}}{\partial t} = \mathcal{N}(\vec{u}). \quad (1)$$

The nonlinear differential operator \mathcal{N} enforces boundary conditions at $\partial\Omega$, the boundaries of the spatial domain Ω . A periodic orbit is a temporally periodic solution of the governing equation,

$$f^T(\vec{u}) - \vec{u} = \vec{0}, \quad (2)$$

where $f^T(\vec{u}) = \vec{u} + \int_t^{t+T} \mathcal{N} dt'$ indicates the nonlinear evolution over the period T .

The shooting method considers solutions of the initial value problem and varies the initial condition $\vec{u}_0(\vec{x})$ until the solution closes on itself and becomes periodic. Equation (2) is thus treated as an algebraic equation for the initial condition and the period. An alternative approach is to consider already time-periodic fields and vary those until they satisfy the governing equations. Instead of

identifying an initial condition as in a shooting method, we consider the entire orbit as a solution of a boundary value problem in the $(d+1)$ -dimensional space-time domain. To ensure periodicity of the solution in time, the boundary conditions in space are augmented by periodic boundary conditions in time. The field $\vec{u}(\vec{x}, t)$ is thus defined on $\Omega \times [0, T)_{\text{periodic}}$.

The length of the domain in time T is unknown and needs to be determined as part of the solution. To convert the problem to a boundary value problem on a fixed domain, we rescale time $t \mapsto s := t/T$, where s denotes the normalized time coordinate. The rescaled field

$$\vec{\tilde{u}}(\vec{x}, s) := \vec{u}(\vec{x}, s \cdot T),$$

is defined on a fixed domain

$$\begin{aligned} \vec{\tilde{u}} : \Omega \times [0, 1)_{\text{periodic}} &\rightarrow \mathbb{R}^n, \\ (\vec{x}, s) &\mapsto \vec{\tilde{u}}(\vec{x}, s). \end{aligned}$$

A periodic orbit is characterized by the space-time field $\vec{\tilde{u}}(\vec{x}, s)$ and the period T satisfying

$$-\frac{1}{T} \frac{\partial \vec{\tilde{u}}}{\partial s} + \mathcal{N}(\vec{\tilde{u}}) = \vec{0}. \quad (3)$$

Boundary conditions in space remain unchanged with respect to the dynamical system (1) and are complemented by periodic boundary conditions in the temporal direction s . To simplify the notation, the overhead tilde is omitted in the remainder of the article.

A periodic orbit is defined by the combination of a field $\vec{u}(\vec{x}, s)$ and a period T that together satisfy the boundary value problem (3). Geometrically the periodic orbit is a closed trajectory in state space. To characterize general closed curves in state space, we define a *loop* $\mathbf{l}(\vec{x}, s)$ as a tuple of a field $\vec{u}(\vec{x}, s)$ and a period T . A loop does not necessarily satisfy the PDE of the boundary value problem (3) but shares all boundary conditions in space and time with periodic orbits. We denote the space of all loops by

$$\mathcal{P} = \left\{ \mathbf{l}(\vec{x}, s) = \left[\begin{array}{c} \vec{u}(\vec{x}, s) \\ T \end{array} \right] \mid \begin{array}{l} \vec{u} : \Omega \times [0, 1)_{\text{periodic}} \rightarrow \mathbb{R}^n, \ T \in \mathbb{R}^+ \\ \vec{u} \text{ satisfies BC at } \partial\Omega \text{ and is periodic in } s \end{array} \right\}. \quad (4)$$

Periodic orbits are specific elements of the loop-space \mathcal{P} that satisfy the PDE (3). A general loop only satisfies the boundary conditions but not the PDE.

The idea of the variational method is to consider an initial loop $\mathbf{l}_0(\vec{x}, s) \in \mathcal{P}$ and to evolve the loop until it satisfies the boundary value problem (3). The loop thereby converges to a periodic orbit. To evolve a loop towards a periodic orbit we minimize the cost function J measuring the deviation of a loop from a solution of the

boundary value problem,

$$J : \mathcal{P} \rightarrow \mathbb{R}^+, \quad (5)$$

$$\mathbf{1} \mapsto J(\mathbf{1}) := \int_0^1 \int_{\Omega} \vec{r} \cdot \vec{r} d\vec{x} ds.$$

where \vec{r} is the residual of Eq. (3):

$$\vec{r} = -\frac{1}{T} \frac{\partial \vec{u}}{\partial s} + \mathcal{N}(\vec{u}). \quad (6)$$

The cost function J penalizes a nonzero residual \vec{r} . For a periodic orbit J is zero otherwise it takes positive values. Thus, absolute minima of J correspond to periodic orbits. The problem of finding periodic orbits has thereby been converted into an optimization over loop-space \mathcal{P} . Instead of fixing the temporal phase of loops by a phase condition, we intentionally define J such that it is invariant under a reparametrization $s \mapsto s' = (s + \sigma) \bmod 1$ corresponding to a phase shift by σ in the temporal periodic direction. Consequently, every periodic orbit is represented by an entire family of equivalent phase-shifted representations. Because the variational method does not have to adapt the temporal phase of a loop to satisfy a phase condition, we expect the arbitrariness in the temporal phase to allow the variational method to converge to periodic orbits more robustly.

Geometrically, minimizing the cost function corresponds to deforming a closed curve, a loop, in the system's state space, the space spanned by all instantaneous fields $\vec{u}(\vec{x})$ satisfying the boundary conditions, until the loop becomes an integral curve of the vector field $\mathcal{N}(\vec{u})$ induced by the dynamical system. The loop thereby becomes a solution of the PDE and represents a periodic orbit. This is schematically shown in Fig. 1. At each point \vec{u} along the loop, the vector field defines the flow direction $\mathcal{N}(\vec{u})$ while $\partial \vec{u} / \partial t = T^{-1} \partial \vec{u} / \partial s$ is the tangent vector along the loop (see panel a). The cost function J measures the misalignment between the vector field and the loop's tangent vectors integrated along the entire loop. Consequently, minimizing J towards its absolute minimum $J = 0$ deforms the loop until the tangent vectors everywhere match the flow and the loop becomes an integral curve of the vector field, as exemplified in panel b. The loop is locally deformed to align with the vector field and no time-marching causing exponential instabilities is required.

III. ADJOINT-BASED METHOD FOR MINIMIZING THE COST FUNCTION J

We recast the problem of finding periodic orbits as a minimization problem in the space of all loops. Absolute minima of the cost function J with value $J = 0$ correspond to periodic orbits. To minimize J without constructing Jacobians we develop an adjoint-based approach inspired by the recently introduced method

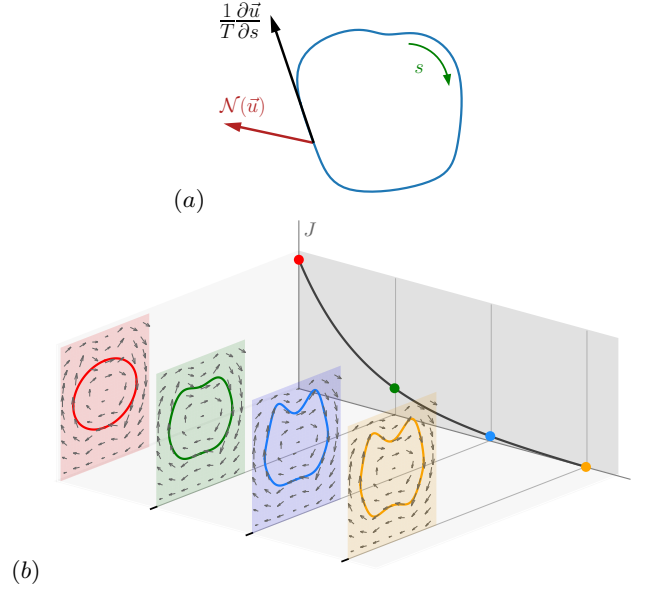


FIG. 1. Schematic of the variational method for finding periodic orbits. (a) An arbitrary closed loop (blue line) parametrized by $s \in [0, 1]$ does not satisfy the governing equations as its loop tangent $\partial \vec{u} / \partial t = T^{-1} \partial \vec{u} / \partial s$ is misaligned relative to the vector field $\mathcal{N}(\vec{u})$ induced by the dynamical system. (b) Minimizing a cost function J measuring the misalignment between the vector field and the loop tangent deforms the loop. When the global minimum of the cost function with $J = 0$ is reached the tangent vectors everywhere match the flow, $\partial \vec{u} / \partial t = \mathcal{N}(\vec{u})$. The loop becomes an integral curve of the vector field and a periodic orbit is identified.

by Farazmand [38] who computes equilibria of a two-dimensional flow. We construct an initial value problem in loop-space \mathcal{P} whose dynamics monotonically decreases the cost function J until a minimum of J is reached.

To derive an appropriate variational dynamics in loop-space, we define the space of generalized loops:

$$\mathcal{P}_g = \left\{ \mathbf{q}(\vec{x}, s) = \begin{bmatrix} \vec{q}_1(\vec{x}, s) \\ q_2 \end{bmatrix} \mid \vec{q}_1 : \Omega \times [0, 1]_{\text{periodic}} \rightarrow \mathbb{R}^n, q_2 \in \mathbb{R}, \vec{q}_1 \text{ is periodic in } s \right\}. \quad (7)$$

Elements $\mathbf{q} \in \mathcal{P}_g$ do not necessarily satisfy the spatial boundary condition of periodic orbits at $\partial \Omega$ and are thus termed generalized loops. Obviously, the space of loops \mathcal{P} is a subset of the space of generalized loops $\mathcal{P} \subset \mathcal{P}_g$. For a loop, the components of the generalized loop have specific meaning, $\vec{q}_1 = \vec{u}$ and $q_2 = T$. Throughout this paper, generalized loops are denoted by boldface letters. The space of generalized loops \mathcal{P}_g carries a real-valued

inner product

$$\langle \cdot, \cdot \rangle : \mathcal{P}_g \times \mathcal{P}_g \rightarrow \mathbb{R},$$

$$\langle \mathbf{q}, \mathbf{q}' \rangle = \left\langle \begin{bmatrix} \vec{q}_1 \\ q_2 \end{bmatrix}, \begin{bmatrix} \vec{q}'_1 \\ q'_2 \end{bmatrix} \right\rangle = \int_0^1 \int_{\Omega} \vec{q}_1 \cdot \vec{q}'_1 d\vec{x} ds + q_2 q'_2, \quad (8)$$

and an L_2 -norm

$$\|\mathbf{q}\| = \sqrt{\langle \mathbf{q}, \mathbf{q} \rangle} = \sqrt{\int_0^1 \int_{\Omega} \vec{q}_1 \cdot \vec{q}_1 d\vec{x} ds + q_2^2}. \quad (9)$$

The objective is to construct a dynamical system in the space of loops \mathcal{P} such that along its solutions the cost function J monotonically decreases and periodic orbits become attracting fixed points of the dynamical system. We parametrize the evolution of loops in \mathcal{P} by a fictitious time τ : $\mathbf{l}(\tau) = [\vec{u}(\vec{x}, s; \tau); T(\tau)]$ and define an evolution equation,

$$\frac{\partial \mathbf{l}}{\partial \tau} = \mathbf{G}(\mathbf{l}) \quad (10)$$

with operator \mathbf{G} chosen such that

$$\frac{\partial J}{\partial \tau} \leq 0 \quad \forall \tau. \quad (11)$$

To construct the operator \mathbf{G} , we follow analogous arguments to those in [38]. Note, however, that in [38] a dynamical system is derived that acts on the space of instantaneous fields to find equilibria while here we determine an operator \mathbf{G} that evolves loops as defined in (4). The rate of change of J along solutions of Eq. (10) is (see Appendix A for details)

$$\frac{\partial J}{\partial \tau} = 2 \langle \mathcal{L}(\mathbf{l}; \mathbf{G}), \mathbf{R} \rangle. \quad (12)$$

where $\mathbf{R} \in \mathcal{P}_g$ is a generalized loop

$$\mathbf{R}(\mathbf{l}) = \begin{bmatrix} \vec{r} \\ 0 \end{bmatrix}, \quad (13)$$

with $\vec{r}(\mathbf{l})$ the residual field (6). $\mathcal{L}(\mathbf{l}; \mathbf{G})$ is the directional derivative of the residual \mathbf{R} in the direction \mathbf{G} , evaluated for the current loop \mathbf{l} :

$$\mathcal{L}(\mathbf{l}; \mathbf{G}) = \lim_{\epsilon \rightarrow 0} \frac{\mathbf{R}(\mathbf{l} + \epsilon \mathbf{G}) - \mathbf{R}(\mathbf{l})}{\epsilon}. \quad (14)$$

Using the adjoint of the directional derivative, we express Eq. (12) as

$$\frac{\partial J}{\partial \tau} = 2 \langle \mathbf{G}, \mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}) \rangle \quad (15)$$

where \mathcal{L}^\dagger is the adjoint operator of \mathcal{L} with

$$\langle \mathcal{L}(\mathbf{q}; \mathbf{q}'), \mathbf{q}'' \rangle = \langle \mathbf{q}', \mathcal{L}^\dagger(\mathbf{q}; \mathbf{q}'') \rangle, \quad (16)$$

for all generalized loops \mathbf{q} , \mathbf{q}' and \mathbf{q}'' . This form allows to enforce the monotonic decrease of the cost function J by explicitly choosing the operator \mathbf{G} :

$$\mathbf{G} = -\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}). \quad (17)$$

With this choice for \mathbf{G} , the cost function evolves as

$$\frac{\partial J}{\partial \tau} = 2 \langle -\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}), \mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}) \rangle = -2 \|\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R})\|^2 \leq 0. \quad (18)$$

Thus, along solutions of $\partial \mathbf{l} / \partial \tau = \mathbf{G}(\mathbf{l}) = -\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R})$ the cost function J is guaranteed to monotonically decrease.

To find a periodic orbit using the adjoint approach, an initial loop is advanced under the dynamical system in loop-space, until a minimum of the cost function, corresponding to an attracting fixed point with $\partial_\tau \mathbf{l} = \mathbf{0}$, is reached. If an absolute minimum, $J = 0$, is reached, the loop satisfies the boundary value problem (3) and represents a periodic orbit. The phase of the minimizing loop is not chosen by the adjoint-based variational method but depends on the initial condition.

IV. APPLICATION TO KURAMOTO-SIVASHINSKY EQUATION

We demonstrate the adjoint-based variational method for the one-dimensional Kuramoto-Sivashinsky equation (KSE) [39, 40]. This nonlinear partial differential equation for a one-dimensional field $u(x, t)$ on a 1D periodic interval $x \in [0, L] = \Omega$ reads

$$\frac{\partial u}{\partial t} = -u \frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \nu \frac{\partial^4 u}{\partial x^4}; \quad x \in [0, L]_{\text{periodic}}, \quad t \in \mathbb{R} \quad (19)$$

with a constant 'superviscosity' $\nu > 0$. The KSE has the general form of Eq. (1) with $n = d = 1$. We denote the scalar spatial coordinate by x . Rescaling the field u by the inverse of L indicates that the only control parameter is $\mathbb{L} = L/\sqrt{\nu}$ the ratio of the domain length and the square-root of the constant ν . Consequently, fixing the domain length L and varying ν is equivalent to fixing ν and treating L as a control parameter. Both scalings are used in literature. Here, we fix $\nu = 1$ and consider L as the control parameter. The equivariance group of the KSE contains continuous shifts in x and the discrete center symmetry,

$$x \rightarrow -x; \quad u \rightarrow -u. \quad (20)$$

We discuss periodic orbits both in the full unconstrained space and in the subspace of fields invariant under the discrete center symmetry.

The trivial solution of the KSE, $u = \text{const}$, is linearly unstable for $L > 2\pi\sqrt{\nu}$ [41]. A series of bifurcations leads to increasingly complex dynamics when L is increased. We consider the parameter value $L = 39$ where the KSE shows spatio-temporally chaotic dynamics reminiscent of turbulence [42].

A. Formulation of the adjoint-based method for the KSE

For the 1D-KSE a loop consists of a one-dimensional field $u(x, s)$ defined over $[0, L) \times [0, 1)$ and the period T . The residual of the boundary value problem for a periodic orbit (6), expressed as generalized loop \mathbf{R} (see Eq. (13)), is

$$\mathbf{R}(\mathbf{l}) = \begin{bmatrix} r(\mathbf{l}) \\ 0 \end{bmatrix} = \begin{bmatrix} -\frac{1}{T} \frac{\partial u}{\partial s} - u \frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4} \\ 0 \end{bmatrix}, \quad (21)$$

where vector notation has been suppressed because the dimension of the field is $n = 1$.

The dynamical system in loop-space for which the cost function monotonically decreases and periodic orbits become attracting fixed points is based on the adjoint operator of the directional derivative of \mathbf{R} . Partial integration directly yields the adjoint operator for the KSE problem (see Appendix B),

$$\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}) = \begin{bmatrix} \frac{1}{T} \frac{\partial r}{\partial s} + u \frac{\partial r}{\partial x} - \frac{\partial^2 r}{\partial x^2} - \frac{\partial^4 r}{\partial x^4} \\ \int_0^1 \int_0^L \frac{1}{T^2} \frac{\partial u}{\partial s} r dx ds \end{bmatrix}. \quad (22)$$

Consequently, the dynamical system in loop-space $\partial \mathbf{l} / \partial \tau = -\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R})$ (see (17)) minimizing the cost function J is

$$\frac{\partial \mathbf{l}}{\partial \tau} = \begin{bmatrix} \frac{\partial u}{\partial \tau} \\ \frac{\partial T}{\partial \tau} \end{bmatrix} = \begin{bmatrix} -\frac{1}{T} \frac{\partial r}{\partial s} - u \frac{\partial r}{\partial x} + \frac{\partial^2 r}{\partial x^2} + \frac{\partial^4 r}{\partial x^4} \\ -\int_0^1 \int_0^L \frac{1}{T^2} \frac{\partial u}{\partial s} r dx ds \end{bmatrix}. \quad (23)$$

The first component of Eq. (23) prescribes the deformation of the field $u(x, s)$, while the second component updates the period T .

The dynamical system in loop-space formulated for the KSE, Eq. (23), is equivariant with respect to the discrete symmetry:

$$\Xi : (x, s) \rightarrow (-x, s); \quad \begin{bmatrix} u \\ T \end{bmatrix} \rightarrow \begin{bmatrix} -u \\ T \end{bmatrix}. \quad (24)$$

If an initial loop is invariant under the action of Ξ , the evolution in τ will preserve the symmetry. Since the transformation of the instantaneous field $x \rightarrow -x; u(\cdot, s) \rightarrow -u(\cdot, s)$ for all $s \in [0, 1)$ corresponds to the center-symmetry (20) of the KSE equation, the dynamical system in loop-space also preserves the center symmetry of the KSE. An initial loop with field component within the center-symmetric subspace of KSE is invariant under Ξ , which is preserved under τ -evolution. Consequently, the adjoint-based variational method preserves the discrete center-symmetry of the KSE.

B. Numerical implementation

Expressing the field component of the dynamical system (23) in terms of u using Eq. (21) yields,

$$\frac{\partial u}{\partial \tau} = G_{1,L} + G_{1,NL}, \quad (25)$$

where the linear and nonlinear terms have the form,

$$\begin{aligned} G_{1,L} &= \frac{1}{T^2} \frac{\partial^2 u}{\partial s^2} - \frac{\partial^8 u}{\partial x^8} - 2 \frac{\partial^6 u}{\partial x^6} - \frac{\partial^4 u}{\partial x^4} \\ G_{1,NL} &= -5 \frac{\partial^4 u}{\partial x^4} \frac{\partial u}{\partial x} - 10 \frac{\partial^3 u}{\partial x^3} \frac{\partial^2 u}{\partial x^2} - 3 \frac{\partial^2 u}{\partial x^2} \frac{\partial u}{\partial x} + u^2 \frac{\partial^2 u}{\partial x^2} \\ &\quad + u \left(\frac{\partial u}{\partial x} \right)^2 + \frac{2u}{T} \frac{\partial^2 u}{\partial x \partial s} + \frac{1}{T} \frac{\partial u}{\partial x} \frac{\partial u}{\partial s}. \end{aligned}$$

The field $u(x, s)$ is defined on a doubly-periodic space-time domain. We thus numerically solve the evolution equation with a pseudospectral method [43] using a Fourier discretization in both space and time. The spectral representation with M modes in space and N modes along the temporal direction is,

$$u(x_m, s_n) = \sum_{j=-\frac{M}{2}}^{\frac{M}{2}-1} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \hat{u}_{j,k} \exp \left\{ 2\pi i \left(\frac{mj}{M} + \frac{nk}{N} \right) \right\}. \quad (26)$$

In physical space, the field is represented by grid values at the Gauss-Lobatto collocation points $\{u(x_m, s_n)\}$ with $(x_m, s_n) = (mL/M, n/N)$ and index ranges $0 \leq m \leq M-1$ and $0 \leq n \leq N-1$. In spectral space, the set of discrete Fourier coefficients $\{\hat{u}_{j,k}\}$ with $-M/2 \leq j \leq M/2-1$ and $-N/2 \leq k \leq N/2-1$ represents the field. In spectral space, the evolution equation (25) for each Fourier coefficient of the field takes the form

$$\begin{aligned} \frac{\partial \hat{u}_{j,k}}{\partial \tau} &= \left[-\left(\frac{2\pi k}{T} \right)^2 - \left(\frac{2\pi j}{L} \right)^8 \right. \\ &\quad \left. + 2 \left(\frac{2\pi j}{L} \right)^6 - \left(\frac{2\pi j}{L} \right)^4 \right] \hat{u}_{j,k} + (\hat{G}_{1,NL})_{j,k}, \quad (27) \end{aligned}$$

where the discrete Fourier transform is indicated by a hat. To evaluate the nonlinear term $\hat{G}_{1,NL}$ derivatives are calculated in spectral space and transformed to physical space, where products are pointwise operations. Transforming the result back to spectral space yields the required terms. In both the spatial and temporal direction dealiasing following the 2/3 rule [43] is applied. To advance the evolution equation (25) in the fictitious time τ we implement a semi-implicit time-stepping method. An implicit-explicit Euler method treats the linear terms implicitly and the nonlinear terms $\hat{G}_{1,NL}$ are discretized explicitly.

The second component of the evolution equation (23) evolves the period of the loop T . We use an explicit

Euler method for time-stepping. The integral defining the right-hand-side is evaluated analogous to the pseudo-spectral treatment of the nonlinear terms in the evolution equation of the field. The integrand is evaluated in physical space followed by transformation to spectral space, where the integral is given by the (0,0) Fourier mode multiplied by L .

Since the purpose of defining the initial value problem in loop-space is to identify attractors corresponding to solutions of the boundary value problem for periodic orbits, stability and simplicity of the implementation are more important than accuracy when choosing a time-stepping scheme. The simple Euler method is only first order accurate in τ but remains stable for the chosen fixed time step $\Delta\tau = 0.15$.

C. Initial guesses and convergence to periodic orbits

The adjoint-based variational method advances some initial loop under the dynamical system that minimizes the cost function J . If a minimum with $J = 0$ is reached the loop satisfies the boundary value problem for a periodic orbit. Initial guesses for the procedure are extracted from chaotic solutions of the KSE (19) $u(x, t)$. The common approach for generating guesses used in conjunction with Newton-GMRES-based shooting methods extracts close recurrences measured in terms of the L_2 -distance from minima of the recurrence map $c(t, T) = \|u(\cdot, t+T) - u(\cdot, t)\|$ [16]. Here, the L_2 -norm is given by

$$\|u\|(t) = \sqrt{\int_0^L u(x, t)^2 dx}. \quad (28)$$

Exploiting the large radius of convergence of the variational method, we here choose a much simpler and computationally significantly cheaper method. Initial guesses are extracted from close recurrences in a one-dimensional projection of the solution. Specifically, we consider subsequent maxima in the time series of $\|u\|(t)$ where $\|u\|(t+T) \approx \|u\|(t)$. The segment of the solution between those subsequent maxima yields the field component of the initial loop. To ensure a smooth closed loop with field component satisfying periodic boundary conditions in the temporal direction, the solution segment is Fourier-transformed in time and high-frequency components are filtered out [34]. The double-periodic field $u_0(x, s)$ complemented by the period defines an initial guess $\mathbf{l}_0 = [u_0(x, s); T]$.

The initial guess \mathbf{l}_0 is evolved under the dynamical system in loop-space (10). Along the evolution the cost function J is guaranteed to monotonically decrease and reach a minimum. Consequently, the adjoint-based variational method is globally convergent. However, it is not guaranteed that an absolute minimum with $J = 0$ is reached but the dynamics may asymptote towards a local minimum with $J > 0$. If a global minimum is reached, a

periodic orbit satisfying the boundary value problem (3) is found. We consider a periodic orbit converged, when $\sqrt{J} < 10^{-12}$ is achieved. The periodic orbit corresponds to an attracting fixed point of the dynamical system in loop-space so that we expect exponential convergence at a rate controlled by the leading eigenvalue of the loop dynamics linearized around the attracting fixed point.

D. Results and discussion

We demonstrate the adjoint-based variational method to construct periodic orbits of the KSE for the parameter value $L = 39$. At this value, the dynamics is chaotic and a large number of unstable periodic orbits are known to exist [23]. Periodic orbits of the KSE are found by evolving initial loops under the dynamical system in loop-space (23). The pseudo-spectral method uses 64×64 Fourier modes in spatial and temporal directions to discretize the field $u(x, s)$. In spatial direction, we double the number of modes compared to [23] where the chaotic attractor of the KSE at the same parameter value $L = 39$ is shown to remain structurally intact with only 32 modes. At the increased spatial resolution of 64 modes, converged periodic orbits are checked to be independent of the number of spatial modes. In the temporal direction, the length of periodic orbits is not known a priori. Instead of adapting the resolution to the period T of the specific orbit, we chose a fixed resolution of 64 modes. Short periodic orbits of the temporally discretized system are found to be independent of the number of temporal modes and thus accurately represent periodic orbits of the continuous system. Due to increasing temporal complexity of periodic orbits with larger T , fully converged periodic orbits of the temporally discretized system represent periodic orbits of the continuous PDE less accurately. However, even for the longest orbits, increasing the temporal resolution showed that solutions of the discretized system faithfully capture the structure of PDE solutions. We consequently focus on the accuracy of convergence to periodic orbits of the temporally discretized system with 64 temporal modes, while noting that extra resolution checks are needed to determine how accurately the longer converged orbits of the discretized system represent solutions of temporally continuous PDE. A fixed time step of $\Delta\tau = 0.15$ leads to stable time-stepping.

Periodic orbits of the KSE are attracting solutions of an initial value problem in the space of loops \mathcal{P} that monotonically decreases the cost function J , as shown in Fig. 2. In the top panel, the square root of the cost function, \sqrt{J} , as a function of the fictitious time τ is shown. After $\tau \approx 1.5 \cdot 10^6$ the convergence criterion $\sqrt{J} \leq 10^{-12}$ is reached. Since the cost function J is the average of $\int_{\Omega} r^2 dx$ over s , the square root of J scales with the L_2 -norm (28) of the residual field r and should be used as the convergence criterion. Along the evolution of the loop with τ the cost function J monotonically decreases. After an initial fast decrease, \sqrt{J} decays exponentially with τ .

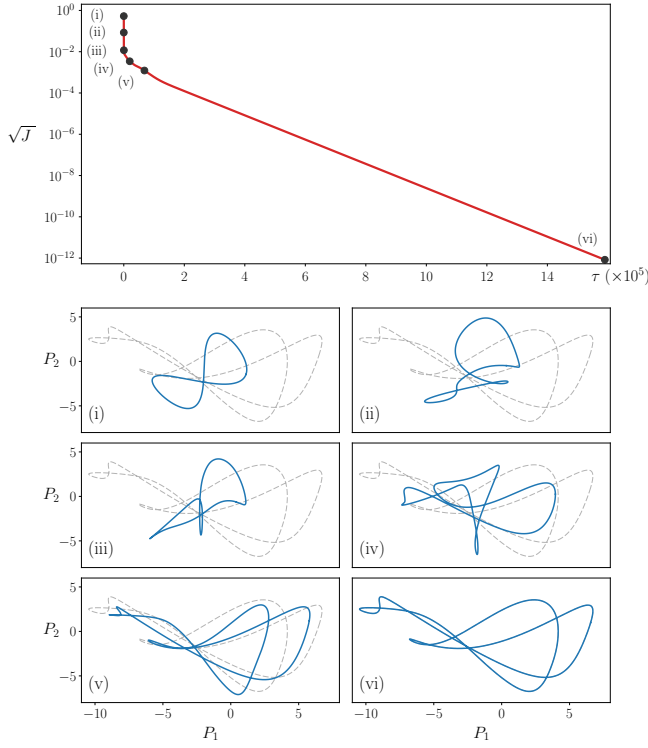


FIG. 2. Convergence of the adjoint-based variational method for finding periodic orbits of the KSE: The initial value problem in loop-space evolves loops such that the cost function J decreases monotonically along the fictitious time τ (top). The exponential decay of J towards zero indicates convergence towards a periodic orbit satisfying $J = 0$. Geometrically, the variational dynamics deforms a closed loop until it becomes an integral curve of the flow and thus a periodic orbit of the KSE. This is shown in the bottom panel, where the evolution of the loop is visualized in a two-dimensional projection of state space. Blue solid lines indicate the evolving loop at times indicated in the top panel. The dashed gray line is the converged periodic orbit. The state space projections $P_1(s)$ and $P_2(s)$ are defined by the imaginary parts of the first and second spatial Fourier coefficients of the field $u(x)$.

Geometrically, the dynamical system in loop-space (23) continuously deforms the initial loop until the loop satisfies the KSE and thereby becomes a periodic orbit. The deformation is visualized in the bottom panel, where the evolution of the loop is shown in a two dimensional projection of the state space. A very substantial deformation of the loop is associated with the fast decrease of J within the initial 10% of the integration time.

In addition to the two-dimensional field defined over the fixed space-time domain $[0, L) \times [0, 1)$, the corresponding period T is required to define a loop. Evolving a loop towards a periodic orbit implies finding the period T , which re-scales the temporal length of the space-time domain $s \rightarrow t = T \cdot s$ and thereby determines the length of extension of the domain in the direction of time t . Figure 3 shows the convergence of T to the period of the periodic orbit together with the space-time contours of




FIG. 3. A periodic orbit is characterized by the combination of the field $u(x, s)$ on a fixed double-periodic space-time domain and the time period T that rescales the temporal direction $s \rightarrow t = T \cdot s$. The variational dynamics adapts T until the period of the periodic orbit is determined (top). Finding the period T corresponds to determining the length of the domain in time t . This is evidenced by space-time contours of the solution $u(x, t = T \cdot s)$ for the initial condition (b) and the converged periodic orbit (c). The period of the initial loop and the periodic orbit are $T = 40$ and $T = 59.59$, respectively.

the corresponding initial loop $u_0(x, t = T \cdot s)$ and the converged periodic orbit $u(x, t = T \cdot s)$. As for the geometry of the loop (Fig. 2) substantial changes in the period T under the adjoint-based variational dynamics are mostly observed within the initial 10% of the integration of the dynamical system in loop-space (23). Already at $\tau = 2 \cdot 10^5$, T is very close to the period of the periodic orbit $T = 59.59$. We omit data beyond $\tau = 4 \cdot 10^5$ from Fig. 3 since changes would not be visible.

The fast initial decrease of the cost function J followed by a slow exponential decay towards zero suggests that the loop approaches the periodic orbit along the leading eigendirection of the loop dynamics linearized around the attracting fixed point. This is evidenced in Fig. 4 where the trajectory of the loop evolving towards the attracting fixed point in loop-space and the leading eigendirection are visualized in a two-dimensional projection. The residual \sqrt{J} along the trajectory at selected points are indicated on the figure. Most of the computational efforts are spent on following the exponential decay until the cost function has reached sufficiently low values, although this part of the dynamics is, at least approximately, linear. Consequently, the convergence of the method can be accelerated by explicitly exploiting the linearized dynamics in the vicinity of the attracting fixed point. A straightforward method reducing the computational costs by approximately 50% is discussed in Appendix C. More sophisticated optimizations can be implemented and will be helpful when applying the adjoint-based variational




FIG. 4. Convergence of the adjoint-based variational method to the periodic orbit shown in Figs. 2 and 3 visualized by a two-dimensional projection of the loop-space. The trajectory of the loop (blue line) evolves towards the attracting fixed point corresponding to the periodic orbit (green circle). The approach follows the leading eigendirection (orange line) of the loop dynamics linearized around the fixed point (see inset). Consequently, the residual \sqrt{J} exponentially decrease with fictitious time τ , as shown in Fig. 2 (top). The two-dimensional projection is spanned by the imaginary parts of the Fourier modes (0, 1) and (1, 0) of the spatio-temporal field $u(x, s)$. Selected residual values \sqrt{J} along the trajectory are indicated.

method to three-dimensional fluid flows.

One major advantage of the adjoint-based variational method is that the successful convergence towards a periodic orbit is independent of the period of the respective orbit. This is in contrast to shooting methods, where the exponential amplification of errors during time-marching along the orbit can hinder computing long orbits. We demonstrate the convergence of orbits of increasing period and complexity in Fig. 5. Six converged periodic orbits with periods ranging from $T = 25.37$ to $T = 147.42$ are shown in terms of state-space projections, together with initial loops extracted from a chaotic time-series of the KSE. The apparent large difference between initial loop and converged orbit demonstrates that the adjoint-based variational method offers a very large radius of convergence and convergence therefore does not depend on an initial condition in the close vicinity of the converged orbit. The evolution of loops under the dynamical system in loop-space converges to minima of the cost function J for any initial condition. While globally convergent, the variational method is not guaranteed to converge to absolute minima of J with $J = 0$, corresponding to periodic orbits, but the dynamics may approach a local minimum with $J > 0$. For initial loops extracted from recurrences in a one-dimensional projection of state space, as discussed in IV C, we observe approximately 70% of all initial conditions to converge to periodic orbits with $J = 0$. An example of a loop approaching a local minimum of J is shown in Appendix D.

Following Lasagna [23], initial loops for the six orbits discussed in Fig. 5 are extracted from a chaotic trajectory of the KSE in the subspace of center-symmetric

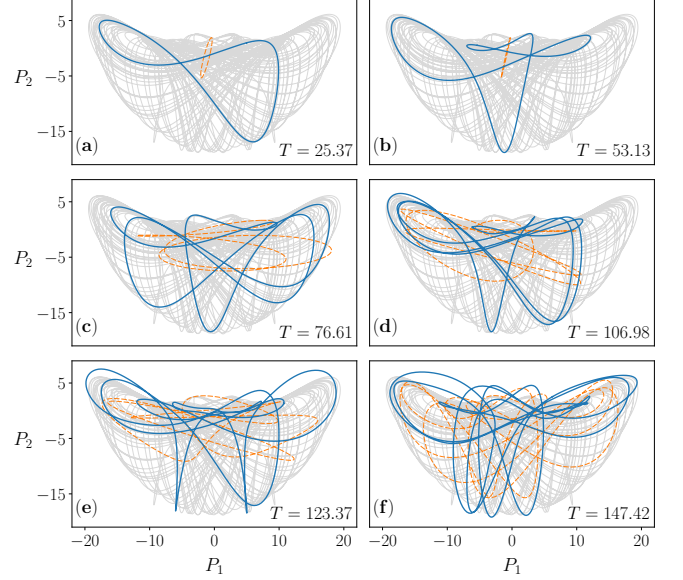


FIG. 5. Periodic orbits of increasing length and complexity converged by the adjoint-based variational method. The two-dimensional projection of state space as in Fig. 2 indicates, the initial loops (dashed orange lines) as well as the converged periodic orbits (solid blue lines). The period of the converged orbits are given in each panel. The gray line in the background of each panel is the trajectory of a long chaotic solution in the center symmetry subspace of the KSE (20). All initial loops are chosen from the center-symmetric subspace. The dynamical system in loop-space preserves the discrete symmetry of the initial loops Ξ so that all converged periodic orbits are also center-symmetric although the symmetry has not been imposed by the method.




FIG. 6. Space-time contours of the converged periodic orbits from Fig. 5 with time periods of (a) $T = 25.37$, (b) $T = 53.13$, (c) $T = 76.61$, (d) $T = 106.98$, (e) $T = 123.37$, and (f) $T = 147.42$. Unlike shooting methods, where exponential error amplification during time-integration along the orbit renders long orbits inaccessible, the adjoint-based variational method deforms orbits locally and thus converges independent of the orbit period.

fields. All initial conditions for the initial value problem in loop-space are therefore center symmetric. The dynamical system in loop-space (23) preserves the symmetry Ξ of loops (24) that corresponds to the center symmetry of instantaneous fields in the KSE system (20). Consequently, all converged periodic orbits also lie in the center symmetry subspace, as confirmed by Fig. 6, where space-time contours of the six periodic orbits are shown. Note that the method does not explicitly enforce the discrete symmetry but preserves the symmetry of the initial condition.

E. Comparing with pre-existing methods

Both the novel adjoint-based and Lan & Cvitanović's existing Newton-based variational methods relax the stringent requirement of the shooting method for the accuracy of initial guesses. However, both variational methods differ significantly in terms of computational costs and memory requirements. While the adjoint-based variational method is matrix-free, the construction of large Jacobian matrices in the Newton-based method renders application of this method to high-dimensional systems practically impossible. In this section, we thus first compare the adjoint-based variational method and the shooting method in terms of their radii of convergence. Then, the matrix-free adjoint-based variational method is compared to the Newton-based variational approach in terms of computational costs and memory requirements.

1. Implementation of existing methods

We implement both the shooting method and the Newton-based variational method for finding periodic orbits of the KSE. The implementation of the Newton-based variational method is analogous to that reported in [34, 35]. Technically we implement the method described in [23] directly in MATLAB.

The shooting method combines a time integrator for the KSE with a Newton-solver for the nonlinear equation $0 = g(u_0, T) = f^T(u_0) - u_0$. For time integrations we implement a pseudo-spectral scheme using Fourier modes in space. A first-order semi-implicit Euler's scheme advances fields in time. The time step is chosen to be close to 0.01 to ensure numerical stability of the integration. With the choice $\Delta t = T/[T/0.01]$, with $[\cdot]$ indicating the integer part of a number, the integration time T is guaranteed to be an integer multiple of Δt . The time integrator allows evaluating $g(u_0, T)$. Its zeroes defining periodic orbits are found using the trust-region-dogleg algorithm [44, 45] provided in MATLAB. When initialized by a guess sufficiently close to a periodic orbit, the method typically requires less than 20 iterations for convergence. If the method does not converge in 100 iterations, we consider the attempt failed.

2. Convergence range compared to the shooting method

The adjoint-based variational method overcomes the exponential error amplification associated with the shooting method, and is thus expected to have a significantly larger radius of convergence. Formally quantifying convergence radii is challenging in general. In addition, here there is no common metric to measure radii with respect to because the mathematical objects characterizing guesses differ fundamentally – an initial condition for the shooting method versus a loop object for the adjoint-based variational method.

To confirm the expected superior convergence range of the suggested variational method, we thus carry out a numerical experiment. We start from an initial condition that converges to a periodic orbit by using the adjoint-based variational method, extract intermediate loops along the convergence path, and test how close to the converged periodic orbit the intermediate state needs to be for the shooting method to also converge. We specifically consider the convergence path of the non-symmetric periodic orbit, discussed previously and shown in Figs. 2, 3 and 4. For different values of the cost function J , we generate initial guesses for the shooting method from instantaneous sections of the loop augmented with the period T . Since the choice of the section is arbitrary, we extract eight different temporal sections equally-spaced in time along the loop and feed all of them to the shooting method.

As expected, for larger values of J the shooting method fails to converge. Only when the cost function is sufficiently small, $\sqrt{J} < 10^{-4}$, and thus the loop is sufficiently close to the periodic orbit, the shooting method initiated with one of the eight temporal sections of the loop converges to the periodic orbit. Note that for $\sqrt{J} < 10^{-4}$, the dynamics of the adjoint-based variational method is already following the leading eigendirection of the periodic orbit, as evidenced by the exponential decrease of the cost function. Moreover, projections shown in Fig. 4 indicate how close to the converged periodic orbit the guess already needs to be for the shooting method to converge. This clearly indicates the vastly larger radius of convergence of the adjoint-based variational method compared to the shooting method.

3. Computational costs and memory requirements compared to the Newton-based variational approach

While both variational methods adapt loops until they satisfy the evolution equation, they fundamentally differ in how the loop is updated, which implies vastly different computational costs. The proposed adjoint-based variational method is matrix-free, and evolves loops directly by the dynamical system in loop-space $\partial \mathbf{l} / \partial \tau = \mathbf{G}(\mathbf{l})$. On the contrary, the existing Newton-based variational method, first, constructs the full Jacobian matrix of an entire loop, and then evolves the loop by solving the sys-

TABLE I. Scaling of computational costs and memory requirements of the proposed adjoint-based variational method in comparison to the existing Newton-based method [34, 35]. The measured time per iteration on a 2.7 GHz Dual-Core Intel Core i5 refers to the KSE example with $N = 64$ spatial degrees of freedom. In the temporal direction, the Newton-based method requires a resolution of $M = 256$ to reach the same accuracy as the adjoint-based one for $M = 64$.

Method	Newton-based	Adjoint-based
G(l) computation	—	$\mathcal{O}(MN \log(MN))$
Jacobian construction	$\mathcal{O}(MN^2 \log(N))$	—
Linear system solution	$\mathcal{O}((MN)^{3/2})$	—
Time per iteration [ms]	4392	2
Memory requirements	$\mathcal{O}(MN^2)$	$\mathcal{O}(MN)$

tem of linear equations for the update vector of the Newton iteration (see [34] for details).

The scalings of the computational costs for both methods, with N the number of degrees of freedom in spatial and M in temporal direction are summarized in Table I. The computational cost of the adjoint-based variational method is dominated by two-dimensional fast Fourier transforms (FFT) requiring $\mathcal{O}(MN \log(MN))$ computations in each iteration. The Newton-based method first constructs a Jacobian matrix of dimension MN -by- MN , which is sparse containing M blocks of N -by- N matrices, 4 diagonal bands and one additional horizontal and vertical band each. The cost of constructing the Jacobian matrix scales like $\mathcal{O}(MN^2 \log(N))$. Solving the sparse system of linear equations for the update vector takes $\mathcal{O}((MN)^{3/2})$ operations so that depending on the size of the system in spatial and temporal direction, either the construction of the Jacobian or the solution of the linear system dominate the computational costs. For large N characteristic of 3D fluid flow problems, the adjoint-based method is expected to reduce the computational costs per iteration by a factor of $\mathcal{O}(N)$.

In addition to the scalings as a function of degrees of freedom, the lower accuracy of finite-difference approximations of the loop tangent in the Newton-based method compared to the spectral representation in the adjoint-based method also implies that equivalent accuracy requires an increased number of temporal sections for discretizing the loop in the Newton-based method. For the periodic orbit discussed in Fig. 2 $M = 256$ temporal sections in the Newton-based variational method are required to reach the same accuracy as $M = 64$ temporal modes in the adjoint-based method. Here we characterize the accuracy by relative changes of the period T when doubling the temporal resolution M . For the equivalent resolution, one iteration of the adjoint-based variational method takes only 2 ms, while each iteration of the Newton-based method takes 4392 ms on a 2.7 GHz Dual-Core Intel Core i5. Note however that while the computational cost per iteration of the adjoint-based variational method even for the KSE are more than three orders of magnitude smaller than the costs

of the Newton-based method, the adjoint-based method converges only exponentially and thus slower than the quadratically converging Newton-based method. Consequently, the adjoint-based variational method requires more iterations to achieve converge. While the adjoint-based method can be accelerated by exploiting the approximately linear loop dynamics around periodic orbits (see Appendix C for an example) the main advantage is not its speed but the vastly reduced memory requirements that allow to tackle larger problems.

Memory requirement are estimated for both the adjoint-based variational method and the Newton-based variational method and reported in Table I. The memory requirement of the Newton-based method is dominated by the number of non-zero elements in the sparse Jacobian matrix, scaling like $\mathcal{O}(MN^2)$. In contrast to the Newton-based method, the adjoint-based method is matrix-free, and only stores spatio-temporal fields with MN elements. Thus, the memory requirement of the adjoint-based method scales with $\mathcal{O}(MN)$, which is $\mathcal{O}(N)$ less than that of the Newton-based method. For the convergence of the discussed KSE test case, the Newton-based method requires more than 2 orders of magnitude more memory than the adjoint-based method. Due to the enormous memory requirements and computational costs, extending the Newton-based variational method to high-dimensional systems such as fluid flows, where N is often more than 10^6 , is practically impossible. On the contrary, owing to the manageable memory requirements and computational costs per iteration, the adjoint-based method can be applied to high-dimensional systems.

V. SUMMARY AND CONCLUSION

Unstable periodic orbits have been recognized as building blocks of the dynamics in driven dissipative spatio-temporally chaotic systems including fluid turbulence. Periodic orbits capture key features of the dynamics and reveal physical processes sustaining the turbulent flow. Constructing a sufficiently large set of periodic orbits moreover carries the hope to eventually yield a predictive rational theory of turbulence, where ‘properties of the turbulent flow can be mathematically deduced from the fundamental equations of hydrodynamics’, as expressed by Hopf in 1948 [3]. Despite the importance of unstable periodic orbits, computing these exact solutions for high-dimensional spatio-temporally chaotic systems remains challenging. Known methods either show poor convergence properties because they are based on time-marching a chaotic system causing exponential error amplification; or they require constructing Jacobian matrices which is prohibitively expensive for high-dimensional problems. We therefore introduce a new matrix-free method for computing periodic orbits that is unaffected by exponential error amplification, shows robust convergence properties and can be applied to high-dimensional

spatio-temporally chaotic systems. As a proof-of-concept we implement the method for the one-dimensional KSE and demonstrate the convergence of periodic orbits underlying spatio-temporal chaos.

The adjoint-based variational method constructs a dynamical system that evolves entire loops such that the value of a cost function measuring deviations of the loop from a solution of the governing equations monotonically decreases. Periodic orbits correspond to attracting fixed points of the variational dynamics. Due to the variational approach, the method provides a large radius of convergence so that periodic orbits can be found from inaccurate initial guesses. For the KSE we demonstrate the robust convergence properties by successfully computing periodic orbits from inaccurate initial guesses. These guesses are extracted from the projection of the free chaotic dynamics on a single scalar quantity, instead from close recurrences based on the L_2 -distance between spatial fields [16]. Reliable convergence to machine precision is observed independent of the period of the orbit.

The large convergence radius of the adjoint-based variational method relaxes accuracy requirements for initial guesses when those are extracted from the chaotic dynamics. Since initial guesses are characterized by an entire loop, one may use fast-to-compute models approximating the full dynamics to construct initial guesses for periodic orbits of the full dynamics. Such an approach would not be reasonable for classical shooting methods where initial guesses are characterized by an instantaneous initial condition and the difference between model and full dynamics would be amplified exponentially by the time-marching. Suitable models that may help provide initial guesses for constructing large sets of periodic orbits for a given chaotic system include under-resolved simulations, spatially filtered equations such as LES in fluids applications [46] and classical POD / DMD based models [47]. In addition, recent breakthroughs in machine learning allow to create data-driven low-dimensional models of the chaotic dynamics that replicate spatio-temporal chaos in one- and two-dimensional systems with remarkable accuracy [48–50].

The feasibility of the proposed method has been demonstrated for a one-dimensional chaotic PDE but the method applies to general autonomous systems and we plan to implement it for the full three-dimensional Navier-Stokes equations. Specifically, we aim for an implementation within our own open-source software Channelflow (channelflow.ch) [51]. In the context of this software not only the identification of periodic orbits but also their numerical continuation will benefit from the adjoint-based variational approach. When transferring the adjoint-based variational approach to three-dimensional fluid turbulence, we envision further optimizations of the method. First, we will exploit that during its approach to the attracting fixed point representing the periodic orbit, the evolution is well approximated by the linearization of the dynamics around the attracting fixed point. This allows to accelerate the time-marching

in loop-space and thereby the exponential convergence, as exemplified for the KSE. Second, one may complement the adjoint dynamics with Newton descent to identify the attracting fixed point in loop-space, following the analogous hybrid approach for identifying equilibrium solutions [38]. Alternatively, we will combine the adjoint-based variational method with a Newton-GMRES-based shooting method. Such a hybrid method offers the large radius of convergence of the adjoint-based variational method in combination with the fast quadratic convergence of Newton’s method. To allow for converging long and unstable periodic orbits, a multi-shooting variant of the standard Newton-GMRES-hook-step method [52] will be used.

ACKNOWLEDGMENTS

This work was supported by the Swiss National Science Foundation (SNSF) under grant no. 200021-160088 and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 865677). We thank Florian Reetz for insightful discussions on the implementation of the proposed method both for the KSE but also for future implementations within Channelflow. SA acknowledges support by the State Secretariat for Education, Research and Innovation SERI via the Swiss Government Excellence Scholarship.

Appendix A: Rate of change of the cost function J

The rate of change of the cost function J with respect to the fictitious time τ is given in Eq. (12). Here we derive this expression including the specific form of \mathbf{R} . With the definition of the cost function J (5)

$$J(\mathbf{l}) = \int_0^1 \int_{\Omega} \bar{\mathbf{r}}(\mathbf{l}) \cdot \bar{\mathbf{r}}(\mathbf{l}) d\vec{x} ds,$$

the rate of change of J with respect to the fictitious time τ is

$$\frac{\partial J}{\partial \tau} = 2 \int_0^1 \int_{\Omega} (\nabla_1 \bar{\mathbf{r}} \cdot \mathbf{G}) \cdot \bar{\mathbf{r}} d\vec{x} ds.$$

where $\partial \mathbf{l} / \partial \tau = \mathbf{G}$ from definition (10) has been used. Using the definition of the inner product in the space of generalized loops (8), we can express the rate of change as

$$\frac{\partial J}{\partial \tau} = 2 \left\langle \begin{bmatrix} \nabla_1 \bar{\mathbf{r}} \cdot \mathbf{G} \\ 0 \end{bmatrix}, \begin{bmatrix} \bar{\mathbf{r}} \\ 0 \end{bmatrix} \right\rangle.$$

Here we choose the second component of both generalized loops to be zero. With this choice, the rate of change of

J is given by

$$\frac{\partial J}{\partial \tau} = 2 \langle \mathcal{L}(\mathbf{l}; \mathbf{G}), \mathbf{R} \rangle,$$

where $\mathcal{L}(\mathbf{l}; \mathbf{G})$ indicates the directional derivative of $\mathbf{R} = [\vec{r}; 0]$ along \mathbf{G} , defined in (14).

Appendix B: Adjoint operator for KSE

We explicitly derive the form of the adjoint operator for the KSE problem given in Eq. (22). In this appendix, subscripts 1 and 2 denote the field component and the scalar component of generalized loops, respectively. The directional derivative of KSE along \mathbf{G} is

$$\mathcal{L}(\mathbf{l}; \mathbf{G}) = \begin{bmatrix} \frac{G_2}{T^2} \frac{\partial u}{\partial s} - \frac{1}{T} \frac{\partial G_1}{\partial s} - \frac{\partial(uG_1)}{\partial x} - \frac{\partial^2 G_1}{\partial x^2} - \frac{\partial^4 G_1}{\partial x^4} \\ 0 \end{bmatrix}$$

To compute the adjoint operator, we expand the inner product of the directional derivative of the residual and the residual itself:

$$\begin{aligned} \langle \mathcal{L}(\mathbf{l}; \mathbf{G}), \mathbf{R} \rangle &= \int_0^1 \int_0^L \mathcal{L}_1 R_1 dx ds + \mathcal{L}_2 R_2 \\ &= \int_0^1 \int_0^L \mathcal{L}_1 R_1 dx ds + 0 \\ &= \int_0^1 \int_0^L \left(\frac{G_2}{T^2} \frac{\partial u}{\partial s} - \frac{1}{T} \frac{\partial G_1}{\partial s} - \frac{\partial(uG_1)}{\partial x} - \frac{\partial^2 G_1}{\partial x^2} - \frac{\partial^4 G_1}{\partial x^4} \right) R_1 dx ds \\ &= \int_0^1 \int_0^L \frac{G_2}{T^2} \frac{\partial u}{\partial s} R_1 dx ds \\ &\quad + \int_0^1 \int_0^L \left(-\frac{1}{T} \frac{\partial G_1}{\partial s} - \frac{\partial(uG_1)}{\partial x} - \frac{\partial^2 G_1}{\partial x^2} - \frac{\partial^4 G_1}{\partial x^4} \right) R_1 dx ds. \end{aligned} \quad (\text{B1})$$

This inner product must be equal to

$$\langle \mathbf{G}, \mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}) \rangle = \int_0^1 \int_0^L \mathcal{L}_1^\dagger G_1 dx ds + \mathcal{L}_2^\dagger G_2, \quad (\text{B2})$$

where the adjoint operator is indicated by a dagger. Direct comparison of equations (B1) and (B2) results in

$$\begin{aligned} \int_0^1 \int_0^L \mathcal{L}_1^\dagger G_1 dx ds &= \int_0^1 \int_0^L \left(-\frac{1}{T} \frac{\partial G_1}{\partial s} - \frac{\partial(uG_1)}{\partial x} - \frac{\partial^2 G_1}{\partial x^2} - \frac{\partial^4 G_1}{\partial x^4} \right) R_1 dx ds \\ &= \int_0^1 \int_0^L \frac{1}{T^2} \frac{\partial u}{\partial s} R_1 dx ds \end{aligned} \quad (\text{B3a})$$

$$\mathcal{L}_2^\dagger G_2 = \left(\int_0^1 \int_0^L \frac{1}{T^2} \frac{\partial u}{\partial s} R_1 dx ds \right) G_2. \quad (\text{B3b})$$

The form of \mathcal{L}_2^\dagger is directly given by (B3b):

$$\mathcal{L}_2^\dagger(\mathbf{q}; \mathbf{R}) = \int_0^1 \int_0^L \frac{1}{T^2} \frac{\partial u}{\partial s} R_1 dx ds.$$

Using integration by parts and the periodicity of the domain in space and time, Eq. (B3a) becomes

$$\begin{aligned} \int_0^1 \int_0^L \mathcal{L}_1^\dagger G_1 dx ds &= \int_0^1 \int_0^L \left(\frac{1}{T} \frac{\partial R_1}{\partial s} + u \frac{\partial R_1}{\partial x} - \frac{\partial^2 R_1}{\partial x^2} - \frac{\partial^4 R_1}{\partial x^4} \right) G_1 dx ds. \end{aligned}$$

Consequently,

$$\mathcal{L}_1^\dagger(\mathbf{l}; \mathbf{R}) = \frac{1}{T} \frac{\partial R_1}{\partial s} + u \frac{\partial R_1}{\partial x} - \frac{\partial^2 R_1}{\partial x^2} - \frac{\partial^4 R_1}{\partial x^4}$$

where $R_1 = r$. The adjoint operator acting on loops therefore has the form

$$\mathcal{L}^\dagger(\mathbf{l}; \mathbf{R}) = \begin{bmatrix} \frac{1}{T} \frac{\partial r}{\partial s} + u \frac{\partial r}{\partial x} - \frac{\partial^2 r}{\partial x^2} - \frac{\partial^4 r}{\partial x^4} \\ \int_0^1 \int_0^L \frac{1}{T^2} \frac{\partial u}{\partial s} r dx ds \end{bmatrix}.$$

Note that, owing to the periodic boundary conditions in both spatial and temporal directions, here, the boundary terms that appear during integration by parts cancel. In general, however, boundary terms do not vanish so that the derivation of adjoint operators for other boundary conditions may be less straightforward.

Appendix C: Acceleration of the convergence by linearized approximation

We demonstrate a straightforward method for accelerating the convergence of the adjoint-based variational method. We iterate between time-stepping of the dynamical system in loop-space (23) and a linear extrapolation along the evolution trajectory of the loops. This extrapolation is based on the assumption that the evolution follows the leading eigendirection of the linearization about the attracting loop. Extrapolations yield the initial conditions of the subsequent advancing of the loop in τ . This procedure is repeated until the periodic orbit is converged. Figure 7 compares the convergence of the periodic orbit shown in Figs. 2 and 3 by continuous integration of the dynamical system in loop-space (23) and the accelerated method iterating between time-stepping of the full dynamics and extrapolations, both from the same initial condition. Vertical drops of the cost function shown in the graph correspond to the extrapolations. In this example the accelerated method reduces the required total number of numerical steps of integration by more than 50%.




FIG. 7. Accelerated convergence of the adjoint-based variational method. Convergence history for the periodic orbit discussed in figures 2 and 3, for the standard method (orange dashed line) and the modified method involving linear extrapolations along the solution trajectory in the loop-space. The linear extrapolations are based on a linear approximation of the loop dynamics around the attracting fixed point in loop-space corresponding to the periodic orbit. The square root of the cost function is shown as a function of the number of fictitious time steps n . The first extrapolation is performed when $\sqrt{J} = 10^{-3}$. Between two consecutive extrapolations, the dynamical system in loop-space is integrated until the value \sqrt{J} is halved. In this example case, extrapolations reduce the total number of fictitious time steps by more than 50%.

Appendix D: Convergence to local and global minima of J

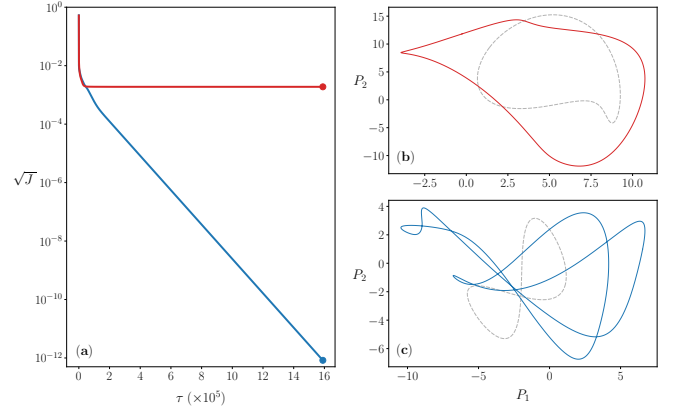


FIG. 8. Minimizing J by the adjoint-based variational method. (a) Evolution of \sqrt{J} with τ for two different initial loops. The blue line shows the convergence for the loop that approaches a periodic orbit with $J = 0$ while the red line shows the convergence for a loop that approaches a local minimum of J with a nonzero value $J > 0$. The corresponding initial (dashed lines) and converged loops (solid lines) in the two-dimensional projection of the state space as in Fig. 2 are visualized for the converged loop with $J > 0$ (panel b) and the periodic orbit with $J \rightarrow 0$ (panel c).

In Figure 8 we show an example of time-stepping of the dynamical system in loop-space where the final loop corresponds to local minimum of J with a nonzero value. Consequently, no periodic orbit is found.

-
- [1] T. Y. Thomas, Qualitative Analysis of the Flow of Fluids in Pipes, *American Journal of Mathematics* **64**, 754 (1942).
 - [2] C. C. Lin, On the Stability of Two-Dimensional Parallel Flows, *Proceedings of the National Academy of Sciences* **30**, 316 (1944).
 - [3] E. Hopf, A mathematical example displaying features of turbulence, *Communications on Pure and Applied Mathematics* **1**, 303 (1948).
 - [4] J. Jiménez, Coherent structures in wall-bounded turbulence, *Journal of Fluid Mechanics* **842**, P1 (2018).
 - [5] R. R. Kerswell, Recent progress in understanding the transition to turbulence in a pipe, *Nonlinearity* **18**, R17 (2005).
 - [6] B. Eckhardt, T. M. Schneider, B. Hof, and J. Westerweel, Turbulence Transition in Pipe Flow, *Annual Review of Fluid Mechanics* **39**, 447 (2007).
 - [7] G. Kawahara, M. Uhlmann, and L. van Veen, The Significance of Simple Invariant Solutions in Turbulent Flows, *Annual Review of Fluid Mechanics* **44**, 203 (2012).
 - [8] J. F. Gibson, J. Halcrow, and P. Cvitanović, Visualizing the geometry of state space in plane Couette flow, *Journal of Fluid Mechanics* **611**, 107 (2008).
 - [9] P. Cvitanović and J. F. Gibson, Geometry of the turbulence in wall-bounded shear flows: periodic orbits, *Physica Scripta* **T142**, 014007 (2010).
 - [10] B. Suri, J. Tithof, R. O. Grigoriév, and M. F. Schatz, Forecasting Fluid Flows Using the Geometry of Turbulence, *Physical Review Letters* **118**, 114501 (2017).
 - [11] F. Reetz, T. Kreilos, and T. M. Schneider, Exact invariant solution reveals the origin of self-organized oblique turbulent-laminar stripes, *Nature Communications* **10**, 2277 (2019).
 - [12] G. Kawahara and S. Kida, Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst, *Journal of Fluid Mechanics* **449**, 291 (2001).
 - [13] H. Poincaré, *Il Nuovo Cimento*, Vol. 10 (Gauthier-Villars, Paris, 1892).
 - [14] D. Ruelle, *Thermodynamic formalism: The mathematical structures of classical equilibrium statistical mechanics*, 1st ed. (Addison-Wesley, 1978).
 - [15] M. C. Gutzwiller, Chaos in classical and quantum mechanics, *Journal of Physics A: Mathematical and Theoretical Interdisciplinary Applied Mathematics*, **43**, 285302 (1990).
 - [16] D. Auerbach, P. Cvitanović, J.-P. Eckmann, G. Gu-

- naratne, and I. Procaccia, Exploring chaotic motion through periodic orbits, *Physical Review Letters* **58**, 2387 (1987).
- [17] P. Cvitanović, Invariant measurement of strange sets in terms of cycles, *Physical Review Letters* **61**, 2729 (1988).
- [18] R. Artuso, E. Aurell, and P. Cvitanović, Recycling of strange sets: I. Cycle expansions, *Nonlinearity* **3**, 325 (1990).
- [19] R. Artuso, E. Aurell, and P. Cvitanović, Recycling of strange sets: II. Applications, *Nonlinearity* **3**, 361 (1990).
- [20] Y. Lan, Cycle expansions: From maps to turbulence, *Communications in Nonlinear Science and Numerical Simulation* **15**, 502 (2010).
- [21] P. Cvitanović, R. Artuso, G. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and quantum* (chaos-book.org, Niels Bohr Institute, Copenhagen, 2016).
- [22] G. J. Chandler and R. R. Kerswell, Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow, *Journal of Fluid Mechanics* **722**, 554 (2013).
- [23] D. Lasagna, Sensitivity Analysis of Chaotic Systems Using Unstable Periodic Orbits, *SIAM Journal on Applied Dynamical Systems* **17**, 547 (2018).
- [24] C. T. Kelley, *Solving Nonlinear Equations with Newton's Method* (Society for Industrial and Applied Mathematics, 2003).
- [25] J. Sanchez, M. Net, B. Garcia-Archilla, and C. Simo, Newton–Krylov continuation of periodic orbits for Navier–Stokes flows, *Journal of Computational Physics* **201**, 13 (2004).
- [26] D. Viswanath, Recurrent motions within plane Couette turbulence, *Journal of Fluid Mechanics* **580**, 339 (2007).
- [27] D. Viswanath, The critical layer in pipe flow at high Reynolds number, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **367**, 561 (2009).
- [28] J. E. Dennis and R. B. Schnabel, *Numerical methods for unconstrained optimization and nonlinear equations* (SIAM, 1996).
- [29] Y. Duguet, C. C. T. Pringle, and R. R. Kerswell, Relative periodic orbits in transitional pipe flow, *Physics of Fluids* **20**, 114102 (2008).
- [30] L. van Veen, A. Vela-Martin, and G. Kawahara, Time-periodic inertial range dynamics, *arXiv preprint* (2018).
- [31] R. H. G. Helleman and T. Bountis, Periodic solutions of arbitrary period, variational methods, in *Stochastic Behavior in Classical and Quantum Hamiltonian Systems* (Springer-Verlag, Berlin/Heidelberg, 1979) pp. 353–375.
- [32] L. Drossos, O. Ragos, M. N. Vrahatis, and T. Bountis, Method for computing long periodic orbits of dynamical systems, *Physical Review E* **53**, 1206 (1996).
- [33] M. VRAHATIS, T. BOUNTIS, and M. KOLLMANN, PERIODIC ORBITS AND INVARIANT SURFACES OF 4D NONLINEAR MAPPINGS, *International Journal of Bifurcation and Chaos* **06**, 1425 (1996).
- [34] Y. Lan and P. Cvitanović, Variational method for finding periodic orbits in a general flow, *Physical Review E* **69**, 016217 (2004).
- [35] Y. Lan and P. Cvitanović, Unstable recurrent patterns in Kuramoto-Sivashinsky dynamics, *Physical Review E* **78**, 026208 (2008).
- [36] L. M. Fardheiro, B. M. Boghosian, P. V. Coveney, and J. Lätt, Unstable periodic orbits in weak turbulence, *Journal of Computational Science* **1**, 13 (2010).
- [37] B. M. Boghosian, A. Brown, J. Lätt, H. Tang, L. M. Fardheiro, and P. V. Coveney, Unstable periodic orbits in the Lorenz attractor., *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* **369**, 2345 (2011).
- [38] M. Farazmand, An adjoint-based approach for finding invariant solutions of Navier-Stokes equations, *Journal of Fluid Mechanics* **795**, 278 (2016).
- [39] Y. Kuramoto and T. Tsuzuki, Persistent Propagation of Concentration Waves in Dissipative Media Far from Thermal Equilibrium, *Progress of Theoretical Physics* **55**, 356 (1976).
- [40] G. I. Sivashinsky, Nonlinear analysis of hydrodynamic instability in laminar flames—I. Derivation of basic equations, *Acta Astronautica* **4**, 1177 (1977).
- [41] 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**, 1 (2010).
- [42] Y. S. Smyrlis and D. T. Papageorgiou, *Computational study of chaotic and ordered solutions of the Kuramoto-Sivashinsky equation*, Tech. Rep. (Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, 1996).
- [43] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang, *New York* (Springer, 2006) p. 563.
- [44] M. J. Powell, *A Fortran subroutine for solving systems of nonlinear algebraic equations*, Tech. Rep. (Atomic Energy Research Establishment, Harwell, England (United Kingdom), 1968).
- [45] A. R. Conn, N. I. M. Gould, and P. L. Toint, *Trust Region Methods* (Society for Industrial and Applied Mathematics, 2000).
- [46] P. Sagaut, *Large eddy simulation for incompressible flows: an introduction* (Springer Science & Business Media, 2006).
- [47] B. J. McKeon, The engine behind (wall) turbulence: Perspectives on scale interactions, *Journal of Fluid Mechanics* **817**, P1 (2017).
- [48] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach, *Physical Review Letters* **120**, 024102 (2018).
- [49] R. S. Zimmermann and U. Parlitz, Observing spatio-temporal dynamics of excitable media using reservoir computing, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **28**, 043118 (2018).
- [50] P. R. Vlachas, W. Byeon, Z. Y. Wan, T. P. Sapsis, and P. Koumoutsakos, Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science* **474**, 20170844 (2018).
- [51] J. F. Gibson, F. Reetz, S. Azimi, A. Ferraro, T. Kreilos, H. Schrobsdorff, M. Farano, A. F. Yesil, S. S. Schütz, M. Culp, and T. M. Schneider, Channelflow 2.0, in preparation (2021).
- [52] J. Sánchez and M. Net, On the Multiple Shooting Continuation of Periodic Orbits By Newton–Krylov Methods, *International Journal of Bifurcation and Chaos* **20**, 43 (2010).