

Hyperbolicity, shadowing directions and sensitivity analysis of a turbulent three-dimensional flow

Angxiu Ni[†]

Department of Mathematics, University of California, Berkeley, Berkeley, CA 94720, USA

(Received 18 May 2018; revised 24 November 2018; accepted 30 November 2018;
first published online 28 January 2019)

This paper uses compressible flow simulation to analyse the hyperbolicity, shadowing directions and sensitivities of a weakly turbulent three-dimensional cylinder flow at Reynolds number 525 and Mach number 0.1. By computing the first 40 covariant Lyapunov vectors (CLVs), we find that unstable CLVs are active in the near-wake region, whereas stable CLVs are active in the far-wake region. This phenomenon is related to hyperbolicity since it shows that CLVs point to different directions; it also suggests that for open flows there is a large fraction of CLVs that are stable. However, due to the extra neutral CLV and the occasional tangencies between CLVs, our system is not uniform hyperbolic. By the non-intrusive least-squares shadowing (NILSS) algorithm, we compute shadowing directions and sensitivities of long-time-averaged objectives. Our results suggest that shadowing methods may be valid for general chaotic fluid problems.

Key words: chaos, chaotic advection, wakes

1. Introduction

The chaotic dynamics of many fluid problems, such as turbulence and vortex streets, are central to many important challenges facing science and technology. These chaotic flows are typically controlled by system parameters, such as the incoming flow conditions or the boundary geometry. This paper studies the three-dimensional (3D) flow past a cylinder at Reynolds number 525 from a dynamical system point of view. More specifically, we study perturbations in flow fields due to perturbations in initial conditions and system parameters, which are covariant Lyapunov vectors and shadowing directions of this flow problem. With the shadowing directions, we can then study perturbations in the long-time-averaged objectives due to perturbations in system parameters.

Chaotic fluid systems depend sensitively on initial conditions: perturbations whose norms grow like exponential functions are called covariant Lyapunov vectors (CLVs), and their exponents are the Lyapunov exponents (LEs) (Katok & Hasselblatt 1997). Unstable CLVs determine the uncertainty in predicting the future behaviour of the dynamical system, provided the past trajectory is known up to some precision

[†] Email address for correspondence: niangxiu@gmail.com

(Young 2003). More importantly, unstable CLVs can be regarded as chaotic degrees of freedom (DOFs) to which the apparently complicated behaviour can be attributed. This reduction in DOFs allows reduced-order numerical methods, such as those developed by Chorin, Hald & Kupferman (2002) and Parish & Duraisamy (2017). In experiments, we can infer some properties of CLVs by reconstructed attractors (Sieber 1987), but more details can be obtained only numerically. In particular, Keefe, Moin & Kim (1992) computed the LE spectrum of turbulent Poiseuille flow. Recently, the LE spectra for some 3D computational fluid dynamics (CFD) problems have been studied (Blonigan *et al.* 2016a; Fernandez & Wang 2017a).

There remain questions about active areas of CLVs in ‘open’ flows, or developing flows such as boundary layers, jets and wakes. Since perturbations in open flows are typically convected downstream, it is difficult to form an intuitive visualization of the CLVs. In contrast, for ‘closed’ flows such as Bénard convection and Taylor–Couette flow, global dynamics are expressed at all spatial locations, and CLVs can be effectively interpreted as eigenmodes (Keefe *et al.* 1992). Fernandez & Wang (2017b) plotted the flow fields of CLVs in an open flow, and they found no significant difference in active areas among CLVs, perhaps because only a few CLVs were computed. However, in this paper, we are able to observe the difference in active areas for different CLVs.

Hyperbolicity is a property of dynamical systems which says that the unstable and stable CLVs are bounded away from each other. Uniform hyperbolicity is assumed in many theoretical results, such as the existence of a steady distribution (Young 2002), linear response formula (Ruelle 1997, 2008) and shadowing lemma (Bowen 1970; Pilyugin 1999). However, many chaotic fluid systems are not uniform hyperbolic. For a Kolmogorov flow simulated with 224 DOFs (Inubushi *et al.* 2012) and for a 3D Boussinesq equation simulated with 5×10^5 DOFs (Xu & Paul 2016), researchers found for both cases that the angles between CLVs can be very small. Currently, there are not many results regarding hyperbolicity for CFD simulated 3D Navier–Stokes fluid systems. In this paper, we conclude that our fluid system is not uniform hyperbolic; however, we find a more robust phenomenon related to hyperbolicity.

It has been conjectured that many theoretical results, although mathematically developed under uniform hyperbolicity, are still valid in most chaotic systems (Ruelle 1980; Gallavotti & Cohen 1995; Gallavotti 2006). Support for such conjectures can be found in Reick (2002), Albers & Sprott (2006) and Ni & Wang (2017). In this paper, we are particularly interested in whether shadowing methods are still valid, even though our system is not uniform hyperbolic.

For a chaotic system, if we coordinate carefully perturbations on both initial conditions and parameters, we can find a shadowing trajectory which remains close to the old trajectory (Bowen 1970). The first-order approximation of the difference between the two trajectories is the shadowing direction (Pilyugin 1997). Shadowing methods enable sensitivity analysis of long-time-averaged objectives (Wang, Hu & Blonigan 2014), which is useful in design (Jameson 1988), control (Bewley 2001), inverse problems (Tromp, Tape & Liu 2005), error estimation (Becker & Rannacher 2001; Giles & Süli 2002; Fidkowski & Darmofal 2011) and data assimilation (Thépaut & Courtier 1991). Shadowing is difficult to observe experimentally, and the only method we are aware of is numerical. The theory for shadowing methods typically assumes uniform hyperbolicity; however, in this paper, we show that shadowing directions also exist for our fluid system.

The least-squares shadowing (LSS) method (Wang 2014; Wang *et al.* 2014) computes sensitivities of long-time-averaged objectives via computing the shadowing

direction. With high cost, LSS has been successfully applied in two-dimensional (2D) CFD problems (Blonigan *et al.* 2016b). The non-intrusive least-squares shadowing (NILSS) method (Ni *et al.* 2016; Ni & Wang 2017) reformulates LSS by constraining computation to the unstable subspace. For real-life problems, such as 2D flow over a step computed in Ni & Wang (2017), NILSS can be hundreds or thousands of times faster than LSS Blonigan *et al.* (2016b). One variant of NILSS is the adjoint, such as one developed by Blonigan (2017), and the non-intrusive least-squares adjoint shadowing (NILSAS) algorithm developed by Ni (2018a). In particular, NILSAS is based on new theoretical progress in adjoint shadowing directions (Ni 2018b). Another variant is the finite difference NILSS (FD-NILSS) algorithm (Ni *et al.* 2018), whose implementation requires only primal solvers. NILSS and its variants have been successfully used for several other turbulent fluid system, such as a 2D flow over a backward step (Ni & Wang 2017), and a minimal flow unit of a channel flow (Blonigan 2017); NILSS has also been successfully applied to easier mathematical models such as the Lorenz 63 and Kuramoto–Sivashinsky (KS) systems. In this paper, we use FD-NILSS to compute shadowing directions and sensitivities of our problem.

We begin the main part of the paper by describing the physical and numerical set-up for our flow problem. Then the rest of the paper consists of two logically connected parts. First, we review LEs and CLVs, and compute LEs and CLVs for our flow problem. Although the CLV result violates the uniform hyperbolicity, it still indicates some form of hyperbolicity, which encourages us to proceed to the second part of the paper, where we review shadowing directions and the FD-NILSS algorithm, using which we compute shadowing directions and sensitivities of several long-time-averaged objectives.

2. Problem set-up and verification of simulation

Our physical problem of the 3D flow past a cylinder is the same as in Ni *et al.* (2018). The front view of the geometry of the entire flow field is shown in figure 1. In the rest of the paper, the units we have in mind are SI units, but since in this paper all values are normalized, and in the software all equations are dimensionless, readers may take any compatible set of units. The diameter of the cylinder is $D = 0.25 \times 10^{-3}$. The spanwise width is $Z = 2D$. The free-stream conditions are: density $\rho_0 = 1.18$, pressure $P_0 = 1.01 \times 10^5$, temperature $T_0 = 298$ and dynamic viscosity $\mu = 1.86 \times 10^{-5}$. The free-stream flow is in the x -direction, with the velocity U being one of the system parameters, and for the base case $U_0 = 33.0$. The flow-through time t_0 , defined as the time for U_0 flowing past the cylinder, is $t_0 = D/U_0 = 7.576 \times 10^{-6}$. The Reynolds number of the base case is $Re = \rho_0 U_0 D / \mu = 525$ and the Mach number is 0.1. The cylinder can rotate around its centre with rotational speed ω , which is the second system parameter for our problem. The ω value is measured in revolutions per unit time, and its positive direction is anticlockwise, as shown in figure 1. For the cylinder to rotate one cycle per flow-through time, the rotation speed $\omega_0 = 1/t_0 = 1.32 \times 10^5$.

We run our simulations on two sets of block-structured hexahedral meshes. The coarser mesh has 3.7×10^5 cells, and finer mesh has 7.5×10^5 cells. A 2D slice of the finer mesh is shown in figure 2. For both meshes, the spanwise direction has 48 cells. The CFD solver we use is CharLES developed at Cascade Technologies (Brès *et al.* 2017), using which we perform under-resolved direct numerical simulation (uDNS) without turbulence model. The accuracy of the solver is formally second order in space and third order in time. The spanwise boundaries use periodic boundary




FIGURE 1. (Colour online) Geometry used in the simulation of a 3D flow past a cylinder. The spanwise extent of the computational domain is $Z = 2D$. The positive direction of the cylinder rotational speed ω is anticlockwise.

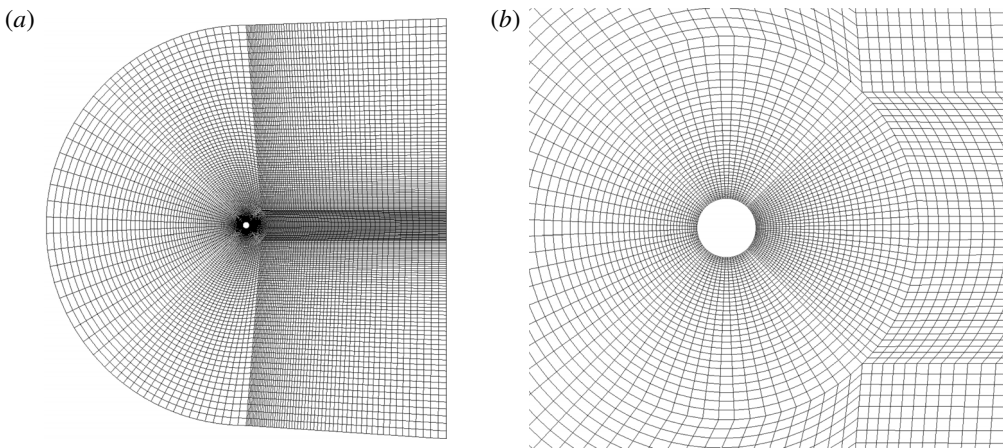


FIGURE 2. (a) A 2D slice of the finer mesh over the entire computational domain. (b) A zoom around the cylinder. This is a block-structured mesh with 7.5×10^5 hexahedra. The spanwise direction has 48 cells.

conditions; the left boundary uses a convective boundary condition (Colonius, Lele & Moin 1993); the right boundary uses the Navier–Stokes characteristic boundary condition (NSCBC) (Poinsot & Lele 1992). The time-step size is $\Delta t = 10^{-8} = 1.32 \times 10^{-3} t_0$.

Some 2D snapshots of the flow field simulated with the above numerical settings on the finer mesh are shown in figure 3. The flow is chaotic and 3D. We also plotted a time-averaged flow field in figure 4. As we can see, the computed flow field generally fits our fluid intuitions in the stagnant area before and after the cylinder, the acceleration area around the cylinder, and the boundary layer and wake structures.

We compare our results with previous literature. The same physical problem has been investigated in experiments by Williamson & Roshko (1990), and in numerical simulations by Mittal & Balachandar (1996). We compare the Strouhal number S_t and




FIGURE 3. Some 2D slices of the flow field on the finer mesh. (a) Vertical cross-section, plotted by the magnitude of velocity. (b) Horizontal cross-section, plotted by the spanwise velocity; this shows that the flow is 3D. All velocities are normalized by the free-stream velocity $U_0 = 33.0$.




FIGURE 4. Flow field on the finer mesh averaged over a time span of $168t_0$, plotted by the ρU component, normalized by $\rho_0 U_0$.

the averaged drag coefficient C_D . Here the Strouhal number is defined by $S_t = fD/U$, where f is the main frequency of the vortex shedding, represented by the history of the lift; the drag coefficient $C_D = D_r / (0.5\rho_0 U^2 DZ)$, where D_r is the drag. As shown in table 1, our results on both meshes match previous experimental and numerical results.

3. Covariant Lyapunov vectors and hyperbolicity

3.1. Definitions of primal and tangent solutions

Both CLVs and shadowing directions are tangent solutions, which describe evolutions of trajectory perturbations caused by perturbations on a dynamical system. There are two kinds of perturbations we can perform on a dynamical system given in (3.1): perturbation on initial conditions and on system parameters. The corresponding tangent solutions are homogeneous and inhomogeneous tangent solutions, respectively. As we shall see, CLVs are special homogeneous tangent solutions, while shadowing directions are special inhomogeneous tangent solutions.

	S_l	C_D
Coarser mesh with 3.7×10^5 cells	0.21	1.22
Finer mesh with 7.5×10^5 cells	0.21	1.19
2D simulation by Mittal & Balachandar (1996)	0.22	1.44
3D simulation by Mittal & Balachandar (1996)	0.22	1.24
Experiment by Williamson & Roshko (1990)	0.21	1.15

TABLE 1. Comparison with previous literature for the Strouhal number S_l and the averaged drag coefficient C_D .

We write the governing equation of the flow field in the form of a general dynamical system, which will be referred to as the primal system:

$$\frac{d\mathbf{u}}{dt} = \mathbf{f}(\mathbf{u}, s), \quad \mathbf{u}|_{t=0} = \mathbf{u}^0 + \mathbf{v}^0 s + \mathbf{w}^0 \phi. \tag{3.1}$$

Here $\mathbf{f}(\mathbf{u}, s) : X \times \mathbb{R} \rightarrow X$ is a smooth function, \mathbf{u} are state variables, and s is the system parameter. The initial condition is \mathbf{u}^0 with a possible perturbation in the direction of \mathbf{v}^0 or \mathbf{w}^0 , controlled by s and another parameter ϕ . A primal solution $\mathbf{u} : \mathbb{R} \rightarrow X$ maps time to a compact phase space X , where a single point represents an entire 3D flow field at an instant; thus $\mathbf{u}(t)$ can as well be imagined as a trajectory in X .

We assume the system is autonomous, that is, \mathbf{f} does not depend on time. We could extend our discussions to some special non-autonomous cases, such as those with some compactness in the time dependence, so that the ergodic theory is still available. For example, for cases whose time dependence can be decomposed into finitely many frequencies, we can extend the phase space by adding finitely many tori and still maintain a compact phase space. On this extended phase space our analysis will still be valid; for example, Antali & Horvath (2018) used the shadowing method to compute sensitivities. However, our discussions in this paper do not extend to non-autonomous systems in the most general setting.

Our dynamical system considered in this paper is given by the semi-discretized Navier–Stokes equation without turbulence models. More specifically, on the spatial direction, the flow field, together with boundary conditions, is represented by the finite-volume method on our meshes. The state variables are conservative variables, that is, ρ , $\rho \mathbf{U}$ and ρE . Hence our phase space $X \subset \mathbb{R}^m$ is finite-dimensional, with m being the DOFs. We further assume that X is a bounded hence compact subset. For the coarser mesh, the number of DOFs is approximately 1.9×10^6 ; for the finer mesh, it is approximately 3.8×10^6 . On the other hand, we regard the temporal direction as being continuous; thus our dynamical system is continuous, and later the theories we use are for continuous dynamics.

We can differentiate (3.1) with respect to ϕ , and define $\mathbf{w} = d\mathbf{u}/d\phi$. Then \mathbf{w} satisfies the so-called homogeneous tangent equation:

$$\frac{d\mathbf{w}}{dt} - \partial_{\mathbf{u}} \mathbf{f} \mathbf{w} = 0, \quad \mathbf{w}(t=0) = \mathbf{w}^0, \tag{3.2}$$

where $\partial_{\mathbf{u}} \mathbf{f} \in \mathbb{R}^{m \times m}$, and the ordinary differential equation (ODE) is called the homogeneous tangent ODE. Here \mathbf{w} reflects the trajectory perturbation caused by perturbing the initial condition in the direction of \mathbf{w}^0 .

On the other hand, we can differentiate (3.1) with respect to s , and define $\mathbf{v} = d\mathbf{u}/ds$. Then \mathbf{v} satisfies the inhomogeneous tangent equation:

$$\frac{d\mathbf{v}}{dt} - \partial_u \mathbf{f} \mathbf{v} = \partial_s \mathbf{f}, \quad \mathbf{v}(t=0) = \mathbf{v}^0, \quad (3.3)$$

where $\partial_s \mathbf{f} \in \mathbb{R}^m$, and the ODE is called the inhomogeneous tangent ODE. Here \mathbf{v} reflects the trajectory perturbation caused by perturbing the system parameter s , which not only affects the governing differential equation, but also potentially impacts the initial condition in the direction of \mathbf{v}^0 . The solution set of an inhomogeneous tangent ODE can be written as a particular inhomogeneous \mathbf{v}^* plus the solution set of the homogeneous tangent ODE, which is a linear subspace due to homogeneity.

3.2. Definitions of covariant Lyapunov vectors and hyperbolicity

A CLV $\boldsymbol{\zeta}(t)$ is a homogeneous tangent solution whose norm behaves like an exponential function of time. That is, there are $C_1, C_2 > 0$ and $\lambda \in \mathbb{R}$ such that, for any t ,

$$C_1 e^{\lambda t} \|\boldsymbol{\zeta}(\mathbf{u}(0))\| \leq \|\boldsymbol{\zeta}(\mathbf{u}(t))\| \leq C_2 e^{\lambda t} \|\boldsymbol{\zeta}(\mathbf{u}(0))\|, \quad (3.4)$$

where the norms are Euclidean norms in \mathbb{R}^m , and λ is defined as the LE corresponding to this CLV. CLVs with positive LEs are said to be unstable, CLVs with negative LEs are stable, and CLVs with zero LEs are neutral. The existence of CLVs is given by the Oseledets theorem (Oseledets 1968), for which a proof in English can be found in Rezakhanlou (2018, chap. 4), and more reference can be found at Oseledets (2008). In this paper, we sort LEs by decreasing order, so the j th largest LE and its corresponding CLV will be called the j th LE and j th CLV, respectively.

Hyperbolicity is roughly defined as being that the tangent space at every state splits into a stable subspace, an unstable space and a neutral subspace. The one most widely used in ergodic theory is uniform hyperbolicity, which has two assumptions: (1) the three subspaces are uniformly bounded away from each other; and (2) the neutral subspace is one-dimensional. Our system cannot be uniform hyperbolic, since it violates the first assumption, that is, it has at least two neutral CLVs: the first one corresponds to the common time translation of continuous dynamical systems, and the second corresponds to spanwise translations due to the periodic boundary conditions.

Still, we want to check if our system violates the second assumption of uniform hyperbolicity. More specifically, we want to bound the norm of the operator which projects to one of the subspaces, i.e. C_α given in Ni (2018b), where we showed that $C_\alpha = (1 - \cos^2 \alpha)^{-1/2}$, with α being the smallest angle between subspaces. Of course, the ideal case would be that these subspaces are orthogonal to each other, leading to $\alpha = 90^\circ$ and $C_\alpha = 1$ takes its minimum. Both the bound of the norm of the shadowing direction (Pilyugin 1997) and the bound of the error in FD-NILSS (Ni & Wang 2017) are proportional to C_α , and we do not want it larger than 10, causing impact of an order of magnitude. The corresponding angle is $\alpha_0 = 5.7^\circ$, which serves as our threshold.

As we will see, the smallest angle between all CLVs can be affected by mesh or time length. However, we find a robust phenomenon related to hyperbolicity, that is, angles between two CLVs become much larger if their indices are further apart. We no longer have uniform hyperbolicity, since there is more than one neutral CLV and there may be tangencies between adjacent CLVs. On the other hand, this phenomenon

is more robust, and reveals more structures of CLVs not indicated even by uniform hyperbolicity. In fact, we can even visualize this phenomenon by showing that active areas of CLVs are different.

More importantly, we will show that for our problem, where we do not have uniform hyperbolicity, tools given by ergodic theories are still valid, although they mostly use uniform hyperbolicity as logical assumptions. Specifically, we can still compute shadowing directions which correctly reflect sensitivities of statistics of the fluid system. Our results confirm the chaotic hypothesis (Gallavotti 2008), which says that many tools logically consequential to the uniform hyperbolicity assumption are in fact valid for a larger class of chaotic systems.

3.3. Numerical methods for computing Lyapunov exponents and covariant Lyapunov vectors

The algorithm we use to compute LEs is given by Benettin *et al.* (1980), and the algorithm for CLVs is by Ginelli *et al.* (2007, 2013). Since these two algorithms share many of the same procedures, we describe them together. For $i = 0, \dots, K - 1$, we define the i th time segment as time span $[t_i, t_{i+1}]$, with $t_i = i\Delta T$. In the algorithm presented below, for quantities defined on the entire segments such as \mathbf{u}_i and \mathbf{W}_i , we use the same subscript as the segment they are defined on. For quantities defined only at interfaces between segments such as \mathbf{Q}_i and \mathbf{R}_i , we use the same subscript as the time point they are defined at.

To begin, we should prescribe: (1) number of LEs/CLVs to compute, M ; (2) length of each time segment, ΔT ; and (3) number of time segments, K . Consequently, the time length of the entire trajectory, $T = K\Delta T$, is determined. Then the algorithm is given by the following procedure.

1. Generate initial conditions for primal solutions and homogeneous tangent solutions.
 - 1.1. Compute the primal solution of (3.1) for sufficiently long time so that the trajectory lands onto the attractor, then set $t = 0$, and set initial condition of the primal system, $\mathbf{u}_0(0)$.
 - 1.2. Randomly generate an $m \times M$ orthogonal matrix $\mathbf{Q}_0 = [\mathbf{q}_{01}, \dots, \mathbf{q}_{0M}]$. This \mathbf{Q}_0 will be used as initial condition for homogeneous tangent solutions.
2. For $i = 0$ to $K - 1$, on segment i , where $t \in [t_i, t_{i+1}]$, do:
 - 2.1. Compute the primal solution $\mathbf{u}_i(t)$ from t_i to t_{i+1} .
 - 2.2. Compute homogeneous tangent solutions $\mathbf{W}_i(t) = [\mathbf{w}_{i1}(t), \dots, \mathbf{w}_{iM}(t)]$.
 - 2.2.1. For each homogeneous tangent solution \mathbf{w}_{ij} , $j = 1, \dots, M$, starting from initial condition $\mathbf{w}_{ij}(t_i) = \mathbf{q}_{ij}$, integrate (3.2) from t_i to t_{i+1} .
 - 2.2.2. Perform QR factorization: $\mathbf{W}_i(t_{i+1}) = \mathbf{Q}_{i+1}\mathbf{R}_{i+1}$, where $\mathbf{Q}_{i+1} = [\mathbf{q}_{i+1,1}, \dots, \mathbf{q}_{i+1,M}]$.
3. The j th largest LE, λ_j , is approximated by

$$\lambda_j = \frac{1}{K\Delta T} \sum_{i=1}^K \log |D_{ij}|, \tag{3.5}$$

where D_{ij} is the j th diagonal element in \mathbf{R}_i . This formula for λ_j converges to the true value as T becomes large.

4. We define an $m \times M$ matrix $\mathbf{V}(t)$ by

$$\mathbf{V}(t) = \mathbf{W}_i(t)\mathbf{R}_{i+1}^{-1} \cdots \mathbf{R}_K^{-1}, \quad t \in [t_i, t_{i+1}]. \tag{3.6}$$

The j th column of $\mathbf{V}(t)$ converges to the direction of the j th CLV when both t and $(T - t)$ become large. Notice that, although $\mathbf{V}(t)$ has different expressions on different segments, its columns are continuous across all segments.

We discuss the inner product we implicitly used in defining the orthogonality of the \mathbf{Q} matrices. The inner product between two flow fields at an instant is a summation of inner products between different components. For consistency of units, we normalize each component by free-stream properties. Moreover, we should prescribe a geometric metric for inner products between the same components. This metric should have bounded total volume, so that inner products are finite; it should also put more weight closer to the cylinder, so that FD-NILSS has better accuracy for the surface objectives we use later. A natural selection of such a metric is simply one where all mesh cells are equally weighted, and the entire space outside the mesh has weight zero. More specifically, with subscript l denoting different components, and k denoting different cells of the mesh, for $\mathbf{v}, \mathbf{w} \in X = \mathbb{R}^m$, we define inner products as

$$\mathbf{v}^T \mathbf{w} := \frac{1}{m'} \sum_{k=1}^{m'} \mathbf{v}_k^T \mathbf{w}_k, \quad \text{where } \mathbf{v}_k^T \mathbf{w}_k := \frac{1}{3} \left(\frac{v_{k1}w_{k1}}{\rho_0^2} + \frac{\sum_{l=2}^4 v_{kl}w_{kl}}{\rho_0^2 U_0^2} + \frac{v_{k5}w_{k5}}{\rho_0^2 E_0^2} \right). \tag{3.7}$$

Here m' is the number of cells. We order the five components of states by $\rho, \rho U, \rho E$, and normalize all components of ρU by ρU_0 . Under our definition, if every point in a flow field has the same state as the free stream, then this flow field has unit norm.

The convergence of CLV in the above algorithm is in terms of the distance between one-dimensional subspaces. Since a CLV is a homogeneous tangent solution, we can multiply it by any factor, and still get a homogeneous tangent solution whose norm behaves like an exponential function with the same LE, which, by definition, is still the same CLV. In other words, it is only the directions of CLVs that are meaningful, but not the magnitudes. Hence we can normalize CLVs by any factor we like. In this paper, we normalize CLVs such that their maximal values are 1.0.

In this paper, we use finite differences to approximate tangent solutions. To approximate a homogeneous solution w with initial condition w^0 , we use the definition $w = du/d\phi \approx \Delta u/\Delta\phi$, where ϕ controls the perturbation in the initial condition in the direction of w^0 . More specifically, we compute a perturbed primal solution u^w with a perturbed initial condition $u^0 + w^0 \Delta\phi$, and the approximation for w is $w \approx (u^w - u)/\Delta\phi$, where u is the base trajectory.

3.4. Results of Lyapunov exponents

The time-step size used for computing LEs and CLVs is the same as in the numerical simulation, that is, $\Delta t = 10^{-8} = 1.32 \times 10^{-3} t_0$. The number of homogeneous tangent solutions we compute is $M = 40$. Each segment has 200 time steps, which gives $\Delta T = 2 \times 10^{-6} = 0.264 t_0$. The total number of time segments is $K = 800$, hence the time length of the entire trajectory is $T = 1.6 \times 10^{-3} = 211 t_0$.

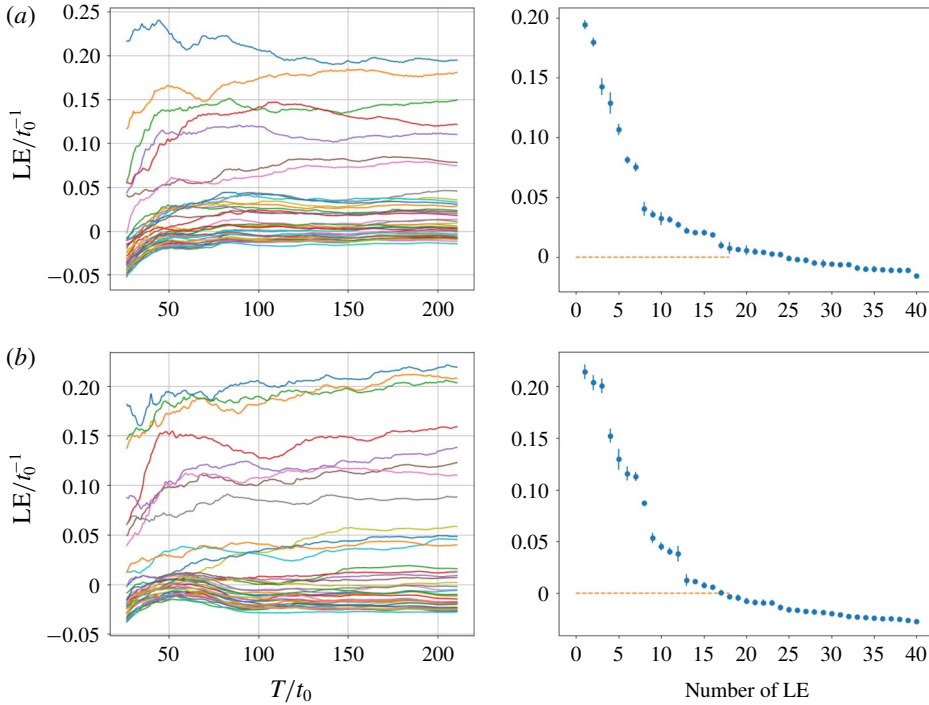


FIGURE 5. (Colour online) Lyapunov exponents (LE) normalized by t_0^{-1} . Left: convergence history of the first 40 different LEs. Right: confidence intervals of LEs. (a) Results on the coarser mesh. The largest LE is $0.20t_0^{-1}$, meaning in one flow-through time t_0 , the norm of the first CLV becomes $e^{0.20} = 1.22$ times larger. (b) Results on the finer mesh. The largest LE is $0.21t_0^{-1}$.

With the above numerical settings, the convergence history of the first 40 LEs is shown on the left of figure 5. As explained in Ni & Wang (2017), the confidence interval of an LE is estimated by the smallest interval which bounds the history of the LE and whose size shrinks as $T^{-0.5}$. The first 40 LEs and their confidence intervals are shown on the right of figure 5.

We can see that LE results depend on the meshes, as similarly discussed by Fernandez & Wang (2017b). In figure 5, although the general shape of the LE spectrum looks similar, a closer look will show that the finer mesh has larger but fewer positive LEs. As a result, the selection of neutral CLVs is ambiguous and changes due to meshes.

However, we will show that there are several important physical phenomena which, although they are logically consequential to the Lyapunov analysis results, exist robustly regardless of the detailed value of the LE spectrum. This agrees with the chaotic hypothesis (Gallavotti 2008), which says that many tools logically consequential to the uniform hyperbolicity assumption are in fact valid for a larger class of dynamical systems. More specifically, those phenomena are as follows:

- (1) the dimension of the attractor is much lower than the dimension of the system;
- (2) angles between far-apart CLVs are large, and, in particular, CLVs are active at different locations;

(3) shadowing directions exist, and reflect the sensitivities of statistics to parameter perturbations.

The focus of this paper will be the last two phenomena. Moreover, for engineering interests, we will show that FD-NILSS can compute shadowing directions and, further, the sensitivities with reasonable cost.

We first show that the dimension of the attractor is much lower than the dimension of the system. Intuitively, to remain its invariant hyper-volume, the attractor manifold should contain enough contracting (stable) CLVs to balance out the expansion caused by unstable CLVs. More specifically, Frederickson *et al.* (1983) estimated the attractor dimension by the so-called Lyapunov dimension:

$$D_\lambda = N + \frac{1}{|\lambda_{N+1}|} \sum_{j=1}^{j=N} \lambda_j \in [N, N + 1], \tag{3.8}$$

where N is such that $\lambda_1 + \lambda_2 + \dots + \lambda_N > 0$ and $\lambda_1 + \lambda_2 + \dots + \lambda_{N+1} < 0$. We want to give a bound to N on the finer mesh. To do this, we remind readers that the algorithm we are using computes LEs by descending order, so we have $\lambda_j \leq \lambda_{40}$ for all $j \geq 40$. Using the first condition of N , we have

$$\sum_{j=1}^{40} \lambda_j + (N - 40)\lambda_{40} = \sum_{j=1}^{40} \lambda_j + \sum_{j=41}^N \lambda_{40} \geq \sum_{j=1}^{40} \lambda_j + \sum_{j=41}^N \lambda_j > 0, \tag{3.9}$$

which further leads to

$$N \leq 40 + \frac{\sum_{j=1}^{40} \lambda_j}{|\lambda_{40}|} = 40 + \frac{1.031}{0.027} = 78. \tag{3.10}$$

On the other hand, since $\lambda_1 + \dots + \lambda_{40} > 0$, by the second property of N , we know $N > 40$. Hence the Lyapunov dimension $41 \leq D_\lambda \leq 79$. In other words, the chaotic dynamic of our flow problem on the finer mesh can be attributed to the interaction of fewer than 79 degrees of freedom. Using the same method, we can estimate that the attractor dimension on the coarser mesh is smaller than 109.

3.5. Results of covariant Lyapunov vectors and hyperbolicity

In figure 6(a,c), we plot the histogram of angles between all pairs of the first 40 CLVs, which are perturbations whose norm grows exponentially, as defined in (3.4). Notice the range of angles is $[0^\circ, 90^\circ]$, since the angle between the directions of two CLVs is the angle between subspaces of dimension one. For the finer mesh, the smallest angle is 16.4° . However, for the coarser mesh, the smallest angle is only 5.8° , almost equal the threshold value of 5.7° . This indicates that the smallest angle and hence the second assumption in uniform hyperbolicity is not a robust criterion to check for fluid systems, since it is plausible that the smallest angle falls below the threshold for another mesh or a longer time span.

The above dependences on meshes call for a closer examination to find some robust physical phenomena not depending on meshes. A deeper look into CLV results finds that tangencies happen only among adjacent CLVs: this result is similar to that




FIGURE 6. Histogram of angles between CLVs on segments 350–449, with the entire trajectory of length $T = 211t_0$ partitioned into $K = 800$ segments. The total area of the histogram is normalized to one. (a) Coarser mesh, all CLVs, smallest angle is 5.8° . (b) Coarser mesh, CLV indices more than five apart, smallest angle is 29.2° . (c) Finer mesh, all CLVs, smallest angle is 16.4° . (d) Finer mesh, CLV indices more than five apart, smallest angle is 40.8° .

obtained by Xu & Paul (2016). Hence we plot the angles between CLVs whose indices are more than five apart. For example, for the 20th CLV, we consider its angles with the 1st to 15th CLVs and the 25th to 40th CLVs. This result is plotted in figure 6(b,d).

As we can see, when the exponents of CLVs are further apart, their angles become larger: this phenomenon is related to hyperbolicity but not formally named yet. So, on one hand, what we observe violates the second assumption of uniform hyperbolicity since occasional tangencies may fail the threshold. But, on the other hand, our observation is even stronger than the uniform hyperbolicity, because it specifies a trend not suggested by the definition of uniform hyperbolicity.

To further illustrate, we plot minimal and averaged angles between all pairs of CLVs in figure 7. Not surprisingly, most tangencies happen between adjacent CLVs, and the angles between CLVs get larger for CLVs whose indices are farther away. In particular, on the coarser mesh, the first seven CLVs are closely coupled, whereas the first several CLVs on the finer mesh are less tangent to each other. In general, tangencies among CLVs are rare events for both meshes, but even less frequent for the finer mesh. Once again, the rareness of the tangencies indicates that smallest angle and hence uniform hyperbolicity is not a robust criterion; on the other hand, the fact that angles between far-apart CLVs are large is robust to meshes, and this may lead to some definition of a new form of hyperbolicity.

In fact, for two CLVs whose exponents are apart, their angles are so large that we can even observe the difference in active areas. To illustrate, for the coarser mesh, we plot the 1st and 40th CLVs in figure 8. For the finer mesh, we plot the 1st, 5th, 17th and 40th CLVs in figure 9. Movies of the 1st, 5th, 17th and 28th CLVs on the finer mesh are movie 1 to movie 4 of the supplementary material, available at <https://doi.org/10.1017/jfm.2018.986>. We remind readers that only the direction of




FIGURE 7. Minimal and averaged angles between all pairs of CLVs. The minimum and average are taken on segments 350–449, with the entire trajectory of length $T = 211t_0$ partitioned into $K = 800$ segments. (a) Coarser mesh, minimal angles between all pairs of CLVs. (b) Coarser mesh, averaged angles between all pairs of CLVs. (c) Finer mesh, minimal angles between all pairs of CLVs. (d) Finer mesh, averaged angles between all pairs of CLVs.

CLVs, that is, only the relative flow pattern of CLVs, is meaningful. Since our primal solution lives in the function space $X \subset \mathbb{R}^m$ whose tangent space is also \mathbb{R}^m , for any time t , a CLV $\xi(t)$ lives in a function space locally the same as the primal solution $\mathbf{u}(t)$. In our 3D flow problem, this means our CLVs also have ρ , $\rho\mathbf{U}$, ρE components, and we plot $|\rho\mathbf{U}|$.

As we can see in both figures 8 and 9, the general trend is that the active areas of CLVs move downstream as exponents become smaller. Intuitively, in boundary layers and near wakes, chaos is generated; perturbations in these places tend to grow, resulting in unstable CLVs. On the other hand, in far wakes, chaos is dissipated; perturbations in this region tend to decay, resulting in stable CLVs. Further, we can rationally conjecture that the location of neutral CLVs is in between stable and unstable CLVs, that is, their active area is the entire wake. This conjecture is also backed by observing the plot of a neutral CLV in figure 9(c).

We conjecture that perturbations in the stagnant area in front of the cylinder will be dissipated, resulting in stable CLVs. This is because the fluid motion in front of the cylinder is almost determined, that is, fluid here should decelerate into stagnancy. If we make a perturbation to this area, the fluid motion in the long term will not be




FIGURE 8. CLVs computed on the coarser mesh on a trajectory of time length $T = 211t_0$: (a) 1st CLV; (b) 40th CLV; and (c) primal flow field. Left: snapshots at $T/2$. Right: averaged over time span $[95t_0, 116t_0]$. Plotted by the magnitude of ρU component, and normalized such that the largest value is 1.00. The primal flow at the bottom is normalized by $\rho_0 U_0$.

changed. Hence, such a perturbation could not last long, and the corresponding LE should be negative. Our conjecture is also partly evidenced by the fact that the first 40 CLVs are not active in this area.

We further conjecture that perturbations in the free stream are also dissipated, resulting in stable CLVs. We provide two arguments supporting this conjecture. First, assume we make a perturbation in the free stream. Then, after changing reference frame, this is equivalent to making a small perturbation in a stagnant fluid domain; the Euclidean norm of such a perturbation decays due to the dissipation in the Navier–Stokes equations. Second, since our metric, induced by (3.7), is compactly supported, a perturbation in the free stream will be transported outside the mesh where the volume weight is zero. Hence under our norm, perturbation in the free stream decays also because of the decay in the volume weight. In fact, because our norm of such perturbation decays to zero, the corresponding LE should be very negative.

To summarize, we make the following conjecture about active areas of CLVs in subsonic open flows. In subsonic open flows, unstable CLVs are active in the instability-generating area such as boundary layers and near wakes; the neutral CLV occupies similar locations as the wake; stable CLVs with moderately negative LEs are active in the dissipative area such as far wakes; stable CLVs active in the stagnant area ahead of the cylinder and in the free stream have very negative LEs. Our conjecture also indicates that subsonic open flows have large angles between far-apart CLVs.

Our conjecture may provide insights into the question about how many CLVs are unstable for fluid systems. Owing to our arguments, for open flows, or more generally fluid systems with non-turbulent regions, there should be a large part of CLVs that are stable; hence unstable CLVs should be a small fraction of all CLVs. At least for these fluid systems, more analysis should be focused on unstable CLVs since they are the major contributor for turbulence. Additionally, there should be more numerical




FIGURE 9. CLV computed on the finer mesh, using same settings as the coarser mesh: (a) 1st CLV; (b) 5th CLV; (c) 17th CLV; (d) 40th CLV; and (e) primal flow field.

methods developed which explicitly exploit the unstable structures, such as NILSS and NILSAS.

We find more structures in hyperbolicity, but our system certainly does not satisfy the uniform hyperbolicity under which shadowing methods were mathematically developed. Yet there is still hope, as the chaotic hypothesis suggested that, although the logical bridge of uniform hyperbolicity breaks down, we may find tools across the river still useful. With this hope, we go on to examine shadowing results on our fluid system. Should we find shadowing methods still valid, we will have more confidence in its generality.

4. Shadowing directions and sensitivities

4.1. Definition of shadowing directions

The shadowing solution v^∞ is an inhomogeneous tangent solution whose orthogonal projection perpendicular to the trajectory, $v^{\infty\perp}$, is uniformly bounded on an infinitely long trajectory. Here $(\cdot)^\perp$ is defined as

$$p^\perp = p - \frac{f^T p}{f^T f} f, \quad (4.1)$$




FIGURE 10. Shadowing directions. The base trajectory has parameter s , the shadowing trajectory has parameter $s + \delta s$. The first-order approximation of the perpendicular distance is $\delta \mathbf{u}^\perp \approx \mathbf{v}^{\infty\perp} \delta s$, where $\mathbf{v}^{\infty\perp}$ is the shadowing direction.

where $\mathbf{p} \in \mathbb{R}^m$ is an arbitrary vector, \mathbf{f} is the trajectory direction defined in (3.1), and the inner product was defined in (3.7). We remark here that, for a continuous dynamical system, the definition of shadowing direction involves a reparametrization of time (Pilyugin 1999, Definition 1.26), or equivalently a ‘time dilation’ (Wang *et al.* 2014), or projecting the tangent solutions onto the subspace perpendicular to \mathbf{f} (Ni & Wang 2017). The $\mathbf{v}^{\infty\perp}$ is the projection of \mathbf{v}^∞ given by the third approach, and $\mathbf{v}^{\infty\perp}$ is bounded, but \mathbf{v}^∞ is not. For convenience, we also call \mathbf{v}^\perp the shadowing direction.

The shadowing direction has the following physical meaning. The definition of inhomogeneous tangent solution $\mathbf{v} = d\mathbf{u}/ds$ tells us that, with perturbed parameter $s + \delta s$, there is a new trajectory $\mathbf{u} + \delta \mathbf{u}$ such that $\delta \mathbf{u} \approx \mathbf{v}^\infty \delta s$. Now the perpendicular distance between the new and the old trajectory is $\delta \mathbf{u}^\perp \approx \mathbf{v}^{\infty\perp} \delta s$; together with the boundedness of $\mathbf{v}^{\infty\perp}$, we see that the new trajectory remains perpendicularly close to the old trajectory for a long time. This new trajectory will be referred to as the shadowing trajectory, and the old one as the base trajectory. This intuition of the shadowing direction is shown in figure 10.

The initial condition of the shadowing direction is not known *a priori*, or equivalently, we do not specify how system parameters would affect initial conditions. In fact, as we shall see, although the shadowing direction solves that same ODE as in (3.3), its determining condition is no longer the initial condition; rather, to determine a shadowing direction, we minimize the L^2 norm of the orthogonal projection of an inhomogeneous tangent solution.

4.2. Sensitivity analysis of long-time averages via shadowing directions

For chaotic dynamical systems such as the 3D flow problem in this paper, the output of the system, such as the drag or lift, is typically aperiodic. As often happens in engineering, the objective is the long-time average $\langle J \rangle_\infty$ of an instantaneous quantity $J(\mathbf{u}, s) : \mathbb{R}^m \times \mathbb{R} \rightarrow \mathbb{R}$. More specifically, we define

$$\langle J \rangle_\infty := \lim_{T \rightarrow \infty} \langle J \rangle_T, \quad \text{where } \langle J \rangle_T := \frac{1}{T} \int_0^T J(\mathbf{u}, s) dt. \tag{4.2}$$

Typically $\langle J \rangle_\infty$ is approximated by a $\langle J \rangle_T$ with a large T . The sensitivity of the objective with respect to system parameters, $d\langle J \rangle_\infty/ds$, is of engineering interest. In this subsection we show how to compute the sensitivity using the shadowing direction.

Assume we have found, on a finitely long trajectory, an approximation $\mathbf{v} \approx \mathbf{v}^\infty$, where \mathbf{v} is also an inhomogeneous tangent solution. We first define the time dilation term (Wang *et al.* 2014) η :

$$\frac{d\mathbf{v}^\perp}{dt} = \partial_u \mathbf{f} \mathbf{v}^\perp + \partial_s \mathbf{f} + \eta \mathbf{f}. \quad (4.3)$$

Intuitively, η describes the relative time spent on the shadowing trajectory in comparison with the base trajectory. In figure 10, if the shadowing trajectory takes less time to travel from point C to D than the base trajectory from A to B, then (4.3) gives $\eta < 0$. On the other hand, if the shadowing trajectory moves slower, then $\eta > 0$.

Once we get \mathbf{v} and η , the sensitivity can be computed by

$$\frac{d\langle J \rangle_\infty}{ds} \approx \frac{1}{T} \int_0^T [\partial_u J \mathbf{v}^\perp + \partial_s J + \eta(J - \langle J \rangle_T)] dt, \quad (4.4)$$

where $\langle J \rangle_T$ is defined in (4.2). Intuitively, the first term inside the integration in (4.4) describes the contribution by the perpendicular distance between the shadowing trajectory and the base trajectory. The second term is due to the fact that the function J may explicitly depend on s . The last term involves η , and it accounts for the fact that the average is taken with respect to time: if in some small region the shadowing trajectory moves faster than the base trajectory, then the time spent is shorter, hence the contribution from this region in the time average should also be smaller.

Another formula for the sensitivity is easier for computer programming:

$$\frac{d\langle J \rangle_\infty}{ds} \approx \frac{1}{T} \left[\int_0^T (\partial_u J \mathbf{v} + \partial_s J) dt + \xi|_0^T \langle J \rangle_T - (\xi J)|_0^T \right], \quad (4.5)$$

where the time difference term, ξ , is a time-dependent scalar function such that

$$\xi \mathbf{f} = \mathbf{v} - \mathbf{v}^\perp. \quad (4.6)$$

Thus ξ is easier to compute than η , since its definition does not involve a time derivative. Note that in (4.5) we use \mathbf{v} instead of its projection \mathbf{v}^\perp . The derivation of (4.4), (4.5) and their relation can be found in the appendix of Ni & Wang (2017).

4.3. The NILSS algorithm for computing shadowing directions

We describe the non-intrusive least-squares shadowing (NILSS) algorithm in this subsection. As in § 3.3, for $i = 0, \dots, K - 1$, we define the i th time segment as $[t_i, t_{i+1}]$, with $t_i = i\Delta T$. In the algorithm presented below, for quantities defined on the entire segments, such as \mathbf{u}_i , \mathbf{W}_i , \mathbf{v}_i^* , \mathbf{C}_i and \mathbf{d}_i , we use the same subscripts as the segments they are defined on. For quantities defined only at interfaces between segments, such as \mathbf{Q}_i , \mathbf{R}_i and \mathbf{b}_i , we use the same subscripts as the time points they are defined at. Again, we use finite difference results to approximate tangent solutions: such a variant is called the finite difference non-intrusive least-squares shadowing (FD-NILSS) algorithm, whose details are given in Ni *et al.* (2018).

To start with, we should prescribe: (1) number of homogeneous tangent solutions, M , which must be larger than the number of unstable CLVs, m_{us} ; (2) length of each time segment, ΔT ; and (3) number of time segments, K . Consequently, the time length of the entire trajectory, $T = K\Delta T$, is also determined. Then, the NILSS algorithm is given by the following procedure.

1. Generate initial conditions for primal solutions, homogeneous tangent solutions and inhomogeneous tangent solutions.
 - 1.1. Compute the primal solution of (3.1) for sufficiently long time so that the trajectory lands on the attractor, then set $t = 0$, and set the initial condition of the primal system, $\mathbf{u}_0(0)$.
 - 1.2. Randomly generate an $m \times M$ orthogonal matrix $\mathbf{Q}_0 = [\mathbf{q}_{01}, \dots, \mathbf{q}_{0M}]$ whose column vectors are orthogonal to $\mathbf{f}(t = 0)$. This \mathbf{Q}_0 will be used as initial conditions for homogeneous tangent solutions.
 - 1.3. Set the initial condition of the particular inhomogeneous tangent solution $\mathbf{v}_0^*(0) = 0$.
2. For $i = 0$ to $K - 1$, on segment i , where $t \in [t_i, t_{i+1}]$, do:
 - 2.1. Compute the primal solution $\mathbf{u}_i(t)$ from t_i to t_{i+1} .
 - 2.2. Compute homogeneous tangent solutions $\mathbf{W}_i(t) = [\mathbf{w}_{i1}(t), \dots, \mathbf{w}_{iM}(t)]$.
 - 2.2.1. For each homogeneous tangent solution \mathbf{w}_{ij} , $j = 1, \dots, M$, starting from initial condition $\mathbf{w}_{ij}(t_i) = \mathbf{q}_{ij}$, integrate (3.2) from t_i to t_{i+1} .
 - 2.2.2. Compute orthogonal projection $\mathbf{W}_i^\perp(t) = [\mathbf{w}_{i1}^\perp(t), \dots, \mathbf{w}_{iM}^\perp(t)]$ via (4.1).
 - 2.2.3. Compute and store the covariant matrix on segment i :

$$\mathbf{C}_i = \int_{t_i}^{t_{i+1}} (\mathbf{W}_i^\perp)^\top \mathbf{W}_i^\perp dt. \tag{4.7}$$
 - 2.2.4. Perform QR factorization: $\mathbf{W}_i^\perp(t_{i+1}) = \mathbf{Q}_{i+1} \mathbf{R}_{i+1}$, where $\mathbf{Q}_{i+1} = [\mathbf{q}_{i+1,1}, \dots, \mathbf{q}_{i+1,M}]$.
 - 2.3. Compute the particular inhomogeneous tangent solution $\mathbf{v}_i^*(t)$.
 - 2.3.1. Starting from initial condition $\mathbf{v}_i^*(t_i)$, integrate the inhomogeneous equation (3.3) from t_i to t_{i+1} .
 - 2.3.2. Compute the orthogonal projection $\mathbf{v}_i^{*\perp}(t)$ via (4.1).
 - 2.3.3. Compute and store

$$\mathbf{d}_i = \int_{t_i}^{t_{i+1}} (\mathbf{W}_i^\perp)^\top \mathbf{v}_i^{*\perp} dt. \tag{4.8}$$
 - 2.3.4. Orthogonalize $\mathbf{v}_i^{*\perp}(t_{i+1})$ with respect to $\mathbf{W}_{i+1}^\perp(t_{i+1}) = \mathbf{Q}_{i+1}$ to obtain the initial condition of the next time segment:

$$\mathbf{v}_{i+1}^*(t_{i+1}) = \mathbf{v}_i^{*\perp}(t_{i+1}) - \mathbf{Q}_{i+1} \mathbf{b}_{i+1}, \tag{4.9}$$

compute and store:

$$\mathbf{b}_{i+1} = \mathbf{Q}_{i+1}^\top \mathbf{v}_i^{*\perp}(t_{i+1}). \tag{4.10}$$

3. Solve the NILSS problem:

$$\left. \begin{aligned} \min_{\{\mathbf{a}_i\}} \sum_{i=0}^{K-1} \frac{1}{2} \mathbf{a}_i^\top \mathbf{C}_i \mathbf{a}_i + \mathbf{d}_i^\top \mathbf{a}_i \\ \text{s.t. } \mathbf{a}_i = \mathbf{R}_i \mathbf{a}_{i-1} + \mathbf{b}_i, \quad i = 1, \dots, K - 1. \end{aligned} \right\} \tag{4.11}$$

This is a least-squares problem with arguments $\{\mathbf{a}_i\}_{i=0}^{K-1}$, where $\mathbf{a}_i \in \mathbb{R}^M$ for each i .

4. Compute v_i within each time segment $t \in [t_i, t_{i+1}]$:

$$\mathbf{v}_i(t) = \mathbf{v}_i^*(t) + \mathbf{W}_i(t)\mathbf{a}_i. \quad (4.12)$$

The orthogonal projections of \mathbf{v}_i on each segment, $\{\mathbf{v}_i^\perp\}_{i=0}^{K-1}$, converge to the orthogonal projection of the shadowing direction, $\mathbf{v}^{\infty\perp}$, as both t and $T - t$ become large. In other words, we have smaller error in the middle section of a long time span.

5. Compute ξ_i at the end of each segment:

$$\xi_i(t_{i+1}) = \frac{(\mathbf{v}_i(t_{i+1}))^T \mathbf{f}(\mathbf{u}(t_{i+1}))}{\mathbf{f}(\mathbf{u}(t_{i+1}))^T \mathbf{f}(\mathbf{u}(t_{i+1}))}. \quad (4.13)$$

6. The derivative can be computed by

$$\frac{d\langle J \rangle_\infty}{ds} \approx \frac{1}{T} \sum_{i=0}^{K-1} \left[\int_{t_i}^{t_{i+1}} (\partial_u J \mathbf{v}_i + \partial_s J) dt + \xi_i(t_{i+1}) (\langle J \rangle_T - J(t_{i+1})) \right]. \quad (4.14)$$

Readers should note that, although at the beginning of each segment, \mathbf{v}_i is perpendicular to $\mathbf{f}(t_i)$, typically \mathbf{v}_i is not perpendicular to $\mathbf{f}(t_i)$ for the rest of the segment. In NILSS, we typically use perpendicularly projected vectors, such as $\mathbf{v}_i^{*\perp}$, \mathbf{W}_i^\perp , \mathbf{v}_i^\perp , for constructing the minimization problem in (4.11) to compute the shadowing directions. Unprojected vectors, such as \mathbf{v}_i , are typically used for computing sensitivities by (4.14). As shown in the appendix of Ni & Wang (2017), we can construct a continuous \mathbf{v} , and its orthogonal projection \mathbf{v}^\perp confined on each segment yields \mathbf{v}_i^\perp . However, \mathbf{v}_i are not confinements of \mathbf{v} , and \mathbf{v}_i are not necessarily continuous across segments. In fact, in (4.11), the minimization is minimizing $\|\mathbf{v}^\perp\|_{L^2}$, while the constraint prescribes that \mathbf{v}^\perp is continuous across all time segments.

Intuitively, at $t = 0$, the difference between $\mathbf{v}^{\infty\perp}$ and \mathbf{v}^\perp contains unstable CLVs; since $\mathbf{v}^{\infty\perp}$ is bounded, as time evolves, the exponential growth of unstable CLVs increases both $\|\mathbf{v}^{\infty\perp} - \mathbf{v}^\perp\|_{L^2}$ and $\|\mathbf{v}^{*\perp}\|_{L^2}$. Also, as time evolves, the span of M homogeneous tangent solutions approximates the span of the first M CLVs, hence $\mathbf{v}^\perp = \mathbf{v}^{*\perp} + \mathbf{W}^\perp \mathbf{a}$ allows us to subtract unstable CLVs from \mathbf{v}^\perp . Now minimizing $\|\mathbf{v}^\perp\|_{L^2}$ removes unstable CLVs from both \mathbf{v}^\perp and $\mathbf{v}^{\infty\perp} - \mathbf{v}^\perp$, thus making $\mathbf{v}_i^\perp \approx \mathbf{v}^{\infty\perp}$. Since the number of unstable CLVs is typically much smaller than the dimension of the dynamical system, NILSS is computationally efficient. A more detailed explanation of NILSS is contained in Ni & Wang (2017).

4.4. Results of shadowing directions

This subsection discusses the shadowing directions computed by the FD-NILSS algorithm. We remind readers that shadowing solutions depend only on the choice of system parameters but not on objectives. In particular, this subsection shows two shadowing solutions with respect to two system parameters: (1) x -directional free-stream velocity U , normalized by U_0 ; and (2) rotation speed of the cylinder ω , measured in cycles per time unit, normalized by ω_0 .

In FD-NILSS we set time-step size 1×10^{-8} , and 200 steps in each time segment; hence the segment length is $\Delta T = 2 \times 10^{-6} = 0.259t_0$. These are the same as what we used for computing LEs and CLVs. Different from those used for computing LEs and CLVs, we set the number of homogeneous tangents $M = 30$ and the number of




FIGURE 11. (Colour online) Norms of shadowing directions on the finer mesh, measured from $66t_0$ to $93t_0$, computed on a trajectory of length $T = 158t_0$: (a) with U as parameter; and (b) with ω as parameter.

segments $K = 600$; hence the time length of the entire trajectory is $T = 1.2 \times 10^{-3} = 158.4t_0$.

We first plot the norm of shadowing directions on the finer mesh in figure 11, where the norm is induced by the inner product defined in (3.7). Since $v^\perp = du^\perp/ds$, with du^\perp already normalized by free-stream constants in the definition of inner product, we further need to normalize shadowing directions by s_0^{-1} . Here s_0 is the unit parameter, that is, U_0 or ω_0 .

By inspecting figure 11, we can verify the uniform boundedness of shadowing directions, as both $\|v^\perp\|$ do not change much from $66t_0$ to $93t_0$. For comparison we may think of the first CLV, whose norm would grow $e^{0.21 \times (93-66)} = 290$ times larger within the same time span.

Orthogonal projections of shadowing directions, v^\perp , are plotted in figure 12. Movies of v^\perp on the finer mesh are movie 5 and movie 6 of the supplementary material. As for CLVs, for any time t , $v^\perp(t)$ lives in the same function space as the primal solution $u(t)$, hence v^\perp also has ρ , ρU and ρE components. In figure 12, for each parameter, we plot the $|\rho U|$ field of $v^\perp(t)$. Unlike CLVs, both the direction and the magnitude of shadowing directions are meaningful, since shadowing directions reflect not only in what direction flow fields will change due to perturbations in system parameters s , but also how large the change will be. Hence, we normalize shadowing directions by free-stream constants and unit parameters to preserve information about magnitudes.

If we perturb system parameters by ΔU , the shadowing solution is a flow field with free-stream property $\rho_0(U_0 + \Delta U)$. Since v^\perp points from the base solution to the shadowing solution, we expect the free-stream area of v^\perp to be $\rho_0 \Delta U / \Delta U = \rho_0$. The above intuition is verified on the left of figure 12(a): with U as the system parameter, the free-stream area of the ρU component of the shadowing direction v^\perp has magnitude 1, since we are normalizing by ρ_0 . On the other hand, perturbing rotation speed ω should have little effect on the free-stream area: this is also verified in figure 12(b), where v^\perp has small magnitude in the free-stream area.

We observe by comparing the left of figure 12(a) and figure 12(b) that, in the wake area, both v^\perp have similar magnitude. This fact is further confirmed by the fact that in figure 11 both $\|v^\perp\|$ are similarly large after normalization. This fits our intuition that, if we perturb cylinder rotation by $\Delta \omega = 0.01\omega_0$, then the circumferential speed is changed by $0.01\omega_0 D/2 = 0.01U_0/2$: this extra rotation should have similar magnitude of impact on the wake as changing the free-stream velocity by $\Delta U = 0.01U_0$.

We compare the shadowing directions with the primal flow field in figure 12. On one hand, shadowing directions are correlated to the primal flow field, since the peak




FIGURE 12. Magnitude of the ρU component of the shadowing direction $\mathbf{v}^\perp = d\mathbf{u}^\perp/ds$ on the finer mesh, which is computed by FD-NILSS on a trajectory of time length $T = 158t_0$: (a) U as the parameter, normalized by $\rho_0 U_0/U_0 = \rho_0$; (b) ω as the parameter, normalized by $\rho_0 U_0/\omega_0$; and (c) primal flow field, normalized by $\rho_0 U_0$. Left: snapshots at $0.5T$; right: averaged from $66t_0$ to $93t_0$. Notice that the left of (a) and (b) are coloured by log scale.

values of shadowing directions are typically located close to the vortices, indicating that vortices structures are the most sensitive to parameter perturbations.

On the other hand, because shadowing directions are perturbations, they are very different from the primal flow. The primal flow has several structures that are insensitive to parameter perturbations, whereas parameter perturbations may bring new structures to the primal flow. For example, when ω is parameter, the stagnant area in front of the cylinder in primal flow changes to a smooth halo around the cylinder in the shadowing directions, due to the fact that perturbing ω leads to a uniform acceleration around the cylinder. Moreover, the shadowing direction has a larger active area than the primal flow. An extreme example is the left of figure 12(a), where the entire free stream of shadowing directions has value one, as we discussed. Indeed, the convolution formula of shadowing directions shows that shadowing directions are affected by both stable and unstable CLVs (Ni 2018b), the union of whose active areas cover the entire flow field.

4.5. Results of sensitivities

With shadowing directions, we can compute sensitivities of long-time-averaged objectives from (4.13) and (4.14) with little extra cost. In fact, as we shall see, using only a trajectory of time length $106t_0$ should be enough to compute accurate sensitivities. In this subsection, we investigate the effect of U on the averaged drag force $\langle D_r \rangle$. We will normalize $\langle D_r \rangle$ by $F_0 = 0.5\rho U_0^2 DZ = 8.031 \times 10^{-5}$. The objective for system parameter ω is averaged lift $\langle L \rangle$, which is also normalized by F_0 .

We plot the history of the sensitivities computed via shadowing directions in figure 13. For each T , we run NILSS from the same initial condition of primal and




FIGURE 13. History plots of sensitivities computed by FD-NILSS: (a) sensitivity of the drag with respect to U , and (b) sensitivity of the lift with respect to ω .

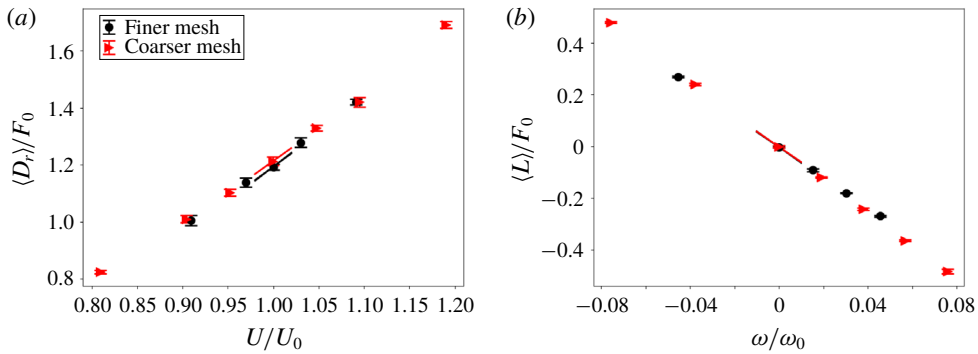


FIGURE 14. (Colour online) Sensitivities for different choices of parameters and objectives: (a) drag versus U , and (b) lift versus ω . Vertical bars indicate 95% confidence intervals of the averaged objectives. The straight lines through the middle data points indicate the sensitivities computed by FD-NILSS.

tangent solutions to T , solve the shadowing direction, using which we compute the sensitivity corresponding to T . The uncertainty of the final sensitivity is estimated by the smallest interval which bounds the history of the sensitivity and whose size shrinks as $T^{-0.5}$. We can see that for both meshes the uncertainty of sensitivities is below 10%.

We explain how to compute the uncertainty of long-time-averaged objectives, which is approximated by $\langle J \rangle_{T'}$, the averaged instantaneous objective over $T' = 7.68 \times 10^{-3} = 1014t_0$. To get the uncertainty due to taking the average over finite time, we divide the history of $J(t)$ into five equally long parts. We denote the objectives averaged over each of the five parts by J_1, \dots, J_5 , whose corrected sample standard deviation is denoted by σ' . We assume that the standard deviation of $\langle J \rangle_{T'}$ is proportional to $T'^{-0.5}$, so we use $\sigma = \sigma' / \sqrt{5}$ as the standard deviation of $\langle J \rangle_{T'}$. We further assume $\pm 2\sigma$ yields the 95% confidence interval, which is indicated by the vertical bar in figure 14.

The sensitivities computed via shadowing directions are shown in figure 14. For both meshes our sensitivities correctly reflect the trend between averaged objectives and system parameters. This confirms the physical meaning of the shadowing direction: it reveals the sensitivity of all physical phenomena in this flow field

with respect to perturbations in parameters, in such a way that the average of the sensitivity is the sensitivity of the average.

The cost of FD-NILSS for computing a sensitivity is mainly in integrating the primal solution over $400 \times 200 \times 32 = 2.6 \times 10^6$ time steps. Here 400 is the number of segments, 200 is the number of time steps in each segment, and 32 is the number of primal solutions computed. In FD-NILSS we need one inhomogeneous tangent and 30 homogeneous tangent solutions, and each tangent solution is approximated by the difference between a perturbed solution and the same base solution: that is 32 primal solutions in total. On the other hand, a sensitivity can be revealed by some regression among the five pairs of objectives and parameters in figure 14, the total cost of which is 3.8×10^6 steps of primal simulation: this cost is similar to that of FD-NILSS.

NILSS can be further accelerated when we have a tangent solver, by taking advantage of the fact that all adjoint solutions use the same Jacobian $\partial_u f$. That is, we can integrate all tangent solutions simultaneously without repeatedly loading $\partial_u f$ into the computer CPU, which is the most time-consuming procedure in the numerical integration. In other words, we can perform, at each time step, one matrix–matrix product, where the second matrix is composed of several tangent solutions, rather than performing several matrix–vector products, where we need to reload the matrix for each product. (This idea was brought up during private discussions with Pablo Fernandez, who co-authored FD-NILSS.)

Additionally, as discussed in Ni & Wang (2017) and Ni *et al.* (2018), FD-NILSS and NILSS have smaller marginal cost for new parameters; for cases with more parameters than objectives or cases when we have only adjoint solvers, the non-intrusive least-squares adjoint shadowing (NILSAS) method (Ni 2018a) can compute the gradient of one objective with respect to many parameters in one run.

From the above discussions, together with the implication of the conjecture in § 3.5 that there should be a small fraction of unstable CLVs for open flows, we believe that NILSS and NILSAS are competitive in terms of efficiency among all sensitivity algorithms.

As we can see, although our system is not uniform hyperbolic, the sensitivities computed by shadowing methods are still valid for both meshes. We thus add one more piece of evidence to the chaotic hypothesis: that is, although shadowing methods are logically consequential to uniform hyperbolicity, they are still valid for our fluid system, which only exhibits some hyperbolic phenomena (angles between far-apart CLVs are large), but is not uniform hyperbolic.

5. Conclusions

In this paper, we first compute the Lyapunov exponents (LEs) of a 3D chaotic flow past a cylinder at Reynolds number 525. Our computation shows that this flow problem has less than 30 positive LEs, but the detailed LE spectrum depends on meshes. For both meshes, we find that the Lyapunov dimension is smaller than 109, meaning that the apparently complicated dynamics of our 3D flow problem can be attributed to the interaction among less than 109 degrees of freedom.

We then compute the first 40 covariant Lyapunov vectors (CLVs) of this 3D flow problem. We find that the angles between CLVs become larger when the indices are further apart, although there may be occasional tangencies between adjacent CLVs. Our conjecture is backed by plots of CLVs, where we find that unstable CLVs are active in the instability-generating area such as boundary layers and near wakes, while

stable CLVs are active in more dissipative areas such as the far wake, the front of the cylinder and the free stream. This difference in active areas is robust to meshes, and it indicates that CLVs point to different directions, and it also implies that for open flows there is a large fraction of CLVs that are stable.

We also conclude that our system is not uniform hyperbolic. Firstly, due to the extra neutral CLV corresponding to translation along the spanwise direction, our problem has at least two neutral directions. Hence at least the first assumption in uniform hyperbolicity is violated on both meshes. Secondly, we found that the smallest angle between CLVs is not robust to meshes and time length, since on the finer mesh the smallest angle is well above the threshold value, whereas on the coarser mesh it almost equals the threshold.

We then use the finite difference non-intrusive least-squares shadowing (FD-NILSS) algorithm to compute the shadowing direction v^\perp of this 3D flow problem. The magnitude of the computed v^\perp remains on the same level as time evolves, thus indicating the existence of shadowing directions of this 3D flow problem. By observing v^\perp , we find that its value in the free-stream area is proportional to perturbations on the x -directional free-stream flow speed U , but not affected by perturbations on the cylinder rotation ω . On the other hand, v^\perp have similar magnitudes in the wake area for unit perturbations on U and ω .

Finally, we use shadowing directions to compute the sensitivities of some long-time-averaged objectives. For both meshes, the computed sensitivities correctly reflect the trend between system parameters and objectives. Our results confirm the physical meaning of the shadowing solution in revealing sensitivities of all physical phenomena in this flow field with respect to perturbations on a parameter. The sensitivity results also suggest that shadowing methods are robust to numerical implementations, and may be valid for more general chaotic fluid systems, although shadowing theories are mathematically developed under uniform hyperbolicity.

Acknowledgements

The author is indebted to Professor Q. Wang at MIT, Dr C. Talnikar at NVIDIA, and N. Chandramoorthy at MIT for both software and hardware support.

Supplementary movies

Supplementary movies are available at <https://doi.org/10.1017/jfm.2018.986>.

REFERENCES

- ALBERS, D. J. & SPOTT, J. C. 2006 Structural stability and hyperbolicity violation in high-dimensional dynamical systems. *Nonlinearity* **19** (8), 1801–1847.
- ANTALI, M. & HORVATH, Z. 2018 Computation of sensitivity of periodically excited dynamical systems. *J. Engng Maths* **113** (1), 123–142.
- BECKER, R. & RANNACHER, R. 2001 An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica* **10**, 1–102.
- BENETTIN, G., GALGANI, L., GIORGILLI, A. & STRELCYN, J.-M. 1980 Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. Part 2. Numerical application. *Meccanica* **15** (1), 21–30.
- BEWLEY, T. R. 2001 Flow control: new challenges for a new renaissance. *Prog. Aerosp. Sci.* **37** (1), 21–58.

- BLONIGAN, P. J. 2017 Adjoint sensitivity analysis of chaotic dynamical systems with non-intrusive least squares shadowing. *J. Comput. Phys.* **348**, 803–826.
- BLONIGAN, P. J., FERNANDEZ, P., MURMAN, S. M. M., WANG, Q., RIGAS, G. & MAGRI, L. 2016a *Towards a Chaotic Adjoint for LES*, pp. 385–394. Center for Turbulence Research, Summer Program.
- BLONIGAN, P. J., WANG, Q., NIELSEN, E. J. & DISKIN, B. 2016b Least squares shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil. *AIAA J.* **56**, 658–672.
- BOWEN, R. 1970 Markov partitions for Axiom A diffeomorphisms. *Amer. J. Math.* **92** (3), 725–747.
- BRÈS, G. A., HAM, F. E., NICHOLS, J. W. & LELE, S. K. 2017 Unstructured large-eddy simulations of supersonic jets. *AIAA J.* **55** (4), 1164–1184.
- CHORIN, A. J., HALD, O. H. & KUPFERMAN, R. 2002 Optimal prediction with memory. *Physica D* **166** (3–4), 239–257.
- COLONIUS, T., LELE, S. K. & MOIN, P. 1993 Boundary conditions for direct computation of aerodynamic sound generation. *AIAA J.* **31** (9), 1574–1582.
- FERNANDEZ, P. & WANG, Q. 2017a Lyapunov spectrum of chaotic flow simulations. Application to chaotic adjoints. *AIAA Paper* 2017–3797.
- FERNANDEZ, P. & WANG, Q. 2017b Lyapunov spectrum of the separated flow around the NACA 0012 airfoil and its dependence on numerical discretization. *J. Comput. Phys.* **350**, 453–469.
- FIDKOWSKI, K. J. & DARMOFAL, D. L. 2011 Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *AIAA J.* **49** (4), 673–694.
- FREDERICKSON, P., KAPLAN, J. L., YORKE, E. D. & YORKE, J. A. 1983 The Liapunov dimension of strange attractors. *J. Differ. Equ.* **49** (2), 185–207.
- GALLAVOTTI, G. 2006 Entropy, thermostats, and chaotic hypothesis. *Chaos* **16** (4), 043114.
- GALLAVOTTI, G. 2008 Chaotic hypothesis. *Scholarpedia* 3(1):5906 **3** (1), 8–11.
- GALLAVOTTI, G. & COHEN, E. G. D. 1995 Dynamical ensembles in stationary states. *J. Stat. Phys.* **80** (5–6), 931–970.
- GILES, M. B. & SÜLI, E. 2002 Adjoint methods for PDEs: a posteriori error analysis and postprocessing by duality. *Acta Numerica* **11**, 145–236.
- GINELLI, F., CHATÉ, H., LIVI, R. & POLITI, A. 2013 Covariant Lyapunov vectors. *J. Phys. A* **46** (25), 254005.
- GINELLI, F., POGGI, P., TURCHI, A., CHATÉ, H., LIVI, R. & POLITI, A. 2007 Characterizing dynamics with covariant Lyapunov vectors. *Phys. Rev. Lett.* **99** (13), 1–4.
- INUBUSHI, M., KOBAYASHI, M. U., TAKEHIRO, S. I. & YAMADA, M. 2012 Covariant Lyapunov analysis of chaotic Kolmogorov flows and time-correlation function. *Proc. IUTAM* **5**, 244–248; Published by Elsevier.
- JAMESON, A. 1988 Aerodynamic design via control theory. *J. Sci. Comput.* **3** (3), 233–260.
- KATOK, A. B. & HASSELBLATT, B. A. 1997 *Introduction to the Modern Theory of Dynamical Systems*, Encyclopedia of Mathematics and its Applications, vol. 54. Cambridge University Press.
- KEEFE, L., MOIN, P. & KIM, J. 1992 The dimension of attractors underlying periodic turbulent Poiseuille flow. *J. Fluid Mech.* **242** (29), 1–29.
- MITTAL, R. & BALACHANDAR, S. 1996 Direct numerical simulation of flow past elliptic cylinders. *J. Comput. Phys.* **367** (124), 351–367.
- NI, A. 2018a Adjoint sensitivity analysis on chaotic dynamical systems by non-intrusive least squares adjoint shadowing (NILSAS), [arXiv:1801.08674](https://arxiv.org/abs/1801.08674), pp. 1–44.
- NI, A. 2018b Adjoint shadowing directions in hyperbolic systems for sensitivity analysis. [arXiv:1807.05568](https://arxiv.org/abs/1807.05568), pp. 1–23.
- NI, A., BLONIGAN, P. J., CHATER, M., WANG, Q. & ZHANG, Z. 2016 Sensitivity analysis on chaotic dynamical system by non-intrusive least square shadowing (NI-LSS). *AIAA Paper* 2016–4399.
- NI, A. & WANG, Q. 2017 Sensitivity analysis on chaotic dynamical systems by non-intrusive least squares shadowing (NILSS). *J. Comput. Phys.* **347**, 56–77.
- NI, A., WANG, Q., FERNANDEZ, P. & TALNIKAR, C. 2018 Sensitivity analysis on chaotic dynamical systems by finite difference non-intrusive least squares shadowing (FD-NILSS), [arXiv:1711.06633](https://arxiv.org/abs/1711.06633), pp. 1–21.

- OSELEDETS, V. I. 1968 A multiplicative ergodic theorem. Characteristic Ljapunov, exponents of dynamical systems. *Tr. Moskov. Mat. Obshchestva* **19**, 179–210.
- OSELEDETS, V. I. 2008 Oseledets theorem. *Scholarpedia* **3** (1), 1846.
- PARISH, E. J. & DURAISAMY, K. 2017 A dynamic subgrid scale model for large eddy simulations based on the Mori–Zwanzig formalism. *J. Comput. Phys.* **349**, 154–175.
- PILYUGIN, S. Y. 1997 Shadowing in structurally stable flows. *J. Differ. Equ.* **140** (2), 238–265.
- PILYUGIN, S. Y. 1999 *Shadowing in Dynamical Systems*, Lecture Notes in Mathematics, vol. 1706. Springer.
- POINSOT, T. J. & LELEF, S. K. 1992 Boundary conditions for direct simulations of compressible viscous flows. *J. Comput. Phys.* **101** (1), 104–129.
- REICK, C. H. 2002 Linear response of the Lorenz system. *Phys. Rev. E* **66** (3), 1–11.
- REZAKHANLOU, F. 2018 *Lectures on Dynamical Systems*. Department of Mathematics, University of California, Berkeley.
- RUELLE, D. 1980 Measures describing a turbulent flow. *Ann. N.Y. Acad. Sci.* **357** (1), 1–9.
- RUELLE, D. 1997 Differentiation of SRB states. *Commun. Math. Phys.* **187**, 227–241.
- RUELLE, D. 2008 Differentiation of SRB states for hyperbolic flows. *Ergod. Theory Dyn. Sys.* **28** (02), 613–631.
- SIEBER, M. 1987 Experiments on the attractor-dimension for turbulent pipe flow. *Phys. Lett. A* **122** (9), 467–470.
- THEPAUT, J.-N. & COURTIER, P. 1991 Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model. *Q. J. R. Meteorol. Soc.* **117** (502), 1225–1254.
- TROMP, J., TAPE, C. & LIU, Q. 2005 Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. *Geophys. J. Intl* **160**, 195–216.
- WANG, Q. 2014 Convergence of the least squares shadowing method for computing derivative of ergodic averages. *SIAM J. Numer. Anal.* **52** (1), 156–170.
- WANG, Q., HU, R. & BLONIGAN, P. 2014 Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *J. Comput. Phys.* **267**, 210–224.
- WILLIAMSON, C. H. K. & ROSHKO, A. 1990 Measurements of base pressure in the wake of a cylinder at low Reynolds numbers. *Z. Flugwiss. Weltraumforsch.* **14**, 38–46.
- XU, M. & PAUL, M. R. 2016 Covariant Lyapunov vectors of chaotic Rayleigh–Bénard convection. *Phys. Rev. E* **93** (6), 1–12.
- YOUNG, L.-S. 2002 What are SRB measures, and which dynamical systems have them? *J. Stat. Phys.* **108** (5), 733–754.
- YOUNG, L.-S. 2003 Entropy in dynamical systems. In *Entropy*, (ed. Greven, Keller & Warnecke), pp. 313–328. Princeton University Press.