

The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds

Francois Lekien^{1,a)} and Shane D. Ross^{2,b)}

¹*École Polytechnique, Université Libre de Bruxelles, B-1050 Brussels, Belgium*

²*Engineering Science and Mechanics, Virginia Tech, Blacksburg, Virginia 24061, USA*

(Received 4 September 2009; accepted 4 December 2009; published online 8 February 2010)

We generalize the concepts of finite-time Lyapunov exponent (FTLE) and Lagrangian coherent structures to arbitrary Riemannian manifolds. The methods are illustrated for convection cells on cylinders and Möbius strips, as well as for the splitting of the Antarctic polar vortex in the spherical stratosphere and a related point vortex model. We modify the FTLE computational method and accommodate unstructured meshes of triangles and tetrahedra to fit manifolds of arbitrary shape, as well as to facilitate dynamic refinement of the FTLE mesh. © 2010 American Institute of Physics. [doi:10.1063/1.3278516]

Riemannian manifolds are ubiquitous in science and engineering, being the more natural mathematical setting for many dynamical systems. For instance, transport along isopycnal surfaces in the ocean and large-scale mixing in the atmosphere are processes taking place on a curved manifold, not a vector space. In this paper, we generalize the notion of finite-time Lyapunov exponent (FTLE) and Lagrangian coherent structures (LCS) to arbitrary Riemannian differentiable manifolds. We show that both notions are independent of the coordinate system but depend on the chosen metric. However, we find that LCS do not depend much on the metric. The FTLE measures separation and tends to be large and positive along LCS and very small elsewhere. For sufficiently large integration times, the steep variations of the FTLE field cannot be modified much by smooth changes in the metric and the LCS remain essentially unchanged. Approximating, or even ignoring, the manifold metric does not influence the result for large integration times, and therefore, computing the FTLE field on manifolds is robust. Aside from these conclusions, we present a general algorithm for computing the FTLE on manifolds covered with meshes of polyhedra. The algorithm requires knowledge of the mesh nodes, the image of the mesh nodes under the flow, as well as information about node neighbors (but not the full connectivity). We also used the same algorithm in Euclidian spaces where the unstructured mesh permits efficient adaptive refinement for capturing sharp LCS features. We illustrate the results and methods on several systems: convection cells in a plane, on a cylinder, and on a Möbius strip, as well as atmospheric transport resulting from the 2002 splitting of the Antarctic ozone hole in the spherical stratosphere and a related point vortex model on the sphere.

I. INTRODUCTION

Consider an Euclidean domain $\Omega \subseteq \mathbb{R}^n$. A typical dynamical system is given in the form of a velocity field $\mathbf{v}(\mathbf{x}, t)$ defined on $\Omega \times \mathbb{R}$. In this case, the trajectories are given by the solutions of the ordinary differential equation $\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t)$. Each trajectory is a function of time but it also depends on the initial position \mathbf{x}_0 and the initial time t_0 . How the trajectory changes as the initial time and the initial position changes are a central interest in this paper. For this reason, the trajectory which is solution of the initial value problem

$$\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t) \quad \text{and} \quad \mathbf{x}(t_0) = \mathbf{x}_0, \quad (1)$$

is denoted as $\mathbf{x}(t; t_0, \mathbf{x}_0)$, which emphasizes the explicit dependence on the initial condition. It is also convenient to define the flow operator to further highlight the dependence on the initial position \mathbf{x}_0 . For a given initial time t_0 and a given final time t , the flow map is the function

$$\phi_{t_0}^t: \Omega \rightarrow \Omega: \mathbf{x}_0 \mapsto \phi_{t_0}^t(\mathbf{x}_0) = \mathbf{x}(t; t_0, \mathbf{x}_0). \quad (2)$$

Autonomous systems do not depend on time. They satisfy $\phi_{t_0}^{t+T} = \phi_{t_0}^t$ for any $T \in \mathbb{R}$. For the associated velocity field, this property translates into $\mathbf{v}(\mathbf{x}, t) = \mathbf{v}(\mathbf{x})$. For two-dimensional autonomous systems, regions of qualitatively different dynamics can be efficiently captured by the stable and unstable invariant manifolds of saddle points in the system.¹⁸ In higher dimensional autonomous systems, separatrices are codimension 1 stable and unstable manifolds of normally hyperbolic manifolds.^{11,40}

In the Poincaré sections of time-periodic systems, lobes between the invariant manifolds govern particle transport across separatrices.^{48,47} In this paper, we consider the generalization of this framework to systems with arbitrary time dependence. In this case, saddle fixed points are nongeneric. The points, where the instantaneous velocity vanishes, are called hyperbolic stagnation points and usually meander in time.⁸ Such points are not trajectories and are not associated with hyperbolic invariant manifolds.³⁶ The concept of LCS

^{a)}Electronic mail: lekien@ulb.ac.be.

^{b)}Electronic mail: sdross@vt.edu.

extends the transport framework delineated by invariant manifolds to time-dependent dynamical systems and to transient processes.^{19,20,52,37}

How can we define coherent structures without relying on stagnation points and without using asymptotic notions such as convergence, hyperbolicity, or exponential dichotomies? Bowman⁴ provided a robust and general answer. He suggested to initiate a grid of virtual particles, integrate them for a time on the scale of the observed process, and look at the final separation between initially nearby particles (finite-strain maps). LCS correspond to curves with locally high relative dispersion.

A remarkable aspect of Bowman’s procedure is its robustness and its ability to uncover structures in the most complex and noisiest data sets. It can be improved by replacing the linear separation between grid neighbors by a better stretching estimate. Two such modern indicators are the finite-time Lyapunov exponent (FTLE)^{19,20} and the finite-size Lyapunov exponent (FSLE).^{27,31} The two quantities are closely related and delineate sharp ridges of high stretching that behave almost like material lines.^{52,37} In the remainder of this paper, we will concentrate on FTLE but most of the results and methods remain valid for other stretching indicators such as FSLE or relative dispersion.

For a given time t_0 and an initial position \mathbf{x}_0 , the quantity $\phi_{t_0}^{t_0+T}(\mathbf{x}_0)$ is the position after an integration time T . To quantify particle separation and sensitivity to the initial condition \mathbf{x}_0 , the FTLE is defined based on the derivative of $\phi_{t_0}^{t_0+T}(\mathbf{x}_0)$ with respect to \mathbf{x}_0 . If we make a small perturbation and we consider the evolution of trajectory starting at the position $\mathbf{x}_0 + \mathbf{u}$ instead of \mathbf{x}_0 , the final position will be

$$\phi_{t_0}^{t_0+T}(\mathbf{x}_0 + \mathbf{u}) = \phi_{t_0}^{t_0+T}(\mathbf{x}_0) + \frac{\partial \phi_{t_0}^{t_0+T}}{\partial \mathbf{x}_0} \mathbf{u} + \mathcal{O}(\|\mathbf{u}\|^2). \tag{3}$$

The point \mathbf{x}_0 is *sensitive to the initial condition* if $\phi_{t_0}^{t_0+T}(\mathbf{x}_0 + \mathbf{u})$ separates from $\phi_{t_0}^{t_0+T}(\mathbf{x}_0)$, even for small perturbations \mathbf{u} . The boundary between regions of qualitatively different dynamics can therefore be detected as the set of points for which $(\partial \phi_{t_0}^{t_0+T} / \partial \mathbf{x}_0) \mathbf{u}$ is very large for some direction \mathbf{u} . To quantify stretching and sensitivity to initial conditions, we define the FTLE as

$$\sigma(\mathbf{x}_0, t_0) = \frac{1}{|T|} \ln \left(\max_{\mathbf{u} \neq 0} \frac{\left\| \frac{\partial \phi_{t_0}^{t_0+T}}{\partial \mathbf{x}_0} \mathbf{u} \right\|}{\|\mathbf{u}\|} \right). \tag{4}$$

In the equation above, the expression in parentheses gives the maximum stretching between infinitesimally close trajectories. The reason for taking the logarithm and dividing by the magnitude of the integration time T is the parallel that we then establish with autonomous systems. Indeed, for a time-independent system, we have $\max_{\mathbf{u} \neq 0} \|(\partial \phi_{t_0}^{t_0+T} / \partial \mathbf{x}_0) \mathbf{u}\| = e^{\gamma} \|\mathbf{u}\|$, where γ is the largest Lyapunov exponent associated with the trajectory starting at \mathbf{x}_0 .

Note that the FTLE $\sigma(\mathbf{x}_0, t_0)$ is a function of position and time but it also depends on the chosen integration time T . The latter is not regarded as very large as our aim is to investigate transient processes and finite-time mixing.^{21,36}

Additionally, $|T|$ is used instead of T because computing the FTLE for $T > 0$ and $T < 0$ produces repulsive LCS and attractive LCS, respectively, and thus this definition facilitates forward and backward time computations.

For a given point \mathbf{x}_0 , the FTLE $\sigma(\mathbf{x}_0, t_0)$, as defined above, can also be seen as the norm of the differential of $\phi_{t_0}^{t_0+T}$ at the point \mathbf{x}_0 , that is

$$\sigma(\mathbf{x}_0, t_0) = \frac{1}{|T|} \ln \|D \phi_{t_0}^{t_0+T}\| \doteq \frac{1}{|T|} \ln \left(\max_{\mathbf{u} \neq 0} \frac{\|D \phi_{t_0}^{t_0+T}(\mathbf{u})\|}{\|\mathbf{u}\|} \right). \tag{5}$$

We will use the expression above to extend the definition of the FTLE to non-Euclidean manifolds and to derive an algorithm for computing FTLE on an unstructured mesh (for Euclidean or non-Euclidean manifolds).

The LCS are ridges in the FTLE field. They are codimension 1 manifolds that maximize the FTLE in the transverse direction. This notion is more formally defined by Shadden *et al.*⁵² and Lekien *et al.*³⁷ The objective is to generalize the concept of stable and unstable invariant manifolds (of hyperbolic invariant objects) to time-dependent systems. Whether in the ocean, in the atmosphere, or any other dynamical system, LCS are barriers to particle transport.⁴⁵

Interlacing of attractive and repulsive LCS gives rise to lobes that govern transport across the regions delineated by the LCS.⁸ The LCS bound coherent patches of particles that evolve independently from the surrounding environment and exchange little matter and energy.¹² Such structures include oceanic eddies,² separation profiles from a coastline or an airfoil,^{36,39} shed vortices over an airfoil,⁶ and wake structures behind swimming animals.⁴⁶ Shadden and Taylor⁵³ showed that LCS indicate stagnation and mixing in blood vessels. Tallapragada and Ross⁵⁷ showed that LCS are also relevant for nonpassive finite-size particles in fluids, where inertial effects are present.

Contrary to classical stable and unstable invariant manifolds, the LCS of time-dependent systems are not perfect barriers to transport. Shadden *et al.*⁵² and Lekien *et al.*³⁷ showed that the flux through a LCS is usually very small but not exactly zero. As a result, one can usually consider LCS as a material barrier with negligible flux, but with the possibility of occasional bifurcations. Invariant manifolds cannot bifurcate and cannot adequately describe transient processes. The weak flux across the LCS of an aperiodic system makes it possible for the LCS to bifurcate and capture finite-time processes.

In the ocean, bifurcation of the LCS indicates a regime change between qualitatively different kinds of motion.^{35,33} Coulliette *et al.*⁷ showed that the position of LCS in coastal regions determines whether contaminants will recirculate along the coast or not. Lipinski *et al.*³⁹ showed that LCS bifurcation at the end of an airfoil is responsible for vortex shedding.

In an atmospheric context, LCS correspond to atmospheric transport barriers which separate different regions of relatively well-mixed air. In this paper, we will look into atmospheric transport and global-scale phenomena, such as the splitting of the Antarctic polar vortex.^{54,61,42} As the atmo-

sphere is global in extent, it becomes necessary to consider the spherical, and therefore non-Euclidean, nature of the flow.

II. FINITE-TIME LYAPUNOV EXPONENTS IN EUCLIDEAN SPACES

The computation of the FTLE can be further simplified in Euclidean spaces by considering the matrix representation of the differential $D\phi_{t_0}^{t_0+T}$. Let us consider an orthonormal basis $\{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ of \mathbb{R}^n and the $n \times n$ matrix $M = (m_{ij})$ defined by $m_{ij} \doteq \hat{e}_i \cdot (D\phi_{t_0}^{t_0+T}(\hat{e}_j))$.

Given the column vector U containing the coordinates of a vector $\mathbf{u} \in \mathbb{R}^n$ in the orthonormal basis $\{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$, the column vector $U' = MU$ gives the coordinates of $D\phi_{t_0}^{t_0+T}(\mathbf{u})$ in the same basis. As a result, we have

$$\begin{aligned} \sigma(\mathbf{x}_0, t_0) &= \frac{1}{|T|} \ln \left(\max_{\mathbf{u} \neq \mathbf{0}} \frac{\|D\phi_{t_0}^{t_0+T}(\mathbf{u})\|}{\|\mathbf{u}\|} \right) \\ &= \frac{1}{2|T|} \ln \left(\max_{U \in \mathbb{R}_0^n} \frac{U^T M^T M U}{U^T U} \right) \\ &= \frac{1}{|T|} \ln (\text{largest singular value } (M)). \end{aligned} \tag{6}$$

We first review how the formula above can be used to compute FTLE and LCS on a Cartesian mesh. We then turn to the implementation for unstructured lists of triangles (in \mathbb{R}^2) and n -tetrahedra (in \mathbb{R}^n).

A. Computation with a Cartesian mesh

Let us consider a Cartesian mesh with grid points $(x_{ijk}, y_{ijk}, \dots, z_{ijk}) \in \mathbb{R}^n$. We compute the final position $\phi_{t_0}^{t_0+T}(x_{ijk}, y_{ijk}, \dots, z_{ijk}) = (x'_{ijk}, y'_{ijk}, \dots, z'_{ijk})$. The matrix representation of the flow differential is given by

$$M = \underbrace{\begin{pmatrix} \frac{x'_{i+1j \dots k} - x'_{i-1j \dots k}}{2\delta_x} & \frac{x'_{ij+1 \dots k} - x'_{ij-1 \dots k}}{2\delta_y} & \dots & \frac{x'_{ij \dots k+1} - x'_{ij \dots k-1}}{2\delta_z} \\ \frac{y'_{i+1j \dots k} - y'_{i-1j \dots k}}{2\delta_x} & \frac{y'_{ij+1 \dots k} - y'_{ij-1 \dots k}}{2\delta_y} & \dots & \frac{y'_{ij \dots k+1} - y'_{ij \dots k-1}}{2\delta_z} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{z'_{i+1j \dots k} - z'_{i-1j \dots k}}{2\delta_x} & \frac{z'_{ij+1 \dots k} - z'_{ij-1 \dots k}}{2\delta_y} & \dots & \frac{z'_{ij \dots k+1} - z'_{ij \dots k-1}}{2\delta_z} \end{pmatrix}}_M + \mathcal{O}(\delta), \tag{7}$$

where $\delta = \max\{\delta_x, \delta_y, \dots, \delta_z\}$ is the maximum step size.

One can then find the largest singular value of the matrix \tilde{M} that approximates M for a sufficiently small grid spacing δ . Similarly, as suggested by Haller,²⁰ one can search for the largest eigenvalue of $\tilde{M}^T \tilde{M}$. The algorithm above is therefore identical to the algorithm proposed as the two-dimensional direct Lyapunov exponent (DLE) of Haller,²⁰ the two-dimensional algorithm of Shadden *et al.*,⁵² or the n -dimensional algorithm of Lekien *et al.*³⁷

The major advantage of this algorithm is its robustness. Regardless of the grid size δ , an existing LCS passing between two grid points creates a visible shade in the approximated FTLE field. As shown in Fig. 1, there is more to this algorithm than just approximating a derivative with finite differences. If we were to approximate the derivative of the flow using arbitrary control points, we may be missing some LCS passing between grid points. Similarly, one could take each grid point separately and integrate the velocity gradient or get the Cauchy–Green strain tensor and the FTLE. This would also lead to possible LCS obliteration. The DLE algo-

rithm estimates total (averaged) stretching in the mesh: Only grid points are used in the computation of the derivatives by finite differences.

The disadvantage of the technique is that the FTLE can be underestimated if the elements are too large: One often needs to refine the mesh in order to avoid saturation. The advantage is that it detects LCS no matter what the distribution of grid points is. Even on coarse meshes, one is able to identify a rough estimate of the LCS structure.

The singular value decomposition of M always exists and is continuous with respect to any parameter, such as the grid size δ .¹⁰ As a result, provided that the largest singular value does not vanish, its logarithm is also continuous. This is, however, the case since the flow is one to one. As a result, none of the column vectors in M are identically zero (this would violate the uniqueness of solutions). This implies that the rank of M is at least 1 and its largest singular value is strictly positive.

Note that the FTLE is typically computed at grid resolutions much higher than the flow resolution.^{52,60} The FTLE indicates long term particle separation and the LCS have a




FIG. 1. Computation of FTLE on a Cartesian mesh: A mesh of points is initialized (left panel), then integrated for a given integration time T (right panel). The derivative of the flow at one grid point is evaluated using finite differences with other grid points. If one were to evaluate the flow derivative using off grid test points (gray dots), a LCS (dashed line) passing between grid points could go unnoticed.

fine, foliated structure, even for flows defined on a coarse grid. In time-periodic systems, the LCS correspond to the invariant manifold of hyperbolic periodic orbit and they are known to form tangles and arbitrary small lobes.^{48,47} For arbitrary time dependence, the LCS can also form meanders of arbitrary small size, provided that the integration time is long enough. To avoid saturation, the resolution of the mesh has to be adapted to the amplitude of the deformation. Consequently, the mesh resolution depends on the chosen integration time, rather than on the resolution or the input velocity. One of the contributions of this paper is to generalize the FTLE algorithm to unstructured meshes that can be efficiently refined by increasing the resolution only near the LCS.

B. Computation with an unstructured mesh (in \mathbb{R}^2)

While robust and efficient, the computation of FTLE using Cartesian meshes does not translate easily to arbitrary manifolds. For manifolds such as spheres, tori, and more complicated shapes, it is usually more convenient to cover the space with a list of connected polygons. Another incentive to use unstructured meshes is the need for higher resolution. To obtain sharper and longer LCS, we need to straddle them with grid points that are as close to each other as possible. With Cartesian meshes, it is difficult to refine the mesh other than uniformly. Using unstructured meshes, one can easily adapt the resolution locally: Smaller elements are needed in regions where FTLE indicate the presence of LCS, while low FTLE areas can be decimated.¹⁴

Another reason for computing FTLE on unstructured meshes rather than Cartesian meshes is to match the format of some input velocity fields. For instance, methods such as normal mode analysis reconstruct coastal surface current maps from radar data and provide reconstructed stream function and velocity potential on meshes of triangles.^{34,29} Computing FTLE on an unstructured mesh whose boundary matches that of the input velocity field greatly improves the smoothness of the FTLE map.

We begin by adapting the algorithm described in Sec. II A for meshes of triangles in \mathbb{R}^2 . We then generalize this method to meshes of n -tetrahedra covering a subspace of \mathbb{R}^n .

Consider a domain $\Omega \subset \mathbb{R}^2$ which is covered by a list of N non-overlapping triangles. For each vertex $\mathbf{p}_i = (\xi_1^{(i)}, \xi_2^{(i)})$, we assume that the image by the flow, $\mathbf{p}_i^t = \phi_{t_0}^{t_0+T}(\mathbf{p}_i) = (\eta_1^{(i)}, \eta_2^{(i)})$, has been computed.

Point \mathbf{p}_i has an arbitrary (≥ 2) number of neighbors from which we must approximate the matrix of the derivative of the flow map. Let us consider the N neighbors \mathbf{p}_j of \mathbf{p}_i and construct the $2 \times N$ matrices

$$Y = \begin{pmatrix} \eta_1^{(1)} - \eta_1^{(i)} & \eta_1^{(2)} - \eta_1^{(i)} & \dots & \eta_1^{(N)} - \eta_1^{(i)} \\ \eta_2^{(1)} - \eta_2^{(i)} & \eta_2^{(2)} - \eta_2^{(i)} & \dots & \eta_2^{(N)} - \eta_2^{(i)} \end{pmatrix} \quad (8)$$

and

$$X = \begin{pmatrix} \xi_1^{(1)} - \xi_1^{(i)} & \xi_1^{(2)} - \xi_1^{(i)} & \dots & \xi_1^{(N)} - \xi_1^{(i)} \\ \xi_2^{(1)} - \xi_2^{(i)} & \xi_2^{(2)} - \xi_2^{(i)} & \dots & \xi_2^{(N)} - \xi_2^{(i)} \end{pmatrix}. \quad (9)$$

There does not usually exist a 2×2 matrix M such that $Y = MX$, but we can compute the least square estimate

$$\tilde{M} = YX^T(XX^T)^{-1}, \quad (10)$$

which minimizes the L_2 error $\|Y - MX\|^2$.

We can then proceed with the computation of the largest singular value of \tilde{M} and obtain an estimate of the FTLE at point \mathbf{p}_i of the unstructured mesh. There is, however, a slight modification that can be made to improve the algorithm. In the procedure described above, each neighbor accounts for two equally weighted data in the least square estimate. Nonetheless, it is not uncommon for a vertex \mathbf{p}_i of an unstructured mesh to have more neighbors in one direction than another. In this case, the reconstructed matrix M is distorted by data recorded in directions where the density of grid points is higher. This distortion can be removed by weighting each equation involved in the least square estimate. Let us consider the Voronoi cell around vertex \mathbf{p}_i that is the polygon




FIG. 2. Deformation of an unstructured mesh under the flow. Left panel: the node \mathbf{p}_i has N neighbors $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_j, \dots, \mathbf{p}_N$. Right panel: under the action of the flow, the node \mathbf{p}_i moves to the position \mathbf{p}'_i . The deformed edges $\mathbf{p}'_i-\mathbf{p}'_1, \mathbf{p}'_i-\mathbf{p}'_2, \dots, \mathbf{p}'_i-\mathbf{p}'_j, \dots, \mathbf{p}'_i-\mathbf{p}'_N$ are used to approximate the deformation tensor.

where all the points are closer to vertex \mathbf{p}_i than any other vertex. The area of the Voronoi cell is

$$\mathcal{V}_i = \frac{1}{8} \sum_{\substack{\mathbf{p}_j \text{ is a} \\ \text{neigh.} \\ \text{of } \mathbf{p}_i}} A_{ij}, \tag{11}$$

where A_{ij} is the area of the two triangles that contain both \mathbf{p}_i and \mathbf{p}_j . We can then define

$$Y = \begin{pmatrix} \frac{A_{1i}}{8\mathcal{V}_i}(\eta_1^{(1)} - \eta_1^{(i)}) & \dots & \frac{A_{Ni}}{8\mathcal{V}_i}(\eta_1^{(N)} - \eta_1^{(i)}) \\ \frac{A_{1i}}{8\mathcal{V}_i}(\eta_2^{(1)} - \eta_2^{(i)}) & \dots & \frac{A_{Ni}}{8\mathcal{V}_i}(\eta_2^{(N)} - \eta_2^{(i)}) \\ \dots & \vdots & \dots \\ \frac{A_{1i}}{8\mathcal{V}_i}(\eta_n^{(1)} - \eta_n^{(i)}) & \dots & \frac{A_{Ni}}{8\mathcal{V}_i}(\eta_n^{(N)} - \eta_n^{(i)}) \end{pmatrix} \tag{12}$$

and

$$X = \begin{pmatrix} \frac{A_{1i}}{8\mathcal{V}_i}(\xi_1^{(1)} - \xi_1^{(i)}) & \dots & \frac{A_{Ni}}{8\mathcal{V}_i}(\xi_1^{(N)} - \xi_1^{(i)}) \\ \frac{A_{1i}}{8\mathcal{V}_i}(\xi_2^{(1)} - \xi_2^{(i)}) & \dots & \frac{A_{Ni}}{8\mathcal{V}_i}(\xi_2^{(N)} - \xi_2^{(i)}) \\ \dots & \vdots & \dots \\ \frac{A_{1i}}{8\mathcal{V}_i}(\xi_n^{(1)} - \xi_n^{(i)}) & \dots & \frac{A_{Ni}}{8\mathcal{V}_i}(\xi_n^{(N)} - \xi_n^{(i)}) \end{pmatrix}. \tag{13}$$

We now have

$$M = \underbrace{YX^T(XX^T)^{-1}}_{\tilde{M}} + \mathcal{O}(\delta) \tag{14}$$

and we approximate the FTLE by

$$\sigma(\mathbf{p}_i, t_0) = \frac{1}{|T|} \ln(\|YX^T(XX^T)^{-1}\|), \tag{15}$$

where the norm of the matrix is defined as its largest singular value. Note that the factors $8\mathcal{V}_i$ cancel out from the equations, hence we can ignore them.

C. Computation with an unstructured mesh (in \mathbb{R}^n)

To generalize the computation of FTLE to unstructured meshes in \mathbb{R}^n , we consider a domain $\Omega \subset \mathbb{R}^n$ that is covered by a mesh of n -dimensional tetrahedra. For each vertex $\mathbf{p}_i = (\xi_1^{(i)}, \xi_2^{(i)}, \dots, \xi_n^{(i)})$ of the mesh, we consider the N neighbors $\mathbf{p}_j = (\xi_1^{(j)}, \xi_2^{(j)}, \dots, \xi_n^{(j)})$. The images of these vertices un-

der the flow map are written as $\mathbf{p}'_i = \phi_{t_0}^{t_0+T}(\mathbf{p}_i) = (\eta_1^{(i)}, \eta_2^{(i)}, \dots, \eta_n^{(i)})$ and $\mathbf{p}'_j = \phi_{t_0}^{t_0+T}(\mathbf{p}_j) = (\eta_1^{(j)}, \eta_2^{(j)}, \dots, \eta_n^{(j)})$. This setting is depicted on Fig. 2.

For each pair of adjacent vertices \mathbf{p}_i and \mathbf{p}_j , we define the hypervolume A_{ij} as the sum of the hypervolumes of all the n -tetrahedra that contain both \mathbf{p}_i and \mathbf{p}_j . We define the $n \times N$ matrices

$$Y = \begin{pmatrix} A_{1i}(\eta_1^{(1)} - \eta_1^{(i)}) & A_{2i}(\eta_1^{(2)} - \eta_1^{(i)}) & \dots & A_{Ni}(\eta_1^{(N)} - \eta_1^{(i)}) \\ A_{1i}(\eta_2^{(1)} - \eta_2^{(i)}) & A_{2i}(\eta_2^{(2)} - \eta_2^{(i)}) & \dots & A_{Ni}(\eta_2^{(N)} - \eta_2^{(i)}) \\ \dots & \vdots & \ddots & \vdots \\ A_{1i}(\eta_n^{(1)} - \eta_n^{(i)}) & A_{2i}(\eta_n^{(2)} - \eta_n^{(i)}) & \dots & A_{Ni}(\eta_n^{(N)} - \eta_n^{(i)}) \end{pmatrix} \tag{16}$$

and

$$X = \begin{pmatrix} A_{1i}(\xi_1^{(1)} - \xi_1^{(i)}) & A_{2i}(\xi_1^{(2)} - \xi_1^{(i)}) & \dots & A_{Ni}(\xi_1^{(N)} - \xi_1^{(i)}) \\ A_{1i}(\xi_2^{(1)} - \xi_2^{(i)}) & A_{2i}(\xi_2^{(2)} - \xi_2^{(i)}) & \dots & A_{Ni}(\xi_2^{(N)} - \xi_2^{(i)}) \\ \dots & \vdots & \ddots & \vdots \\ A_{1i}(\xi_n^{(1)} - \xi_n^{(i)}) & A_{2i}(\xi_n^{(2)} - \xi_n^{(i)}) & \dots & A_{Ni}(\xi_n^{(N)} - \xi_n^{(i)}) \end{pmatrix} \tag{17}$$

and we approximate the FTLE by

$$\sigma(\mathbf{p}_i, t_0) = \frac{1}{|T|} \ln(\|YX^T(XX^T)^{-1}\|), \tag{18}$$

where $YX^T(XX^T)^{-1}$ is a nonsingular $n \times n$ matrix from which we can extract the strictly positive maximum singular value.

D. Example 1: Rayleigh–Bénard convection cells

To illustrate the use of the algorithm, we consider a two-dimensional model of Rayleigh–Bénard convection cells introduced by Solomon and Gollub.^{56,55} The stream function is given by

$$\psi(x, y, t) = \sin(\pi(x - g(t)))\sin(\pi y). \tag{19}$$

We consider that the lines $y=0$ and $y=1$ are material boundaries and that the stream function delineates an infinite horizontal array of convection cells. For $g(t)=0$, the cells are bounded by vertical segments $x=k \in \mathbb{Z}$ and there is not any




FIG. 3. (Color online) Using elements that are too large can lead to an underestimated FTLE. Starting with a square element at $t=t_0$, the stretching is well captured by the motion of the four grid points at a later time $t=t_1$. Nevertheless, the deformation is nonlinear and the element can fold at a later time $t=t_f$. At this time, the algorithm that uses the position of the four grid points to evaluate FTLE will underestimate the actual stretching.

cell transport. When $g(t)$ is an oscillating function, the boundaries become more complex and transport is possible through the lobes of oscillating LCS.^{5,33,36}

In this paper, we used a quasiperiodic roll motion $g(t)=0.3 \sin(4\pi t)+0.1 \sin(2t)$. The model can be equally used with a fully aperiodic function $g(t)$. For instance, one can define $g(t)$ as a realization of a random process with zero mean and Gaussian covariance in time (see Ref. 36 or <http://www.lekien.com/~francois/software/rndfieldgen>).

To capture long, sharp LCS, one typically needs very dense meshes. Indeed, as illustrated in Fig. 3 for a Cartesian mesh, the FTLE can be severely underestimated when the grid size is too small (or the integration step T is too large). The initial element is stretched in a nonlinear way and the folding of the element can lead to an improper estimate of the stretching since only grid points are used to evaluate the deformation.

To solve this problem and obtain sharp LCS, it is therefore necessary to use elements with very small sizes. Such a high resolution cannot usually be achieved with Cartesian meshes since one would need to refine the entire mesh. An advantage of computing FTLE on an unstructured mesh is the ability to refine the mesh only where needed. In this case, we seek high resolution near the FTLE ridges, but not everywhere. Based on an estimate of the FTLE field on a rough mesh, one can generate a new mesh of triangles and increase the resolution only near the LCS.

Figure 4 illustrates such a refinement procedure for the Rayleigh–Bénard convection cells using refinement rules similar to that of Garth *et al.*^{13,14} The first row of Fig. 4 illustrates the preliminary computation. The domain is covered with a rough mesh containing about 3000 triangles, each with roughly the same size. The upper right panel of Fig. 4 shows the resulting FTLE field for the rough mesh. To ease the analysis of Fig. 4, each triangle is plotted with a uniform color corresponding to the maximum FTLE value of the triangle nodes.

Using the FTLE estimate in the upper right panel of Fig. 4, a more appropriate mesh can be designed. We identify all the nodes of the rough mesh for which the FTLE estimate is above 2.5. All the triangles containing at least one such node are labeled for refinement. We then require a minimum

triangle diameter of 0.02 inside the selected zone. Note that the triangles can still be rather small in regions where the FTLE is below the threshold. The actual triangle size varies depending on the proximity to the refined zone and the properties of the meshing algorithm. The resulting mesh has 8000 triangles and are shown, along with the corresponding FTLE, on the second row of Fig. 4.

Note that the maximum value of the FTLE is higher in the refined mesh. This is a consequence of the saturation phenomena depicted in Fig. 3: The stretching induced by a LCS passing between grid points is underevaluated by the finite mesh size. By decreasing the distance between grid points near the LCS, we refine the estimated FTLE value on the LCS and, hence the maximum FTLE value for the domain. As a result, the right panels of Fig. 4 have different level sets. To compare the FTLE distribution, the ranges are, however, kept constant: When the maximum level set is increased, the minimum level set is increased by the same amount.

Based on the new FTLE estimate in the second row of Fig. 4, we can continue the refinement process. We selected a new FTLE threshold of 2.9 and set a minimum triangle diameter of 0.01 for triangles above the threshold. The resulting mesh is shown on the third row of Fig. 4 and has 20 000 triangles. The resulting high quality FTLE field should be compared to the relatively low quality result that a 100×200 Cartesian mesh would give (for the same number of elements). The last refinement of Fig. 4 uses a threshold of 4.0 and a minimum length scale of 0.004. The final mesh provides very high FTLE details on a mesh with only 55 000 triangles. As a comparison, Fig. 6 of Shadden *et al.*⁵² shows FTLE on a similar system based on a 160×320 Cartesian mesh (51 200 elements). The adaptive mesh obviously provides a tighter estimate of the LCS position.

An advantage of the refinement method described above is that it increases the resolution near the steep FTLE ridge. As shown by Shadden *et al.*,⁵² the flux through a LCS is as small as the ridge is steep (sharp slope in a transverse direction) and the FTLE value is constant (along the ridge). This situation is hardly observable with Cartesian meshes as the resolution of the mesh must be extremely high to avoid spurious FTLE oscillation along the ridge. With the unstructured




FIG. 4. (Color) Computation of FTLE for the aperiodic convection cells of Solomon and Gollub (Ref. 56) using an unstructured mesh and dynamic mesh refinement. The first row is the initial computation (rough estimate). A mesh with 3000 triangles of equal size (upper left panel) is used to evaluate the FTLE (upper right panel). The second row shows the first refinement step. The domain is covered by 8000 triangles whose maximum diameter is set to 0.02 wherever the FTLE estimate (upper right panel) is above 2.5. The third row uses the FTLE estimate from the second row to further refine the mesh. The new mesh has 20 000 triangles and the maximum triangle diameter is set to 0.01 for all the points where the previous FTLE estimate was above 2.9. In the last row, a final mesh of 55 000 triangles is created by setting the maximum diameter to 0.004 for all points, where the previous FTLE estimate is above 4.0. The resulting FTLE field shows a crisp LCS with roughly constant FTLE values (lower right panel).

mesh, we can concentrate small elements along the ridge and keep coarser elements elsewhere.

Note that the method described above is not designed to compute FTLE in any subset of an Euclidean space. It works only for compact submanifolds of \mathbb{R}^n , with the same dimension n . For instance, it cannot be used to compute FTLE on a sphere (codimension 1 in \mathbb{R}^3). In Sec. III, we extend this method to arbitrary manifolds and, hence, make it possible to compute FTLE on regions such as spheres, cylinders, and Möbius bands.

III. FINITE-TIME LYAPUNOV EXPONENTS IN RIEMANNIAN MANIFOLDS

A n -dimensional differentiable manifold \mathcal{M} is a space in which every point has a local neighborhood that looks like \mathbb{R}^n . More specifically there exist diffeomorphisms

$$\beta_k: U_k \subseteq \mathcal{M} \rightarrow \mathbb{R}^n, \tag{20}$$

such that the finite sequence of open sets U_k covers \mathcal{M} . Each couple (U_k, β_k) is called a chart and for any point $\mathbf{p} \in \mathcal{M}$, we

can find at least one diffeomorphism from an open neighborhood of \mathbf{p} to an open set of \mathbb{R}^n . A set of diffeomorphisms β_k whose domains cover \mathcal{M} is called an atlas.

Example: Consider the sphere $\mathcal{M} = \{(x, y, z) \in \mathbb{R}^3 | x^2 + y^2 + z^2 = 1\}$. We consider its partition into the following three overlapping open sets:

$$\begin{aligned}
 U_1 &= \{(x, y, z) \in \mathbb{R}^3 | x^2 + y^2 + z^2 = 1 \text{ and } z > \frac{1}{4}\}, \\
 U_2 &= \{(x, y, z) \in \mathbb{R}^3 | x^2 + y^2 + z^2 = 1 \text{ and } -\frac{1}{2} < z < \frac{1}{2}\}, \\
 U_3 &= \{(x, y, z) \in \mathbb{R}^3 | x^2 + y^2 + z^2 = 1 \text{ and } z < -\frac{1}{4}\},
 \end{aligned}
 \tag{21}$$

and the diffeomorphic maps

$$\begin{aligned}
 \beta_1: U_1 &\rightarrow \mathbb{R}^2: (x, y, z) \mapsto \beta_1(x, y, z) \doteq (x, y), \\
 \beta_2: U_2 &\rightarrow \mathbb{R}^2: (x, y, z) \mapsto \beta_2(x, y, z) \doteq (\text{atan2}(y, x), z), \\
 \beta_3: U_3 &\rightarrow \mathbb{R}^2: (x, y, z) \mapsto \beta_3(x, y, z) \doteq (x, y).
 \end{aligned}
 \tag{22}$$

Each diffeomorphism β_k maps a subset U_k of the manifold to an open set of \mathbb{R}^2 and the union of the subsets U_k is the sphere. This shows that the unit sphere \mathcal{M} is a two-dimensional differentiable manifold.

Let us consider a manifold \mathcal{M} . A dynamical system on \mathcal{M} is defined by a flow $\phi_{t_0}^t: \mathcal{M} \rightarrow \mathcal{M}$ which, for each point $\mathbf{p} \in \mathcal{M}$, returns its image $\mathbf{p}' = \phi_{t_0}^t(\mathbf{p})$. The differential of $\phi_{t_0}^t$ at a point $\mathbf{p} \in \mathcal{M}$ is a linear map from $T_{\mathbf{p}}\mathcal{M}$, the tangent space at point \mathbf{p} , to $T_{\mathbf{p}'}\mathcal{M}$, the tangent space at point $\mathbf{p}' = \phi_{t_0}^t(\mathbf{p})$.

To define FTLE, we need a norm in $T_{\mathbf{p}}\mathcal{M}$ and in $T_{\mathbf{p}'}\mathcal{M}$. The computations in the next sections will require also an inner product in the tangent spaces (to construct orthonormal bases). For this reason, we consider that \mathcal{M} is a Riemannian manifold, i.e., it is a differentiable manifold where each tangent space $T_{\mathbf{p}}\mathcal{M}$ is equipped with an inner product (and an associated norm) that varies smoothly with \mathbf{p} . Note that this is not a limitation as one can always find such a smooth system of inner products for a given differentiable manifold \mathcal{M} . One way to find such an inner product and make a manifold Riemannian is to exploit the Whitney embedding theorem: Any n -dimensional manifold can be embedded in \mathbb{R}^m , provided that $m \geq n$ is large enough.^{62,24} Given an embedding of the manifold in \mathbb{R}^m , the inner product in the tangent spaces can then be derived from the Euclidean inner product in \mathbb{R}^m . Furthermore, we will later show that the choice of the inner product does not influence much the position of the resulting LCS. Changing the metric results in a different FLTE field but, while low FTLE values can change by much, the ridges remain mostly the same.

Given a Riemannian manifold \mathcal{M} , the definition of the FTLE carries over easily. We have

$$\begin{aligned}
 \sigma(\mathbf{p}, t_0) &\doteq \frac{1}{|T|} \ln(\|D\phi_{t_0}^{t_0+T}\|) \\
 &= \frac{1}{|T|} \ln \left(\max_{\substack{\mathbf{u} \in T_{\mathbf{p}}\mathcal{M} \\ \|\mathbf{u}\| \neq 0}} \frac{\|D\phi_{t_0}^{t_0+T}(\mathbf{u})\|_{T_{\mathbf{p}'}\mathcal{M}}}{\|\mathbf{u}\|_{T_{\mathbf{p}}\mathcal{M}}} \right),
 \end{aligned}
 \tag{23}$$

where $\|\cdot\|_{T_{\mathbf{p}}\mathcal{M}}$ and $\|\cdot\|_{T_{\mathbf{p}'}\mathcal{M}}$ are the norms in the tangent spaces $T_{\mathbf{p}}\mathcal{M}$ and $T_{\mathbf{p}'}\mathcal{M}$.

A. FTLE in manifold coordinates

Note that the definition of the FTLE on a manifold does not depend on the atlas (or coordinate system) used. Nevertheless, when one evaluates the FTLE, it is often more convenient to use local coordinates. By the definition of a differentiable manifold, there must be an open set $U_{\mathbf{p}} \subseteq \mathcal{M}$ containing \mathbf{p} on which there exists a diffeomorphism

$$\beta_{\mathbf{p}}: U_{\mathbf{p}} \subseteq \mathcal{M} \rightarrow \mathbb{R}^n.
 \tag{24}$$

Since $\beta_{\mathbf{p}}$ is a diffeomorphism, it has a C^1 inverse

$$\beta_{\mathbf{p}}^{-1}: V_{\mathbf{p}} \subseteq \mathbb{R}^n \rightarrow \mathcal{M}: \boldsymbol{\xi} \mapsto \beta_{\mathbf{p}}^{-1}(\boldsymbol{\xi}),
 \tag{25}$$

where $V_{\mathbf{p}} = \beta_{\mathbf{p}}(U_{\mathbf{p}})$ is an open set. Since \mathbf{p} is in the domain of $\beta_{\mathbf{p}}$, it must belong to the image of $\beta_{\mathbf{p}}^{-1}$. We denote by $\boldsymbol{\xi}^* = (\xi_1^*, \xi_2^*, \dots, \xi_n^*)$ the unique point of $V_{\mathbf{p}} \subseteq \mathbb{R}^n$ satisfying $\beta_{\mathbf{p}}^{-1}(\boldsymbol{\xi}^*) = \mathbf{p}$.

Note that by definition of the diffeomorphism, the derivative $D\beta_{\mathbf{p}}^{-1}$ is continuous, takes values in $T_{\mathbf{p}}\mathcal{M}$, and is full rank. The n vectors

$$\begin{cases}
 \mathbf{e}_1 \doteq D\beta_{\mathbf{p}}^{-1}(1, 0, \dots, 0) = \left. \frac{\partial \beta_{\mathbf{p}}^{-1}}{\partial \xi_1} \right|_{\boldsymbol{\xi} = \boldsymbol{\xi}^*}, \\
 \mathbf{e}_2 \doteq D\beta_{\mathbf{p}}^{-1}(0, 1, \dots, 0) = \left. \frac{\partial \beta_{\mathbf{p}}^{-1}}{\partial \xi_2} \right|_{\boldsymbol{\xi} = \boldsymbol{\xi}^*}, \\
 \vdots \\
 \mathbf{e}_n \doteq D\beta_{\mathbf{p}}^{-1}(0, 0, \dots, 1) = \left. \frac{\partial \beta_{\mathbf{p}}^{-1}}{\partial \xi_n} \right|_{\boldsymbol{\xi} = \boldsymbol{\xi}^*},
 \end{cases}
 \tag{26}$$

can therefore be used as a basis of $T_{\mathbf{p}}\mathcal{M}$. The basis vectors do not necessarily have unit norm and are not necessarily orthogonal. Nevertheless, \mathcal{M} is a Riemannian manifold and there is an inner product in $T_{\mathbf{p}}\mathcal{M}$, which gives the ability to use the Gram-Schmidt process and derive an orthonormal basis $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n\}$ from the basis $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$.

We can proceed similarly for the image $\mathbf{p}' = \phi_{t_0}^{t_0+T}(\mathbf{p})$ and $T_{\mathbf{p}'}\mathcal{M}$, the tangent space at point \mathbf{p}' . We denote by $\{\hat{\mathbf{f}}_1, \hat{\mathbf{f}}_2, \dots, \hat{\mathbf{f}}_n\}$ the orthonormal basis of $T_{\mathbf{p}'}\mathcal{M}$. We then build the $n \times n$ matrix $M = (m_{ij})$ defined by

$$m_{ij} \doteq \hat{\mathbf{f}}_i \cdot (D\phi_{t_0}^{t_0+T}(\hat{\mathbf{e}}_j)).
 \tag{27}$$

Since the two bases have been orthonormalized with the appropriate inner products, the norm of any vector in $T_{\mathbf{p}}\mathcal{M}$ or in $T_{\mathbf{p}'}\mathcal{M}$ can be computed with its coordinates and we get the expected FTLE formula

$$\begin{aligned} \sigma(\mathbf{x}_0, t_0) &= \frac{1}{|T|} \ln \left(\max_{U \in \mathbb{R}_0^n} \frac{\|MU\|}{\|U\|} \right) \\ &= \frac{1}{2|T|} \ln \left(\max_{U \in \mathbb{R}_0^n} \frac{U^T M^T M U}{U^T U} \right) \\ &= \frac{1}{|T|} \ln (\text{largest singular value } (M)). \end{aligned} \tag{28}$$

The only remaining step to compute the FTLE on the Riemannian manifold is therefore the evaluation of the matrix M using the manifold coordinates. Indeed, in practice, one will not compute the actual flow map $\phi_{t_0}^{t_0+T}$ but the ‘‘coordinate flow map’’ defined by

$$\Phi_{t_0}^{t_0+T}: V_{\mathbf{p}} \rightarrow V_{\mathbf{p}'}: \boldsymbol{\xi} \mapsto \boldsymbol{\xi}' \doteq \beta_{\mathbf{p}'}(\phi_{t_0}^{t_0+T}(\beta_{\mathbf{p}}^{-1}(\boldsymbol{\xi}))). \tag{29}$$

In the equation above, one starts with manifold coordinates $\boldsymbol{\xi}$ and computes the corresponding manifold point \mathbf{p} using the chart $\beta_{\mathbf{p}}^{-1}$. The flow $\phi_{t_0}^{t_0+T}$ is applied to \mathbf{p} which returns the image \mathbf{p}' . We then get the coordinates $\boldsymbol{\xi}'$ of the image using the chart $\beta_{\mathbf{p}'}$ (which might be different from the first chart $\beta_{\mathbf{p}}$). Combining all the steps above gives a ‘‘coordinate map’’ $\Phi_{t_0}^{t_0+T}$ that returns the final coordinates $\boldsymbol{\xi}'$ as a function of the initial coordinates $\boldsymbol{\xi}$. The map derivative $D\Phi_{t_0}^{t_0+T}$ is the most readily available deformation tensor as it only uses coordinates. How can we derive the matrix M and the FTLE from $D\Phi_{t_0}^{t_0+T}$?

Note that by definition, we have $\phi_{t_0}^{t_0+T} = \beta_{\mathbf{p}'}^{-1} \circ \Phi_{t_0}^{t_0+T} \circ \beta_{\mathbf{p}}$, hence

$$D\phi_{t_0}^{t_0+T} = D\beta_{\mathbf{p}'}^{-1} \Big|_{\boldsymbol{\xi}'} \circ D\Phi_{t_0}^{t_0+T} \Big|_{\boldsymbol{\xi}} \circ D\beta_{\mathbf{p}} \Big|_{\mathbf{p}}. \tag{30}$$

The Gram–Schmidt process was used to orthonormalize the column vectors of $D\beta_{\mathbf{p}}^{-1}$ and is, hence, equivalent to the QR decomposition

$$D\beta_{\mathbf{p}}^{-1} = Q_{\mathbf{p}} \circ R_{\mathbf{p}}, \tag{31}$$

where $Q_{\mathbf{p}}$ is a unitary linear operator (that is $Q_{\mathbf{p}}: \mathbb{R}^n \rightarrow T_{\mathbf{p}}\mathcal{M}$ such that $\forall \boldsymbol{\xi}, \boldsymbol{\eta} \in \mathbb{R}^n: Q_{\mathbf{p}}(\boldsymbol{\xi}) \cdot Q_{\mathbf{p}}(\boldsymbol{\eta}) = \boldsymbol{\xi} \cdot \boldsymbol{\eta}$) and $R_{\mathbf{p}}$ is an upper diagonal matrix (linear operator from $\mathbb{R}^n \rightarrow \mathbb{R}^n$). From this point of view, the Gram–Schmidt process is equivalent to defining the new basis by

$$\forall j = 1, \dots, n: \hat{\mathbf{e}}_j = Q_{\mathbf{p}}(\mathbf{1}_j), \tag{32}$$

where $\mathbf{1}_j$ denotes the vector of \mathbb{R}^n that has all components equal to zero except the j th that is equal to one.

Similarly, in $T_{\mathbf{p}'}\mathcal{M}$, the Gram–Schmidt process is equivalent to the QR decomposition

$$D\beta_{\mathbf{p}'}^{-1} = Q_{\mathbf{p}'} \circ R_{\mathbf{p}'}, \tag{33}$$

where $Q_{\mathbf{p}'}$ is unitary and $R_{\mathbf{p}'}$ is the upper diagonal. The orthonormal basis in $T_{\mathbf{p}'}\mathcal{M}$ is then given by

$$\forall i = 1, \dots, n: \hat{\mathbf{f}}_i = Q_{\mathbf{p}'}(\mathbf{1}_i). \tag{34}$$

Direct computation gives

$$D\phi_{t_0}^{t_0+T}(\hat{\mathbf{e}}_j) = D\beta_{\mathbf{p}'}^{-1} \Big|_{\boldsymbol{\xi}'} \circ D\Phi_{t_0}^{t_0+T} \Big|_{\boldsymbol{\xi}} (R_{\mathbf{p}}^{-1} \mathbf{1}_j) \tag{35}$$

and

$$m_{ij} = \hat{\mathbf{f}}_i \cdot (D\phi_{t_0}^{t_0+T}(\hat{\mathbf{e}}_j)) = \mathbf{1}_i \cdot R_{\mathbf{p}'} D\Phi_{t_0}^{t_0+T} \Big|_{\boldsymbol{\xi}} R_{\mathbf{p}}^{-1} \mathbf{1}_j. \tag{36}$$

As a result, the matrix M is given by

$$M = R_{\mathbf{p}'} D\Phi_{t_0}^{t_0+T} \Big|_{\boldsymbol{\xi}} R_{\mathbf{p}}^{-1} = R_{\mathbf{p}'} \begin{pmatrix} \frac{\partial \xi'_1}{\partial \xi_1} & \frac{\partial \xi'_1}{\partial \xi_2} & \dots & \frac{\partial \xi'_1}{\partial \xi_n} \\ \frac{\partial \xi'_2}{\partial \xi_1} & \frac{\partial \xi'_2}{\partial \xi_2} & \dots & \frac{\partial \xi'_2}{\partial \xi_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \xi'_n}{\partial \xi_1} & \frac{\partial \xi'_n}{\partial \xi_2} & \dots & \frac{\partial \xi'_n}{\partial \xi_n} \end{pmatrix} R_{\mathbf{p}}^{-1}. \tag{37}$$

The equation above states that we can compute the FTLE in coordinates. From the ‘‘coordinate deformation tensor,’’ $D\Phi_{t_0}^{t_0+T}$, one needs to multiply by $R_{\mathbf{p}'}$ on the left and by $R_{\mathbf{p}}^{-1}$ on the right to obtain the matrix M whose largest singular value gives the FTLE. Note that the matrix $R_{\mathbf{p}}$ is obtained from the QR decomposition of the manifold charts. As shown by Dieci and Eirola,¹⁰ the matrix $R_{\mathbf{p}}$ must then vary smoothly as a function of \mathbf{p} .

B. Orthogonal coordinates

The equation above gives the matrix M from the coordinate deformation tensor but it requires the QR decomposition of the derivative of the coordinate map $D\beta^{-1}$. If the manifold coordinates are orthogonal, the vectors in the basis $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$ are mutually orthogonal and the matrix $R_{\mathbf{p}}$ is diagonal. We have

$$M = \begin{pmatrix} \gamma_1 \frac{\partial \xi'_1}{\partial \xi_1} \delta_1^{-1} & \gamma_1 \frac{\partial \xi'_1}{\partial \xi_2} \delta_2^{-1} & \dots & \gamma_1 \frac{\partial \xi'_1}{\partial \xi_n} \delta_n^{-1} \\ \gamma_2 \frac{\partial \xi'_2}{\partial \xi_1} \delta_1^{-1} & \gamma_2 \frac{\partial \xi'_2}{\partial \xi_2} \delta_2^{-1} & \dots & \gamma_2 \frac{\partial \xi'_2}{\partial \xi_n} \delta_n^{-1} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_n \frac{\partial \xi'_n}{\partial \xi_1} \delta_1^{-1} & \gamma_n \frac{\partial \xi'_n}{\partial \xi_2} \delta_2^{-1} & \dots & \gamma_n \frac{\partial \xi'_n}{\partial \xi_n} \delta_n^{-1} \end{pmatrix}, \tag{38}$$

with $\delta_i = \|\partial \beta_{\mathbf{p}}^{-1} / \partial \xi_i\|$ and $\gamma_i = \|\partial \beta_{\mathbf{p}'}^{-1} / \partial \xi'_i\|$.

C. Orthonormal coordinates

If the coordinates are orthonormal, we have $M = D\Phi_{t_0}^{t_0+T}$. In other words, the FTLE can be computed directly as the largest singular value of the coordinate deformation tensor. We will later use this property to derive an efficient algorithm for FTLE computation.

D. Example 2: Convection cells on a torus

We consider a system similar to the example in Sec. II D but instead of creating an infinite array of plane convection cells, we wrap four cells around a cylinder. More specifically, we consider the cylinder parametrized by

$$\beta^{-1}: \left[0, 4\pi \right] \times \left[-\frac{\pi}{2}, \frac{\pi}{2} \right] \rightarrow \mathbb{R}^3$$

$$(\xi_1, \xi_2) \mapsto \left(\cos \frac{\xi_1}{2}, \sin \frac{\xi_1}{2}, \frac{\xi_2}{\pi} \right) \tag{39}$$

and the same stream function as in Sec. II D that is

$$\psi(\xi_1, \xi_2, t) = \sin(\xi_1 - g(t)) \sin\left(\xi_2 + \frac{\pi}{2}\right). \tag{40}$$

In other words, for each point (x, y, z) on the cylinder, one can use the diffeomorphism β to compute the manifold coordinates (ξ_1, ξ_2) , then compute the stream function or the velocity field using the formula above where the coordinates x and y of Solomon and Gollub have been replaced by ξ_1 and ξ_2 .

Since we have the parametrization of the manifold, we can determine analytically the local basis. For given ξ_1 and ξ_2 , we have

$$\mathbf{e}_1 = D\beta^{-1}(1, 0) = \frac{1}{2} \left(-\sin \frac{\xi_1}{2}, \cos \frac{\xi_1}{2}, 0 \right) \tag{41}$$

and

$$\mathbf{e}_2 = D\beta^{-1}(0, 1) = \left(0, 0, \frac{1}{\pi} \right). \tag{42}$$

The vectors above are orthogonal, hence we can compute FTLE using Eq. (38) with $\delta_1 = \gamma_1 = \|\mathbf{e}_1\| = \frac{1}{2}$ and $\delta_2 = \gamma_2 = \|\mathbf{e}_2\| = \frac{1}{\pi}$.

If (ξ'_1, ξ'_2) represents the coordinates of a particle that started at time t_0 at position (ξ_1, ξ_2) , the deformation in “manifold coordinates” reads

$$D\Phi_{t_0}^{t_0+T} = \begin{pmatrix} \frac{\partial \xi'_1}{\partial \xi_1} & \frac{\partial \xi'_1}{\partial \xi_2} \\ \frac{\partial \xi'_2}{\partial \xi_1} & \frac{\partial \xi'_2}{\partial \xi_2} \end{pmatrix}. \tag{43}$$

According to Sec. III B, the FTLE is given by the largest singular value of the matrix

$$M = \begin{pmatrix} \frac{\partial \xi'_1}{\partial \xi_1} & \frac{\pi}{2} \frac{\partial \xi'_1}{\partial \xi_2} \\ \frac{2}{\pi} \frac{\partial \xi'_2}{\partial \xi_1} & \frac{\partial \xi'_2}{\partial \xi_2} \end{pmatrix}. \tag{44}$$

To compute the FTLE, we can then cover the manifold with a mesh of points, each with given ξ_1 and ξ_2 coordinates. We

then differentiate numerically the final positions with respect to the initial positions as if the mesh was in \mathbb{R}^2 . The only differences with a computation in Euclidean space are the factors $\pi/2$ and $2/\pi$ that scale the off-diagonal elements of M . The FTLE and corresponding LCS on the cylinder are shown in Fig. 5.

E. Example 3: Convection cells on a Möbius strip

Instead of wrapping an even number of convection cells around a cylinder as in the previous example, we can arrange an odd number of cells into a continuous velocity field on a Möbius strip. This construction does not correspond to a known physical setting but provides a test case to illustrate and test the methods developed above, as well as a first instance of an FTLE field for a non-orientable manifold. We consider the Möbius strip parametrized by

$$\beta^{-1}: [0, 3\pi] \times \left[-\frac{\pi}{2}, \frac{\pi}{2} \right] \rightarrow \mathbb{R}^3$$

$$(\xi_1, \xi_2) \mapsto \begin{pmatrix} \left(1 + \frac{\xi_2}{\pi} \cos \frac{\xi_1}{3} \right) \cos \frac{2\xi_1}{3} \\ \left(1 + \frac{\xi_2}{\pi} \cos \frac{\xi_1}{3} \right) \sin \frac{2\xi_1}{3} \\ \frac{\xi_2}{\pi} \sin \frac{\xi_1}{3} \end{pmatrix} \tag{45}$$

and the same streamfunction as in the previous example. Direct computation shows that the coordinates are also orthogonal. In this case, we have $\delta_2 = \gamma_2 = \|\mathbf{e}_2\| = 1/\pi$. The norm of \mathbf{e}_1 depends on the position on the manifold and we have

$$\delta_1 = \|\mathbf{e}_1(\xi_1, \xi_2)\| = \frac{\xi_2^2}{9\pi^2} + \frac{4}{9} \left(1 + \frac{\xi_2}{\pi} \cos \frac{\xi_1}{3} \right)^2 \tag{46}$$

and

$$\gamma_1 = \|\mathbf{e}_1(\xi'_1, \xi'_2)\| = \frac{\xi_2'^2}{9\pi^2} + \frac{4}{9} \left(1 + \frac{\xi_2'}{\pi} \cos \frac{\xi_1'}{3} \right)^2. \tag{47}$$

As a result, the deformation can be computed in manifold coordinates ξ_1 and ξ_2 and the FTLE is given as the singular value of

$$M = \begin{pmatrix} \frac{\xi_2'^2}{9\pi^2} + \frac{4}{9} \left(1 + \frac{\xi_2'}{\pi} \cos \frac{\xi_1'}{3} \right)^2 \frac{\partial \xi'_1}{\partial \xi_1} & \left(\frac{\xi_2'^2}{9\pi^2} + \frac{4}{9} \left(1 + \frac{\xi_2'}{\pi} \cos \frac{\xi_1'}{3} \right)^2 \right) \frac{\partial \xi'_1}{\partial \xi_2} \pi \\ \frac{\xi_2^2}{9\pi^2} + \frac{4}{9} \left(1 + \frac{\xi_2}{\pi} \cos \frac{\xi_1}{3} \right)^2 \frac{\partial \xi_1}{\partial \xi_1} & \frac{\partial \xi_1}{\partial \xi_2} \pi \\ \frac{1}{\pi} \frac{\xi_2'^2}{9\pi^2} + \frac{4}{9} \left(1 + \frac{\xi_2'}{\pi} \cos \frac{\xi_1'}{3} \right)^2 \frac{\partial \xi_2'}{\partial \xi_1} & \frac{\partial \xi_2'}{\partial \xi_2} \end{pmatrix}. \tag{48}$$




FIG. 5. (Color) FTLE on a cylinder for a positive integration time (upper panel) and a negative integration time (middle panel). Corresponding attractive and repulsive LCS (lower panel) (enhanced online). [URL: <http://dx.doi.org/10.1063/1.3278516.1>]

Figure 6 shows the corresponding FTLE for positive integration time, as well as the attractive and repulsive LCS. Note that the Möbius strip is not an orientable surface. FTLE and LCS are well defined on any Riemannian differentiable manifold, whether orientable or not.

IV. FINITE-TIME LYAPUNOV EXPONENTS ON DIFFERENTIABLE MANIFOLDS EMBEDDED IN R^m

The materials in Sec. III can, in principle, be used to compute FTLE in any Riemannian differentiable manifold, but it requires knowledge of an explicit atlas. In practice, however, many manifolds of interest are only described numerically. The maps in the atlas, and hence their derivatives, might not be known analytically. For this reason, this section presents a more general method for computing FTLE on differentiable manifolds.

The rationale for the method is that any differentiable manifold of dimension n can be *embedded* in the Euclidian




FIG. 6. (Color) FTLE (upper panel) LCS (lower panel) on a Möbius strip (enhanced online). [URL: <http://dx.doi.org/10.1063/1.3278516.2>]

space R^m , provided that m is sufficiently large.^{62,24} If we begin with a Riemannian manifold, that is a manifold together with an inner product at each point, Nash and Kuiper’s embedding theorem, combined with Whitney’s embedding theorem, guarantees that the manifold can be isometrically embedded in R^m .^{41,32} The inner product on the manifold then corresponds to the Euclidian inner product in R^m . In other words, it is sufficient to consider only manifolds embedded in R^m .

A. Finding the local basis

Consider the node \mathbf{p}_i of the unstructured mesh. As shown in Fig. 7, we seek the effect of the flow derivative $D\phi_{t_0}^{t_0+T}$ on vectors in the tangent space $T_{\mathbf{p}_i}\mathcal{M}$ but we only have information about the motion of the mesh node \mathbf{p}_i and its neighbors $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_j, \dots, \mathbf{p}_N$ that are all located on the manifold \mathcal{M} .

We consider that the manifold \mathcal{M} is embedded in R^m and, accordingly, we view the node \mathbf{p}_i and its neighbors $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_j, \dots, \mathbf{p}_N$ as vectors of R^m . The objective is to determine a basis of $T_{\mathbf{p}_i}\mathcal{M}$ (seen as embedded in $T_{\mathbf{p}_i}R^m = R^m$). First we consider the N vectors $\mathbf{u}_j = \mathbf{p}_j - \mathbf{p}_i$ ($j=1, \dots, N$). If the grid size was infinitesimally small, the N vectors of \mathbf{u}_j would exactly span the n -dimensional manifold $T_{\mathbf{p}_i}\mathcal{M}$. In this case, we could find an orthonormal basis of $T_{\mathbf{p}_i}\mathcal{M}$ using the Gram–Schmidt process. Starting with the N vectors $\mathbf{u}_j \in R^m$, we would obtain n nonzero orthonormal vectors \mathbf{e}_k (since $T_{\mathbf{p}_i}\mathcal{M}$ has dimension n , during the Gram–Schmidt process, $N-n$ vectors would vanish and be eliminated).




FIG. 7. Approximation of the deformation tensor (map from $T_{p_i}\mathcal{M}$ to $T_{p_j}\mathcal{M}$) using an unstructured mesh on a manifold \mathcal{M} embedded in \mathbb{R}^m .

When the grid size is finite, the N vectors \mathbf{u}_j are not exactly lying in $T_{p_i}\mathcal{M}$. An orthonormal set of vectors that approximately spans $T_{p_i}\mathcal{M}$ can, however, still be obtained using a variant of the Gram–Schmidt process.

Let us begin by normalizing the N vectors \mathbf{u}_j and defining $\hat{\mathbf{u}}_j = \mathbf{u}_j / \|\mathbf{u}_j\|$. In a regular Gram–Schmidt, the first vector of the orthonormal basis is arbitrary. To reconstruct the local basis from the N vectors $\hat{\mathbf{u}}_j$, we need to avoid the “directions pointing out of $T_{p_i}\mathcal{M}$.” More specifically, for each basis vector candidate $\hat{\mathbf{u}}_j$, we consider how well the other vectors also align in this direction. The vector $\hat{\mathbf{u}}_k - (\hat{\mathbf{u}}_k \cdot \hat{\mathbf{u}}_j)\hat{\mathbf{u}}_j$ represents the portion of $\hat{\mathbf{u}}_k$ that cannot be captured by a projection on $\hat{\mathbf{u}}_j$. The quantity

$$\sum_{k=1}^N \|\hat{\mathbf{u}}_k - (\hat{\mathbf{u}}_k \cdot \hat{\mathbf{u}}_j)\hat{\mathbf{u}}_j\|^2 \tag{49}$$

quantifies how well the set of vectors $\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_2, \dots, \hat{\mathbf{u}}_N$ can be represented by the single vector $\hat{\mathbf{u}}_j$. Accordingly, we define the first basis vector as

$$\tilde{\mathbf{u}}_1 = \arg \min_{\hat{\mathbf{u}}_j} \left(\sum_{k=1}^N \|\hat{\mathbf{u}}_k - (\hat{\mathbf{u}}_k \cdot \hat{\mathbf{u}}_j)\hat{\mathbf{u}}_j\|^2 \right). \tag{50}$$

To get a second basis vector, we proceed by recursion. We have the direct sum

$$T_{p_i}\mathcal{M} = \text{span}\{\tilde{\mathbf{u}}_1\} \otimes T_{n-1}, \tag{51}$$

where T_{n-1} is a subspace of dimension $n-1$ that best fits the remaining $N-1$ vectors $\hat{\mathbf{u}}_j$ (all but the one that was selected as $\tilde{\mathbf{u}}_1$). As a result, we can subtract the projection on $\tilde{\mathbf{u}}_1$ from each vector $\hat{\mathbf{u}}_j$ and consider the new set of vectors

$$\hat{\mathbf{u}}_j^{(2)} = \begin{cases} 0, & \text{if } \|\hat{\mathbf{u}}_j - (\hat{\mathbf{u}}_j \cdot \tilde{\mathbf{u}}_1)\tilde{\mathbf{u}}_1\| = 0, \\ \frac{\hat{\mathbf{u}}_j - (\hat{\mathbf{u}}_j \cdot \tilde{\mathbf{u}}_1)\tilde{\mathbf{u}}_1}{\|\hat{\mathbf{u}}_j - (\hat{\mathbf{u}}_j \cdot \tilde{\mathbf{u}}_1)\tilde{\mathbf{u}}_1\|}, & \text{otherwise.} \end{cases} \tag{52}$$

The second basis vector is then selected as

$$\tilde{\mathbf{u}}_2 = \arg \min_{\hat{\mathbf{u}}_j^{(2)}} \left(\sum_{k=1}^N \|\hat{\mathbf{u}}_k^{(2)} - (\hat{\mathbf{u}}_k^{(2)} \cdot \hat{\mathbf{u}}_j^{(2)})\hat{\mathbf{u}}_j^{(2)}\|^2 \right). \tag{53}$$

Note that the search above should be restricted to only $N-1$ indices k and exclude the index corresponding to the first basis vector. Nevertheless the first basis vector leads to a vanishing $\hat{\mathbf{u}}_j^{(2)}$ and will never correspond to the minimum in the equation above. As a result, we can keep the N candidates in the formula above.

Applying the method recursively, we derive the n step algorithm below.

Initialization:

$$\forall j = 1, \dots, N: \hat{\mathbf{u}}_j^{(1)} = \frac{\mathbf{u}_j}{\|\mathbf{u}_j\|}.$$

Loop for $i = 1, \dots, n$:

$$\tilde{\mathbf{u}}_i = \arg \min_{\hat{\mathbf{u}}_j^{(i)}} \left(\sum_{k=1}^N \|\hat{\mathbf{u}}_k^{(i)} - (\hat{\mathbf{u}}_k^{(i)} \cdot \hat{\mathbf{u}}_j^{(i)})\hat{\mathbf{u}}_j^{(i)}\|^2 \right),$$

$$\forall j = 1, \dots, N:$$

$$\hat{\mathbf{u}}_j^{(i+1)} = \begin{cases} 0, & \text{if } \|\hat{\mathbf{u}}_j^{(i)} - (\hat{\mathbf{u}}_j^{(i)} \cdot \tilde{\mathbf{u}}_i) \tilde{\mathbf{u}}_i\| = 0, \\ \frac{\hat{\mathbf{u}}_j^{(i)} - (\hat{\mathbf{u}}_j^{(i)} \cdot \tilde{\mathbf{u}}_i) \tilde{\mathbf{u}}_i}{\|\hat{\mathbf{u}}_j^{(i)} - (\hat{\mathbf{u}}_j^{(i)} \cdot \tilde{\mathbf{u}}_i) \tilde{\mathbf{u}}_i\|}, & \text{otherwise.} \end{cases}$$

Termination: After the n steps, the set $\{\tilde{\mathbf{u}}_1, \tilde{\mathbf{u}}_2, \dots, \tilde{\mathbf{u}}_n\}$ is an orthonormal basis that spans $T_{\mathbf{p}_i} \mathcal{M}$ in the limit of infinitesimal mesh elements. The residual error

$$\rho = \sum_{k=1}^N \|\hat{\mathbf{u}}_k^{(n+1)}\|^2 \tag{54}$$

represents how well the n vectors of the orthonormal basis capture the N input vectors. ρ vanishes when the number of neighbors N is equal to n or when the N vectors \mathbf{u}_j all lie in a n -dimensional subspace. If the residual error ρ is too high, the mesh covering the manifold must be refined.

Given the N neighbors $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N$ of a node \mathbf{p}_i , the algorithm above provides an orthonormal basis of an approximation of $T_{\mathbf{p}_i} \mathcal{M}$. According to Sec. III C, since the basis is orthonormal, the FTLE can be computed directly in coordinates. We can then process the data gathered on the unstructured mesh using a method similar to that described in Sec. II C. In this case the local coordinates are given by the projection of the vectors $\mathbf{p}_j - \mathbf{p}_i$ on the basis vectors. We write

$$\forall j = 1, \dots, N: \quad \forall k = 1, \dots, n: \quad \xi_k^{(j)} = (\mathbf{p}_j - \mathbf{p}_i) \cdot \tilde{\mathbf{u}}^{(k)} \tag{55}$$

and the matrix of initial coordinates is given by

$$X = \begin{pmatrix} \frac{A_{1i}}{8\mathcal{V}_i} \xi_1^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \xi_1^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \xi_1^{(N)} \\ \frac{A_{1i}}{8\mathcal{V}_i} \xi_2^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \xi_2^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \xi_2^{(N)} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{A_{1i}}{8\mathcal{V}_i} \xi_n^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \xi_n^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \xi_n^{(N)} \end{pmatrix}, \tag{56}$$

where A_{ji} is the hypervolume of all the n -tetrahedra that contain both \mathbf{p}_i and \mathbf{p}_j and \mathcal{V}_i is the hypervolume of the Voronoi cell of node \mathbf{p}_i (as in Sec. II C).

Let us denote the image of \mathbf{p}_j by \mathbf{p}'_j . We can consider the N vectors $\mathbf{u}'_j = \mathbf{p}'_j - \mathbf{p}'_i$ and apply the variant of the Gram-Schmidt to obtain an orthonormal basis $\{\tilde{\mathbf{u}}'_1, \tilde{\mathbf{u}}'_2, \dots, \tilde{\mathbf{u}}'_n\}$ of $T_{\mathbf{p}'_i} \mathcal{M}$. Defining

$$\forall j = 1, \dots, N: \quad \forall k = 1, \dots, n: \quad \eta_k^{(j)} = (\mathbf{p}'_j - \mathbf{p}'_i) \cdot \tilde{\mathbf{u}}'^{(k)}, \tag{57}$$

and the matrix of initial coordinates is given by

$$Y = \begin{pmatrix} \frac{A_{1i}}{8\mathcal{V}_i} \eta_1^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \eta_1^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \eta_1^{(N)} \\ \frac{A_{1i}}{8\mathcal{V}_i} \eta_2^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \eta_2^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \eta_2^{(N)} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{A_{1i}}{8\mathcal{V}_i} \eta_n^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \eta_n^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \eta_n^{(N)} \end{pmatrix}. \tag{58}$$

Since the two basis are orthonormal, the FTLE at point \mathbf{p}_i is approximated by the same formula as in Sec. II C, that is

$$\begin{aligned} \sigma(\mathbf{p}_i, t_0) &= \frac{1}{|T|} \ln(\|YX^T(XX^T)^{-1}\|) \\ &= \frac{1}{|T|} \ln(\text{largest singular value}(YX^T(XX^T)^{-1})). \end{aligned} \tag{59}$$

B. Exploiting the embedding

The algorithm described in Sec. IV A evaluates FTLE on any n -dimensional manifold embedded in \mathbb{R}^m . Reconstructing an orthonormal basis in $T_{\mathbf{p}_i} \mathcal{M}$ using the Gram-Schmidt variant is usually straightforward provided that the mesh elements are small enough and well conditioned. If the residual error is too large, it is always possible to refine the mesh covering the manifold.

The same remark does not apply to the reconstruction of the local basis in $T_{\mathbf{p}'_i} \mathcal{M}$. Indeed, the nodes \mathbf{p}'_j are defined as the advection of the points \mathbf{p}_j . The triangles can be strongly deformed by the flow and the vectors $\mathbf{u}'_j = \mathbf{p}'_j - \mathbf{p}'_i$ might not be small. The LCS are characterized by locally high values of FTLE and stretching. Consequently, the deformation of the mesh is high in the very regions where we need the highest accuracy.

It is possible, however, to skip the reconstruction of the basis in $T_{\mathbf{p}'_i} \mathcal{M}$ and replace the $n \times N$ matrix Y of Sec. IV A by the $m \times N$ matrix containing the coordinates in \mathbb{R}^m (in which the manifold is embedded). More specifically, we consider the coordinates $(\tilde{\eta}_1^{(j)}, \tilde{\eta}_2^{(j)}, \dots, \tilde{\eta}_m^{(j)})$ of the vectors $\mathbf{u}'_j = \mathbf{p}'_j - \mathbf{p}'_i$ and we do not project on a basis of $T_{\mathbf{p}'_i} \mathcal{M}$. The corresponding $n \times N$ matrix is given by

$$\tilde{Y} = \begin{pmatrix} \frac{A_{1i}}{8\mathcal{V}_i} \tilde{\eta}_1^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \tilde{\eta}_1^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \tilde{\eta}_1^{(N)} \\ \frac{A_{1i}}{8\mathcal{V}_i} \tilde{\eta}_2^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \tilde{\eta}_2^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \tilde{\eta}_2^{(N)} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{A_{1i}}{8\mathcal{V}_i} \tilde{\eta}_m^{(1)} & \frac{A_{2i}}{8\mathcal{V}_i} \tilde{\eta}_m^{(2)} & \dots & \frac{A_{Ni}}{8\mathcal{V}_i} \tilde{\eta}_m^{(N)} \end{pmatrix}. \tag{60}$$

Since the n -dimensional manifold is embedded in \mathbb{R}^m , we have $m \geq n$ and the Whitney embedding theorem guarantees that we can select $m \leq 2n$.^{62,24} The $m \times N$ matrix \tilde{Y} is therefore larger than the $n \times N$ matrix Y . Nevertheless, the matrix \tilde{Y} does not require the computation of the local basis at the final position \mathbf{p}'_i and is simpler to compute. In addition, using

\tilde{Y} leads to a more precise FTLE value as it does not rely on the reconstruction of the local basis in $T_{\mathbf{p}'_i}\mathcal{M}$, which can be imprecise due to the stretching of the mesh near the LCS.

The computation of the FTLE based on \tilde{Y} is identical to the computation based on Y . The FTLE is given by

$$\sigma(\mathbf{p}_i, t_0) = \frac{1}{|T|} \ln (\text{largest singular value } (\tilde{Y}X^\top (XX^\top)^{-1})), \tag{61}$$

where we substituted \tilde{Y} for Y .

To prove the formula above, recall that \mathcal{M} is embedded in \mathbb{R}^m . As a result, the tangent space can be decomposed as

$$\mathbb{R}^m = T_{\mathbf{p}'_i}\mathcal{M} \otimes N_{\mathbf{p}'_i}\mathcal{M}, \tag{62}$$

where $N_{\mathbf{p}'_i}\mathcal{M}$ is the normal space. Y is the matrix representation of $D\phi_{t_0}^{t_0+T}: T_{\mathbf{p}'_i}\mathcal{M} \rightarrow T_{\mathbf{p}'_i}\mathcal{M}$. In the limit of infinitesimally small elements, \tilde{Y} is the matrix representation of a map $\mathbb{R}^m \rightarrow \mathbb{R}^m$ that corresponds to the direct sum

$$D\phi_{t_0}^{t_0+T} \otimes O_{m-n}, \tag{63}$$

where $O_{m-n}: T_{\mathbf{p}'_i}\mathcal{M} \rightarrow N_{\mathbf{p}'_i}\mathcal{M}$ is the null operator. For a mesh whose elements have a finite diameter h , the null operator must be replaced by hW , with $\lim_{h \rightarrow 0} hW = O_{m-n}$. As a result, there must be a unitary $m \times m$ matrix R such that

$$R\tilde{Y}R^\top = \begin{pmatrix} Y \\ hW \end{pmatrix}. \tag{64}$$

Let us define

$$M = YX^\top (XX^\top)^{-1} \tag{65}$$

and

$$\tilde{M} = \tilde{Y}X^\top (XX^\top)^{-1} = R^\top \begin{pmatrix} YX^\top (XX^\top)^{-1} \\ hWX^\top (XX^\top)^{-1} \end{pmatrix} R. \tag{66}$$

Direct computation gives

$$R\tilde{M}^\top \tilde{M}R^\top = M^\top M + h^2 (XX^\top)^{-1} XW^\top WX^\top (XX^\top)^{-1}, \tag{67}$$

hence $\tilde{M}^\top \tilde{M}$ and $M^\top M$ have the same eigenvalues up to a term that vanishes as the square of the maximum element diameter in the mesh (i.e., proportional to the element area). Provided that the mesh is sufficiently fine, we do not need to reconstruct the local basis at the point \mathbf{p}'_i . We can just list the coordinates in \mathbb{R}^m and use \tilde{Y} instead of Y in the FTLE formula.

Note that the same trick does not apply to the initial position. We cannot avoid the reconstruction of the local basis at \mathbf{p}'_i and list the Cartesian coordinates in a $m \times N$ matrix \tilde{X} . Indeed the product $\tilde{X}\tilde{X}^\top$ would have $(m-n)$ eigenvalues then vanishes with h and it would be close to singular.

It is also worth noting that the product $M^\top M$ we use in the definition of the FTLE is an approximation of the *right* finite-time Cauchy–Green deformation tensor. Although this option has not been exploited yet, the FTLE can also be derived using the *left* Cauchy–Green deformation tensor

MM^\top .⁵² The same technique described in this section would apply to the left Cauchy–Green deformation tensor. Indeed, we have

$$\begin{aligned} \tilde{M}\tilde{M}^\top &= \begin{pmatrix} MM^\top & hM(XX^\top)^{-1}XW^\top \\ hWX^\top(XX^\top)^{-1}M^\top & h^2WX^\top(XX^\top)^{-2}XW^\top \end{pmatrix} \\ &= \begin{pmatrix} MM^\top & hMP^\top \\ hPM^\top & h^2PP^\top \end{pmatrix}, \end{aligned} \tag{68}$$

with $P = WX^\top(XX^\top)^{-1}$. Compared to MM^\top , the matrix $\tilde{M}\tilde{M}^\top$ has $m-n$ extra eigenvalues but for sufficiently small h , they are all smaller than the eigenvalues of MM^\top . Using the identity

$$\det \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \det(A) \det(D - CA^{-1}B), \tag{69}$$

we can then show that the matching eigenvalues of $\tilde{M}\tilde{M}^\top$ and MM^\top differ by a term that vanishes as h^2 . As a result, we have also

$$\begin{aligned} \text{largest singular value } (\tilde{M}) &= \text{largest singular value } (M) \\ &\quad + \mathcal{O}(h^2). \end{aligned} \tag{70}$$

C. Example 4: The Antarctic polar vortex splitting event of 2002

In late September of 2002, over a period of a few days, the Antarctic ozone hole split in two, with one of the daughter fragments subsequently reasserting its position over the pole while the other spread into the midlatitudes. The split is clearly evident in the ozone concentrations shown in Fig. 8 while the splitting/reformation sequence is shown in Fig. 9. A split ozone hole implies a split vortex and thus a sudden stratospheric warming.⁵⁴

These types of warming occur in roughly half of all winters in the Arctic²³ and are thought to be produced by the dynamical momentum forcing resulting from the breaking and dissipation of planetary-scale Rossby waves in the stratosphere. Prior to 2002, however, no stratospheric sudden warming had been observed in the Antarctic, where reliable records go back half a century. The relatively early warming (in September) followed by a reforming of the circumpolar vortex was an unprecedented event. Because of the importance of the coherence of the stratospheric polar vortex for stratospheric ozone depletion, the event made it onto the front pages of newspapers worldwide and became a focus of the atmospheric science community.⁶¹ Figure 8 shows the Antarctic ozone hole in September 2001 (1 year before the split), September 2002 (split ozone hole), and September 2003 (1 year after the split).

We view the event as one that caused a clear and identifiable topological bifurcation in the stratospheric flow field. Thus it serves as a prototype atmospheric event on which to study the effects of topological bifurcations on mixing and transport of atmospheric tracers. Figure 9 shows the LCS associated with the splitting event based on the 650 K isentrope of the National Centers for Environmental Prediction/ National Center for Atmospheric Research (NCEP/NCAR)




FIG. 8. (Color) The splitting of the Antarctic ozone hole, September 2002, as evident in total column ozone concentrations. September 2001 and 2003 are shown for comparison (from Ref. 61). An example of the computed LCS around the time of the splitting and reformation event is shown in Fig. 9 (enhanced online). [URL: <http://dx.doi.org/10.1063/1.3278516.3>]

reanalysis data,²⁸ an analysis (and forecast) system performing data assimilation using past data from 1948 to the present.

To compute the FTLE, we consider the projection of the points of the Southern hemisphere on the tangent plane at the equator. In other words, for each point $\mathbf{p}=(x,y,z)$ in the Southern hemisphere, we consider the manifold coordinates (x,y) and the chart

$$\beta^{-1}: D \rightarrow \mathcal{M}: (x,y) \mapsto \begin{pmatrix} x \\ y \\ -\sqrt{R^2-x^2-y^2} \end{pmatrix}, \quad (71)$$

where R is the radius of the Earth and D is a disk of radius R . The above coordinates are not orthogonal and the full equation at the end of Sec. III A must be implemented to obtain the FTLE. However, instead of computing the 2×2 matrix representation of $D\beta^{-1}$, we will exploit the embedding. The sphere \mathcal{M} is embedded in \mathbb{R}^3 and the chart derivative $D\beta^{-1}: T_{\mathbf{p}}\mathcal{M} \rightarrow T_{\mathbf{p}}\mathbb{R}^3$ can also be seen as a function $D\beta^{-1}: T_{\mathbf{p}}\mathcal{M} \rightarrow \mathbb{R}^3$. From this point of view, the matrix representation of the embedded chart is of size 3×2 and its QR decomposition is of the form

$$D\beta^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ \frac{x}{\sqrt{R^2-x^2-y^2}} & \frac{y}{\sqrt{R^2-x^2-y^2}} \end{pmatrix} = QR, \quad (72)$$

where Q is a matrix of size 3×2 satisfying $Q^T Q = I$ and R is an upper triangular matrix of size 2×2 . Note that the “reduced QR decomposition” above is inspired by the Gram–Schmidt process and does not match the usual QR decomposition, where Q would be of size 3×3 and R would be of size 3×2 . The desired matrices above can, however, easily be derived either from the Gram–Schmidt process, the modified Gram–Schmidt process or from the standard (full) QR factorization.¹⁶

The upper triangular matrix R can be computed analytically and depends on the coordinates x and y . Its inverse can also be computed analytically and, according to Sec. III A, the FTLE is given by the largest singular value of the matrix

$$R(x,y)MR^{-1}(x_0,y_0), \quad (73)$$

where (x_0,y_0) are the initial coordinates, (x,y) are the final coordinates, and M is the matrix containing the finite differences of manifold coordinates. In other words, we keep track of the manifold coordinates (x,y) and we compute finite differences as if the space was a plane. The only modification due to the manifold chart is the multiplication by $R(x,y)$ (on the left) and by $R^{-1}(x_0,y_0)$ (on the right) prior to computing the singular value.

Figure 9 reveals that the computed LCS match the material border of the ozone hole. It has been recognized that transport in the stratosphere is dominated by advection from large-scale structures and that on a time scale of days to weeks, the transport is quasi-horizontal, along isentropic surfaces,²² i.e., on two-dimensional spherical layers like the one shown in Fig. 9.

The bifurcating LCS structure is noteworthy in several ways. The ridges show details of the lobe dynamics associated with the structures, something that would be more difficult to obtain from an analysis of potential vorticity fields because vortex boundaries defined, for example, as potential vorticity contours or regions of enhanced potential vorticity gradients are less accurate. Although qualitatively the picture of the vortex breakup and of the large-scale filamentation in the LCS figure is similar to what can be seen in the potential vorticity field, the LCS are more suitable for calculations of transport (e.g., by lobe dynamics), for which accurate boundaries and small-scale details of the time-dependent vortex boundaries are needed. As showed by Green *et al.*¹⁷ and d’Ovidio *et al.*,⁹ defining LCS based on Lagrangian criteria instead of potential vorticity or other Eulerian quantities becomes essential for fluids with strong time dependence or in turbulent regions.

Figure 10 shows how a new LCS appears in the center of the polar vortex prior to the splitting and separates particles based on which daughter vortex they will end up in.

We finish this study by discussing the effect of the matrix R in the computation of the FTLE. For this example, ignoring the matrices $R(x,y)$ and $R(x_0,y_0)$ does not make any visible difference in the computed LCS. The value of the




FIG. 9. (Color) The left column shows the superposition of the attractive and repulsive LCS on the 650 K isentrope on the days surrounding the Antarctic polar vortex splitting event of September 2002 (based on NCEP/NCAR reanalysis data). The attracting (repelling) curves, analogous to unstable (stable) manifolds, are shown in blue (red). Before and after the splitting event in late September, we see an isolated blob of air, bounded by LCS curves, slowly rotating over Antarctica. In the days leading up to the splitting, LCS curves form inside the vortex. The vortex pinches off, sending the northwestern part of the ozone hole off into the midlatitudes while the southwestern portion goes back to its regular position over Antarctica. Note the formation of lobe at the edges where chaotic stirring occurs across the LCS. The right column shows the corresponding daily ozone concentration (based on NASA TOMS satellite data) (enhanced online). [URL: <http://dx.doi.org/10.1063/1.3278516.4>]

FTLE does change, but the variations are small for points where the FTLE is high; the resulting ridges are almost identical. The differences get much higher when the chart derivative has higher variations, for example, considering the flow near surface topography or along subsurfaces (e.g., invariant manifolds) of high curvature. However as our calculations were limited to one chart, specifically the southern hemi-

isphere below the 10S latitude, we did not see large variations (but note that the FTLE in the chosen chart is clearly incorrect as we approach the equator).

This effect can be understood by analyzing the structure of the matrices R and M . Near a LCS, the finite-time coordinate deformation tensor M has an exponentially growing eigenvalue and an exponentially decaying eigenvalue. The separation between the eigenvalues grows exponentially in time and there are several orders of magnitude between the two when we integrate particles for 8 days. In contrast, the matrices R are bounded (regardless of the integration time). Consequently, a large eigenvalue in M induces a large eigenvalue in $R(x,y)MR(x_0,y_0)$. The actual value of the FTLE changes if we ignore R , but, provided that the integration time is sufficiently large, the ridges will remain the same. An important corollary is that the LCS computation is robust to error in manifold modelization.

D. Example 5: A multivortex flow on a sphere

In Sec. IV C, we studied the 2002 splitting of the Antarctic polar vortex using an explicit atlas of the spherical Earth. An alternate method for computing FTLE in this system is to cover the Earth with an unstructured mesh that we consider embedded in \mathbb{R}^3 (see, e.g., left panel of Fig. 11). As we follow the mesh nodes in their motion, the algorithm in Sec. IV B provides an estimate of the FTLE for each of the nodes. In this setting, the finite differences are taken with respect to the edges of the two-dimensional meshes but the differentiated functions are unprojected coordinates in \mathbb{R}^3 . In other words, the matrix elements of $X (2 \times N)$ are differences between manifold coordinates (projection on the reconstructed local basis) but the matrix elements of $Y (3 \times N)$ are differences in Cartesian $x, y,$ and z coordinates.

By and large, this technique is more convenient and efficient. It does not require knowledge of analytical equations of the charts nor the analytical computation of the tangent bundle. Furthermore, it is simpler to implement and optimize. When applied to the NCEP/NCAR reanalysis data in Sec. IV C, we obtain the same FTLE field in less computation time.

To illustrate the method on a different, but related system, we consider chaotic advection in a multivortex flow on a sphere. This system is a semi-analytical model that offers a useful paradigm for understanding how complicated spatial structures can arise and evolve, such as the polar vortex split. As a first attempt to model this event, we consider the motion of a passive tracer due to velocity field produced by four self-advecting point vortices, all constrained to the surface of the sphere (for additional details, see, e.g., Ref. 42). We consider this the simplest model for capturing some of the dynamical features of the event.

Additional important geophysical effects such as rotation^{26,43} or vertical density stratification further complicate these dynamical processes but are not considered here.

The dynamics of point vortices moving on the surface of a sphere is not as well understood as the corresponding planar problem.³ However the model is very relevant both in




FIG. 10. (Color) Prior to the splitting event, repulsive LCS appear in the core of the polar vortex indicating the onset of a bifurcation and a separation line between particles that will end up in different vortices. To verify the dynamics we initiated two parcels of particles on 20 September, one on each side of the nascent LCS. The simulation shows that the green parcel (northwest of the LCS) will remain in the core vortex while the purple parcel (southeast of the LCS) is dragged into the secondary vortex and disintegrate at higher latitude (enhanced online). [URL: <http://dx.doi.org/10.1063/1.3278516.5>]

atmospheric and oceanographic settings when one considers large-scale phenomena where the spherical geometry of Earth’s surface becomes important. The full spherical geometry, as opposed to tangent plane approximations, is particularly important when considering global streamline patterns generated by a given vorticity distribution since the Poincaré index theorem provides an important constraint on allowable patterns.³⁰ These patterns, in turn, provide the dynamical templates by which one can begin to understand the chaotic advection of particles in a vortex-dominated flow, a topic closely related to the dynamics of the point vortices.

Our interest is the simple case of four identical vortices, initially on a constant latitudinal ring of colatitude θ_0 with respect to the pole. When the vortices are evenly spaced along the ring, the configuration is known to be a relative equilibrium configuration.³⁸

We perturb the configuration symmetrically. We define the perturbation parameter ε , which takes the configuration from its initial square shape ($\varepsilon=0$), to a rectangle ($\varepsilon>0$), ultimately to the singular limit of a two-vortex configuration ($\varepsilon=1$). The vortex motion consists of two frequencies which, in general, are incommensurate.




FIG. 11. (Color) Left: Unstructured mesh on the Earth. Center: FTLE for $T=10T_{rel}$ (forward time). Right: FTLE for $T=-10T_{rel}$ (backward time).




FIG. 12. (Color) Superposition of the attractive and repulsive LCS for a perturbed four vortex ring near the South Pole. As in the LCS shown in Fig. 9 from atmospheric reanalysis data, the model shows a hyperbolic region and the formation of filamentary structures that extend over large portions of the sphere (enhanced online). [URL: <http://dx.doi.org/10.1063/1.3278516.6>]

For a given θ_0 , the parameter ε naturally divides the phase space into two regimes in which one would expect the advected particle motion to share certain dynamical characteristics.

For an unperturbed ring of N vortices on the sphere, for $4 \leq N \leq 6$, there is a colatitudinal region of instability which opens up in the equatorial zone, whereas for $N \geq 7$ the ring is unstable at all latitudes.³⁸ In particular, for $N=4$, the ring is stable if $\theta_0 < \cos^{-1}(1/\sqrt{3})$, i.e., $\theta_0 < 54.7^\circ$.

The dynamics of the four identical vortices can be decomposed into a superposition of relative vortex motion along with a global rotation of the whole system around the center of vorticity. The relative motion is periodic with period T_{rel} and frequency $\omega_{\text{rel}} = 2\pi/T_{\text{rel}}$. The global rotation can be characterized by measuring the angular displacement $\Delta\phi_0$ about the center of vorticity between two configurations separated by time T_{rel} . The frequency of the global rotation can be defined as $\omega_{\text{glob}} = \Delta\phi_0/T_{\text{rel}}$. The pair of frequencies $(\omega_{\text{rel}}, \omega_{\text{glob}})$ depend on the parameters (θ_0, ε) and are generally incommensurate, implying that the overall vortex motion is quasiperiodic.

Without loss of generality, we align the center of vorticity with the south polar axis. The global rotation of the vortices is then around the z -axis with frequency ω_{glob} . In the rotating frame, the particle motion corresponds to a periodically forced Hamiltonian dynamical system that depends, in general, on the pair of parameters (θ_0, ε) .⁴²

For illustration, we consider a ring of vortices initially at latitude $\theta_0 = 30^\circ$ and with $\varepsilon = 0.54$, which roughly approximates the extent of the Antarctic polar vortex. In Fig. 11, we

show the FTLE field for the resulting particle velocity field and the superimposed attraction and repulsion LCS in Fig. 12. One notices the prominent rotating saddle point at the pole with undulating stable and unstable manifolds, which is also apparent in the real data (Fig. 9).

Advection in a multivortex flow on a sphere provides an important link between simple dynamical systems models and much more complicated models of particle advection in global geophysical flows, such as the polar vortex.^{23,54} Taking the point of view of building dynamically consistent simple models, we can add additional vortices of various strengths, along with realistic rotation models, all of which avoids the trouble of interpreting results belonging only to a kinematic model.²² Techniques developed and tested in these simpler models for quantifying—and perhaps even influencing—transport can be useful for more realistic models of global-scale geophysical phenomena, such as pollution dispersion in the atmosphere and ocean,^{35,2} polar vortex break-up,^{61,23} and large-scale transport of aerobiota, including airborne plant pathogens and other invasive species.^{1,25,51}

V. CONCLUSIONS

This paper introduces numerical techniques for resolving and representing evolving, and likely convoluted, $(n-1)$ -dimensional LCS in nonautonomous vector fields on n -dimensional Riemannian manifolds, where the vector field is known either analytically or from measured data.

We have shown that LCS exist for systems on any Riemannian differentiable manifold, whether orientable or not. While our first algorithm for computing FTLE requires knowledge of charts and chart derivatives and is more theoretical in nature, we have developed a general algorithm for computing FTLE on Riemannian manifolds covered by meshes of polyhedra. Only the initial positions, the final positions, and a link list of node neighbors are required to use the algorithm.

In addition to providing FTLE fields on Riemannian manifolds, we also showed that the unstructured algorithm can be efficiently used in Euclidean spaces where it gives the opportunity to refine the mesh near the LCS while not increasing the number of elements in other regions. Adaptive meshing is an essential step in achieving FTLE computation in high-dimensional spaces at a sufficiently high resolution. High resolution FTLE fields are important for accurate automated ridge extraction algorithms.

In this paper, we have applied the developed methods to compute FTLE in two-dimensional Riemannian manifolds. However, the methods also apply to the computation of FTLE for three-dimensional flows. In this case, the LCS are two-dimensional surfaces. More generally, computation of FTLE for a dynamical system on a n -dimensional manifold leads to the identification of LCS of dimension $n-1$.³⁷ A library encapsulating all the methods described in this paper is available at <http://www.lekien.com/~francois/software/libunsdle>. The library can compute FTLE on unstructured meshes in spaces of any dimension (using an embedding or not).

We also note that the present paper focuses on computing the FTLE but the definitions and methods extend easily to FSLE.^{27,31,9} In the finite time approach, we evaluate the stretching of the elements for a prescribed finite horizon in time. In the finite size scheme, one marks the time it takes to reach a prescribed stretching ratio. The techniques developed in this paper can be applied equally to FTLE or FSLE.

Further developments will include application to higher dimensional manifolds relevant in mechanics such as products of the Lie groups $SO(3)$, $SE(3)$, and their tangent bundles. With the appropriate visualization tools, computations in these phase spaces can aid the search for dynamical structure in mechanical and biomechanical data from experiments or computational models, such as separatrices between stable and unstable motions.^{44,15,49,58,59,50} Additionally, it may be of interest to detect LCS structure on high-dimensional invariant manifolds of possibly high curvature, which are commonly found in phase space, such as energy manifolds, center manifolds, and stable and unstable manifolds.

- ¹Aylor, D. E., "Spread of plant disease on a continental scale: Role of aerial dispersal of pathogens," *Ecology* **84**, 1989–1997 (2003).
- ²Beron-Vera, F. J., Olascoaga, M. J., and Goni, G. J., "Oceanic mesoscale eddies as revealed by Lagrangian coherent structures," *Geophys. Res. Lett.* **35**, L12603, doi:10.1029/2008GL033957 (2008).
- ³Boatto, S. and Pierrehumbert, R., "Dynamics of a passive tracer in the velocity field of four identical point vortices," *J. Fluid Mech.* **394**, 137–174 (1999).
- ⁴Bowman, K. P., "Manifold geometry and mixing in observed atmospheric flows," preprint, 1999.
- ⁵Camassa, R. and Wiggins, S., "Chaotic advection in a Rayleigh–Bénard flow," *Phys. Rev. A* **43**, 774–797 (1991).
- ⁶Cardwell, B. M. and Mohseni, K., "Vortex shedding over a two-dimensional airfoil: Where the particles come from," *AIAA J.* **46**, 545–547 (2008).
- ⁷Coulliette, C., Lekien, F., Paduan, J. D., Haller, G., and Marsden, J. E., "Optimal pollution mitigation in Monterey Bay based on coastal radar data and nonlinear dynamics," *Environ. Sci. Technol.* **41**, 6562–6572 (2007).
- ⁸Coulliette, C. and Wiggins, S., "Intergyre transport in a wind-driven, quasigeostrophic double gyre: An application of lobe dynamics," *Nonlinear Processes Geophys.* **8**, 69–94 (2001).
- ⁹d'Ovidio, F., Isern-Fontanet, J., López, C., Hernández-García, E., and García-Ladona, E., "Comparison between Eulerian diagnostics and finite-size Lyapunov exponents computed from altimetry in the Algerian basin," *Deep-Sea Res., Part I* **56**, 15–31 (2009).
- ¹⁰Dieci, L. and Eirola, T., "On smooth decompositions of matrices," *SIAM J. Matrix Anal. Appl.* **20**, 800–819 (1999).
- ¹¹Fenichel, N., "Persistence and smoothness of invariant manifolds for flows," *Indiana Univ. Math. J.* **21**, 193 (1971).
- ¹²Froyland, G. and Padberg, K., "Almost-invariant sets and invariant manifolds—Connecting probabilistic and geometric descriptions of coherent structures in flows," *Physica D* **238**, 1507–1523 (2009).
- ¹³Garth, C., Gerhardt, F., Tricoche, X., and Hagen, H., "Efficient computation and visualization of coherent structures in fluid flow applications," *IEEE Trans. Vis. Comput. Graph.* **13**, 1464–1471 (2007).
- ¹⁴Garth, C., Wiebel, A., Tricoche, X., Joy, K., and Scheuermann, G., "Lagrangian visualization of flow-embedded surface structures," *Comput. Graph. Forum* **27**, 1007–1014 (2008).
- ¹⁵Gawlik, E. S., Marsden, J. E., Du Toit, P. C., and Campagnola, S., "Lagrangian coherent structures in the planar elliptic restricted three-body problem," *Celest. Mech. Dyn. Astron.* **103**, 227–249 (2009).
- ¹⁶Golub, G. H. and Van Loan, C. F., *Matrix Computations*, 3rd ed. (Johns Hopkins University Press, Baltimore, MD, 1996).
- ¹⁷Green, M. A., Rowley, C. W., and Haller, G., "Detection of Lagrangian coherent structures in three-dimensional turbulence," *J. Fluid Mech.* **572**, 111–120 (2007).
- ¹⁸Guckenheimer, J. and Holmes, P. J., *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields* (Springer-Verlag, New York, 1983).
- ¹⁹Haller, G., "Finding finite-time invariant manifolds in two-dimensional velocity fields," *Chaos* **10**, 99–108 (2000).
- ²⁰Haller, G., "Distinguished material surfaces and coherent structures in 3D fluid flows," *Physica D* **149**, 248–277 (2001).
- ²¹Haller, G. and Poje, A. C., "Finite-time transport in aperiodic flows," *Physica D* **119**, 352–380 (1998).
- ²²Haynes, P., "Transport, stirring and mixing in the atmosphere," in *Proceedings of the NATO Advanced Study Institute on Mixing: Chaos and Turbulence*, Cargese, Corse, France, 7–20 July 1996, edited by Chaté, H. and Villermaux, E. (Kluwer, Dordrecht, 1999), pp. 229–272.
- ²³Haynes, P., "Stratospheric dynamics," *Annu. Rev. Fluid Mech.* **37**, 263–293 (2005).
- ²⁴Hirsch, M. W., *Differential Topology* (Springer, New York, 1976), p. 33.
- ²⁵Isard, S. A., Gage, S. H., Comtois, P., and Russo, J. H., "Principles of the atmospheric pathway for invasive species applied to soybean rust," *BioScience* **55**, 851–861 (2005).
- ²⁶Jamaloodeen, M. I. and Newton, P. K., "The N -vortex problem on a rotating sphere. II. Heterogeneous Platonic solid equilibria," *Proc. R. Soc. London, Ser. A* **462**, 3277–3299 (2006).
- ²⁷Joseph, B. and Legras, B., "Relation between kinematic boundaries, stirring and barriers for the Antarctic polar vortex," *J. Atmos. Sci.* **59**, 1198–1212 (2002).
- ²⁸Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., Mo, K. C., Ropelewski, C., Wang, J., Leetmaa, A., Reynolds, R., Jenne, R., and Joseph, D., "The NCEP/NCAR 40-year reanalysis project," *Bull. Am. Meteorol. Soc.* **77**, 437–471 (1996).
- ²⁹Kaplan, D. M. and Lekien, F., "Spatial interpolation and filtering of surface current data based on open-boundary modal analysis," *J. Geophys. Res., [Oceans]* **112**, C12007, doi:10.1029/2006JC003984 (2007).
- ³⁰Kidambi, R. and Newton, P. K., "Streamline topologies for integrable vortex motion on a sphere," *Physica D* **140**, 95–125 (2000).
- ³¹Koh, T. and Legras, B., "Hyperbolic lines and the stratospheric polar vortex," *Chaos* **12**, 382–394 (2002).
- ³²Kuiper, N. H., "On C^1 -isometric embeddings," *Proc. K. Ned. Akad. Wet., Ser. A: Math. Sci.* **58**, 545–556 (1955).
- ³³Lekien, F. and Coulliette, C., "Chaotic stirring in quasi-turbulent flows," *Philos. Trans. R. Soc. London, Ser. A* **365**, 3061–3084 (2007).
- ³⁴Lekien, F., Coulliette, C., Bank, R., and Marsden, J. E., "Open-boundary modal analysis: Interpolation, extrapolation, and filtering," *J. Geophys. Res., [Oceans]* **109**, C12004, doi:10.1029/2004JC002323 (2004).
- ³⁵Lekien, F., Coulliette, C., Mariano, A. J., Ryan, E. H., Shay, L. K., Haller, G., and Marsden, J. E., "Pollution release tied to invariant manifolds: A case study for the coast of Florida," *Physica D* **210**, 1–20 (2005); see www.lekien.com/~francois/papers/rsmas.
- ³⁶Lekien, F. and Haller, G., "Unsteady flow separation on slip boundaries," *Phys. Fluids* **20**, 097101 (2008).
- ³⁷Lekien, F., Shadden, S. C., and Marsden, J. E., "Lagrangian coherent structures in n -dimensional systems," *J. Math. Phys.* **48**, 065404 (2007); see www.lekien.com/~francois/papers/lcs3d.
- ³⁸Lim, C. C., Montaldi, J., and Roberts, M., "Relative equilibria of point vortices on the sphere," *Physica D* **148**, 97–135 (2001).
- ³⁹Lipinski, D., Cardwell, B., and Mohseni, K., "A Lagrangian analysis of a two-dimensional airfoil with vortex shedding," *J. Phys. A: Math. Theor.* **41**, 344011 (2008).
- ⁴⁰Mañé, R., "Persistent manifolds are normally hyperbolic," *Trans. Am. Math. Soc.* **246**, 261–283 (1978).
- ⁴¹Nash, J., " C^1 -isometric imbeddings," *Ann. Math.* **60**, 383–396 (1954).
- ⁴²Newton, P. K. and Ross, S. D., "Chaotic advection in the restricted four-vortex problem on a sphere," *Physica D* **223**, 36–53 (2006).
- ⁴³Newton, P. K. and Shokraneh, H., "The N -vortex problem on a rotating sphere. I. Multi-frequency configurations," *Proc. R. Soc. London, Ser. A* **462**, 149–169 (2006).
- ⁴⁴Norris, J. A., Marsh, A. P., Granata, K. P., and Ross, S. D., "Revisiting the stability of 2D passive biped walking: Local behavior," *Physica D* **237**, 3038–3045 (2008).
- ⁴⁵Olascoaga, M. J., Rypina, I. I., Brown, M. G., Beron-Vera, F. J., Kocak, H., Brand, L. E., Halliwell, G. R., and Shay, L. K., "Persistent transport barrier on the West Florida Shelf," *Geophys. Res. Lett.* **33**, L22603, doi:10.1029/2006GL027800 (2006).
- ⁴⁶Peng, J. F., Dabiri, J. O., Madden, P. G., and Lauder, G. V., "Non-invasive

- measurement of instantaneous forces during aquatic locomotion: A case study of the bluegill sunfish pectoral fin," *J. Exp. Biol.* **210**, 685–698 (2007).
- ⁴⁷Rom-Kedar, V. and Wiggins, S., "Transport in two-dimensional maps: Concepts, examples, and a comparison of the theory of Rom-Kedar and Wiggins with the Markov model of MacKay, Meiss, Ott, and Percival," *Physica D* **51**, 248–266 (1991).
- ⁴⁸Rom-Kedar, V. and Wiggins, S., "Transport in 2-dimensional maps," *Arch. Ration. Mech. Anal.* **109**, 239–298 (1990).
- ⁴⁹Ross, S. D., "The interplanetary transport network," *Am. Sci.* **94**, 230–237 (2006).
- ⁵⁰Ross, S. D., Tanaka, M. L., and Senatore, C., "Detecting dynamical boundaries from kinematic data in biomechanics," *Chaos* **20**, 017507 (2010).
- ⁵¹Schmale, D. G., Shields, E. J., and Bergstrom, G. C., "Night-time spore deposition of the *Fusarium* head blight pathogen, *Gibberella zeae*," *Can. J. Plant Pathol.* **28**, 100–108 (2006).
- ⁵²Shadden, S. C., Lekien, F., and Marsden, J. E., "Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows," *Physica D* **212**, 271–304 (2005).
- ⁵³Shadden, S. C. and Taylor, C. A., "Characterization of coherent structures in the cardiovascular system," *Ann. Biomed. Eng.* **36**, 1152–1162 (2008).
- ⁵⁴Simmons, A., Hortal, M., Kelly, G., McNally, A., Untch, A., and Uppala, S., "ECMWF analyses and forecasts of stratospheric winter polar vortex breakup: September 2002 in the Southern Hemisphere and related events," *J. Atmos. Sci.* **62**, 668–689 (2005).
- ⁵⁵Solomon, T. H. and Gollub, J. P., "Chaotic particle-transport in time-dependent Rayleigh–Bénard convection," *Phys. Rev. A* **38**, 6280–6286 (1988).
- ⁵⁶Solomon, T. H. and Gollub, J. P., "Passive transport in steady Rayleigh–Bénard convection," *Phys. Fluids* **31**, 1372–1379 (1988).
- ⁵⁷Tallapragada, P. and Ross, S. D., "Particle segregation by Stokes number for small neutrally buoyant spheres in a fluid," *Phys. Rev. E* **78**, 036308 (2008).
- ⁵⁸Tanaka, M. L. and Ross, S. D., "Separatrices and basins of stability from time series data: An application to biodynamics," *Nonlinear Dyn.* **58**, 1–21 (2009).
- ⁵⁹Tanaka, M. L., Nussbaum, M. A., and Ross, S. D., "Evaluation of the threshold of stability for the human spine," *J. Biomech.* **42**, 1017–1022 (2009).
- ⁶⁰Tew Kai, E., Rossi, V., Sudre, J., Weimerskirch, H., Lopez, C., Hernandez-Garcia, E., Marsac, F. and Garçon, V., "Top marine predators track Lagrangian coherent structures," *Proc. Natl. Acad. Sci. U.S.A.* **106**, 8245–8250 (2009).
- ⁶¹*The Antarctic Vortex Splitting Event*, edited by Shepherd, T. and Plumb, A., special issue of *J. Atmos. Sci.* **62** (3) (2005).
- ⁶²Whitney, H., "The self-intersections of a smooth n -manifold in $2n$ -space," *Ann. Math.* **45**, 220–293 (1944).