

Accepted to the *Journal of Computational Dynamics* on 14 February 2020.

UNCERTAINTY IN FINITE-TIME LYAPUNOV EXPONENT COMPUTATIONS

SANJEEVA BALASURIYA*

School of Mathematical Sciences
University of Adelaide
Adelaide SA 5005, Australia

—and—

Department of Mechanical Engineering
Technische Universiteit Eindhoven
5612 AZ Eindhoven, The Netherlands

(Communicated by the associate editor name)

ABSTRACT. The Finite-Time Lyapunov Exponent (FTLE) is a well-established numerical tool for assessing stretching rates of initial parcels of fluid, which are advected according to a given time-varying velocity field (which is often available only as data). When viewed as a field over initial conditions, the FTLE's spatial structure is often used to infer the nonhomogeneous transport. Given the measurement and resolution errors inevitably present in the unsteady velocity data, the computed FTLE field should in reality be treated only as an approximation. A method which, for the first time, is able for attribute spatially-varying errors to the FTLE field is developed. The formulation is, however, confined to two-dimensional flows. Knowledge of the errors prevent reaching erroneous conclusions based only on the FTLE field. Moreover, it is established that increasing the spatial resolution does not improve the accuracy of the FTLE field in the presence of velocity uncertainties, and indeed has the opposite effect. Stochastic simulations are used to validate and exemplify these results, and demonstrate the computability of the error field.

1. Introduction. In many situations (computational fluid dynamics simulations, satellite observations of the ocean, particle image velocimetry of industrial flows) fluid velocities are available as data in time and space. Understanding the spatiotemporal transport resulting from this data is important in determining, for example, how a pollutant spreads, how energy is transported, and which regions remain unmixed. There are a multitude of approaches which have been suggested for this, reviewed for example in [?, ?, ?, ?, ?, ?]. A crucial point is that Eulerian (viewed in time-slices) information is inadequate when the velocity is time-dependent; Lagrangian evolution of trajectories is necessary to reach any conclusions regarding transport.

2010 *Mathematics Subject Classification.* Primary: 37B5, 76M25; Secondary: 60H10.

Key words and phrases. Finite-Time Lyapunov Exponents, Lagrangian coherent structures,

Partial support from the Australian Research Council (grants FT130100484 and DP200101764) and the Australian Department of Education and Training (Endeavour Research Leadership Award) are gratefully acknowledged.

One well-established tool in this endeavor is the Finite-Time Lyapunov Exponent (FTLE), which is computed as a field on the space of initial conditions [?, ?]. It measures the elongation of a tiny fluid parcel placed at each initial location, due to the flow over the given time-duration. In computing the FTLE field, the gradient of the flow map is usually approximated; this provides information on infinitesimal changes in the initial location. There continue to be methods put forward to improve the efficiency of FTLE computations [?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, e.g.] and discuss their legitimacy and relevance [?, ?, ?, ?, ?]. A related quantity is the Finite-Size Lyapunov Exponent (FSLE), which does the computation differently: each initial condition is evolved until nearby trajectories get further away than a specified threshold, at which point the relative elongation is computed [?, ?, ?, ?, ?, ?, ?]. Once a field has been obtained, an important characteristic sought is whether there are codimension-1 ridges present in the field (for example for two-dimensional flows, whether there are curves in the field at which the value is sharply larger than just off the curve). The precise identification of a ‘ridge’ is not always easy, and has resulted in various computational techniques being advanced [?, ?, ?, ?]. The expectation—though this has been shown to be not always true [?, ?, ?]—is that these ridges are analogous to stable/unstable manifolds in smooth infinite-time flows (which are difficult to compute for unsteady flows [?, ?]), and therefore form flow separators. Given the unsteadiness of realistic velocity fields, these ridges move as time evolves, and thereby the hope is to capture the moving flow barriers (which demarcate moving coherent regions of fluids). While persuasive indications of this abound in the literature, slight problematic issues cannot be ignored: ridges are not guaranteed to evolve with the flow [?], can have a nonzero flux across them [?], and in any case the ambiguity of stable/unstable manifolds in finite-time situations [?, ?] renders the comparison questionable. Indeed, an exact understanding of what a ‘flow barrier’ means in this situation continues to be debated. Irrespective of this connection to flow barriers, though, the FTLE field indubitably quantifies the rates of elongation of fluid elements, and therefore offers valuable information on the flow evolution.

While there are many articles in oceanography [?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?, e.g.], atmospheric science [?, ?, ?, ?, ?, ?, e.g.], industrial and experimental flows [?, ?, ?, ?, ?, ?, ?, e.g.] and biology [?, ?, e.g.] which use FTLE fields to reach conclusions regarding transport, the robustness of the procedure is not always clear. Indeed, several studies have pointed to the fact that the FTLE field can be uncertain [?, ?, ?, ?, ?, ?, ?], often demonstrated by methods such as artificially downsampling extant data and comparing with the ‘true’ FTLE field for the highly-sampled data. This highly cautionary evidence is difficult to put into practical usage in quantifying an error when using the best available data and consequently, in most situations, computed FTLE fields at the best resolution are treated as the ground truth in reaching conclusions from the field. However, in all realistic situations uncertainties in the (best available) Eulerian velocity data are inevitable due to *many* reasons, which include observational uncertainty, resolution limitation, modelling uncertainty in transforming proxy data into velocities, etc. How do such uncertainties affect the FTLE field, and are conclusions reached legitimate? In particular, is it possible to *quantify* an error in the computed FTLE field? A recent article by Guo et al [?] takes a step in answering this by computing a range of FTLE fields arising from introducing stochasticity. Each field is computed via many stochastic simulations, and statistically analyzed. This process is computationally

expensive and does not give uncertainty thresholds on the base FTLE field. The current article instead seeks *theoretical error estimates* for the FTLE field, based on parameters such as the spatial resolution which are known to the practitioner, and which has been identified as a definite cause for error in many cases [?, ?, ?, ?, ?, ?, e.g.].

The basis for quantifying the FTLE error is a recent article which enables the characterization of an uncertainty in Lagrangian particle positions due to the presence of uncertainties in Eulerian velocity fields [?]. Given an initial condition (‘a particle’), the stochastic differential equations approach in [?] assigns an error, projected in each general direction, for the particle’s eventual location after advecting over a finite time by a two-dimensional velocity field. In this paper, the theory from [?] is leveraged to compute an uncertainty in the FTLE field. In addition to being able to estimate errors for the particular instance of FTLEs, this constitutes an in-principle approach of adapting *location* uncertainties (as computed by [?]) to uncertainties in *spatial fields* generated from the corresponding flows.

This paper is organized as follows. In Section 2, the aforementioned theory [?] is adapted to compute a spatially-varying uncertainty range for the FTLE field in unsteady flows. It is only possible to assess the FTLE uncertainty for *two-dimensional* flows, since the particle uncertainty theory [?] has this restriction. In formulating the theory, the Eulerian velocities are assumed to possess uncertainties which are Gaussian. By quantifying the probabilistic impact of small stochastic variation in the velocity data, ‘error-bars’ are obtained. These errors form a *nonhomogeneous* field, highlighting that the FTLE computations have spatially-varying errors. Since *variations* in the computed FTLE field are often used as diagnostics for transport, these errors imply that any conclusions need to be tempered appropriately. The inaccuracies in identifying coherent structures by simply computing FTLEs (by effectively treating the velocity field as highly-resolved and deterministic) are exemplified via several numerical examples in Section 3. Moreover, the performance of the theoretical error is validated by comparing with error measures generated from stochastic simulations. Some concluding remarks are given in Section 4.

2. The FTLE and its error. Consider a flow in two dimensions from time t_0 to $t_0 + T$, where $T > 0$ is the time-of-flow. Assuming that Eulerian velocity data $\mathbf{u}(\mathbf{x}, t)$ is available on a spatiotemporal grid, particle trajectories are obtained by solving

$$\frac{d}{dt}\mathbf{x} = \mathbf{u}(\mathbf{x}, t) \quad ; \quad \mathbf{x}(t_0) = \mathbf{x}_0 \tag{1}$$

for $\mathbf{x}_0 \in \Omega_0$, where Ω_0 is the space of initial conditions, and the time $t \in (t_0, t_0 + T)$. The ordinary differential equation (1) is numerically solved till time $t_0 + T$ by interpolating \mathbf{u} as necessary.

Let \mathbf{F} be the flow map from t_0 to $t_0 + T$, i.e.,

$$\mathbf{F}(\mathbf{x}_0) = \mathbf{x}(t_0 + T) \tag{2}$$

where $\mathbf{x}(t)$ is the solution to (1). The Finite-Time Lyapunov Exponent (FTLE) associated with the flow from time t_0 to $t_0 + T$ is a scaled field on Ω_0 which captures the stretching σ of fluid blobs due to the flow. If \mathbf{y} is a small quantity, such stretching is given by

$$\sigma(\mathbf{x}_0) = \frac{|\mathbf{F}(\mathbf{x}_0 + \mathbf{y}) - \mathbf{F}(\mathbf{x}_0)|}{|\mathbf{y}|} \approx \frac{|\nabla \mathbf{F}(\mathbf{x}_0)\mathbf{y}|}{|\mathbf{y}|}, \tag{3}$$

where the stretching must be positive according to this definition (‘shrinking’ has a stretching value of less than unity). The FTLE is defined by

$$\Lambda(\mathbf{x}_0) := \frac{1}{|T|} \ln \left[\sup_{\mathbf{y} \neq \mathbf{0}} \frac{|\nabla \mathbf{F}(\mathbf{x}_0) \mathbf{y}|}{|\mathbf{y}|} \right], \quad (4)$$

where sup stands for the supremum taken over all directions associated with the small vector \mathbf{y} . The FTLE has dimensions of reciprocal time. Given the numerical resolution, computing the local stretching rate—the quantity in square brackets in (4)—is limited by an estimate such as (3). An alternative is that it is equivalent to the square-root of the largest eigenvalue of the Cauchy-Green tensor

$$C(\mathbf{x}_0) := [\nabla \mathbf{F}(\mathbf{x}_0)]^\top \nabla \mathbf{F}(\mathbf{x}_0). \quad (5)$$

It is clear that in computing the FTLE field Λ , the gradient of the flow map must be computed. Typically, one takes the initial conditions \mathbf{x}_0 on a spatial grid in Ω_0 , advects them forward to time $t_0 + T$ (using spatiotemporal interpolation as needed because \mathbf{u} is only available on spatial and temporal grids). The eventual location of the nearest neighbours of each point \mathbf{x}_0 enables an estimate of $\nabla \mathbf{F}(\mathbf{x}_0)$ —once again, essentially optimising the discretisation σ as given in (3) over all directions. Various methods in improving the efficiency of the general FTLE computation described here have been proposed [?, ?, ?, ?, ?, ?, ?, ?, ?, ?, ?], as have been methods [?, ?, ?, ?] for extracting ridges in the field (which, analogous to stable manifolds in infinite-time flows, strongly separate blobs of fluid perched on them in forward time, and consequently are barriers between coherent structures). Moreover, the implications of the FTLE field (and the Finite-Size Lyapunov Exponent, which is an alternative for computing the same quantity but leads to some differences) on fluid transport continues to be an active area of study [?, ?, ?].

The above is based on *deterministic* knowledge of the velocity field \mathbf{u} . However, in realistic situations the Eulerian velocity \mathbf{u} must have uncertainties, because

- There are *measurement* errors in \mathbf{u} , whether one is using a computational fluid dynamics method, or observations using instrumentation;
- If \mathbf{u} is computed from a proxy (e.g., oceanic sea-surface height from satellite measurements converted to a velocity using the geostrophic approximation, or estimated using data assimilation techniques), then there are *modelling* errors;
- Given that \mathbf{u} is on a spatiotemporal grid, since interpolation must be used to approximate \mathbf{F} , there are *resolution* errors.

What are the impacts of these uncertainties on the computed FTLE field?

Imagine two nearby initial conditions, evolved by the given velocity field as if it were deterministic (using smooth interpolation between gridpoints to estimate the velocity at intermediate times as needed). Their final locations will be close, and this means (see 3) that the gradient estimate will be small. However, imagine now that each of these trajectories is evolved with an ongoing uncertainty overlaid on the deterministic velocity field. In this case, the accumulated effect of the uncertainties can potentially make their eventual locations far apart, and thus *gradients can be large*. This means that estimates for the FTLE field may have very large uncertainties. The idea now is to *quantify* these uncertainties, to enable informed decisions to be taken from a computed FTLE field.

To proceed, some assumptions on the uncertainties in the Eulerian velocity field \mathbf{u} are needed. Gaussian uncertainties are an obvious first-case scenario. This assumption means that fluid trajectories will evolve according to the stochastic differential

equation

$$d\mathbf{x}_t = \mathbf{u}(\mathbf{x}_t, t)dt + \varepsilon d\mathbf{W}_t \quad ; \quad \mathbf{x}_{t_0} = \mathbf{x}_0 \quad (6)$$

instead of (1). The standard stochastic differential equations notation \mathbf{x}_t is used for flow trajectories rather than $\mathbf{x}(t)$. The Wiener process $d\mathbf{W}_t$ in two-dimensions—with independence in each of the two spatial dimensions— captures the Brownian motion in time, i.e., the Gaussian nature of the Eulerian velocity uncertainties. The quantity ε is small, indicating that there is some confidence in the velocity data.

Now, the eventual location \mathbf{x}_{t_0+T} obtained from (6) is *random*, and so it makes sense to examine means and variances in relation to many implementations of the Brownian motion. Recent work [?] (which is in a more general framework than (6)) quantifies this, in the limit of small ε :

- The leading-order—to $\mathcal{O}(\varepsilon)$ —mean of the location of \mathbf{x}_{t_0+T} is the deterministic location $\mathbf{x}(t_0 + T)$ obtained from (1);
- The leading-order—to $\mathcal{O}(\varepsilon)$ — variance $\mathbf{x}_{t_0+T} - \mathbf{x}(t_0 + T)$, optimized over all directions, is given by the *stochastic sensitivity* field $S^2(\mathbf{x}_0)$, whose numerical computation from the Eulerian velocity data is outlined in Appendix A. (The form of the distribution is not specified by the theory; only its variance in the sense described.)

It should be noted that the divergence of \mathbf{u} is permitted to be nonzero in this theory; there are many instances in which, even though the fluid is incompressible, the divergence of \mathbf{u} computed from the data may be nonzero. This is no surprise given the uncertainties mentioned previously, but also can occur because of other reasons (e.g., flow on ocean’s surface is not strictly two-dimensional, methodology used to obtain proxies for velocities, etc). How to numerically compute the stochastic sensitivity field S^2 is also discussed in Appendix A. This field will now be used in the quantification of error bounds for the FTLE field.

In the context of the current problem, it makes sense to have ε dimensional, and given by

$$\varepsilon = \sqrt{L_r v_r} \quad (7)$$

where L_r is the spatial resolution lengthscale, and v_r is the velocity-scale associated with the velocity error. This scaling arises from dimensional analysis, bearing in mind the standard convention [?, ?] that $dW_t \sim \sqrt{dt}$, as used in the derivations [?] of the S^2 field. Thus, two main uncertainties are captured within ε : the first, L_r , is known from the data, while an informed guess for the other, v_r , is assumed available (e.g., based on apparatus precision of velocity measurement, or an estimate for deviation from geostrophy for oceanic velocities obtained from satellite sea-surface height measurements).

The interpretation would be, then, that the data that is available results from *one* realization of the stochastic system (6). In the absence of any additional information on this realization, it makes sense to think of this as the ‘most likely’ realization in that Brownian motion has expectation zero over each of the time subintervals considered. Thus, the ‘most likely’ realization of (6) is (1), which is used to compute the FTLE field. However, there must be a ‘true’ flow map which is in general different from the flow map of (1). This uncertainty will lead to an error in the computed FTLE field. Unadorned variables will be used for the *computed* states from the data, and tilde variables for the *true* states (which are unknown):

$$\mathbf{F}(\mathbf{x}_0) = \mathbf{x}(t_0 + T) \quad \text{and} \quad \tilde{\mathbf{F}}(\mathbf{x}_0) = \mathbf{x}_{t_0+T}. \quad (8)$$

The *true* flow map $\tilde{\mathbf{F}}$ needs to be obtained from a particular (but unknown) realisation of (6). The stochastic sensitivity theory [?] described above therefore enables an estimate of the uncertainty in final location as

$$\left| \mathbf{F}(\mathbf{x}_0) - \tilde{\mathbf{F}}(\mathbf{x}_0) \right| \lesssim \varepsilon \sqrt{S^2(\mathbf{x}_0)}, \quad (9)$$

because S^2 is the variance. The vector \mathbf{y} in (4) can be taken to have a length L_r (smaller values are not measurable), and so by setting $\mathbf{y} = L_r \hat{\mathbf{n}}$ where $\hat{\mathbf{n}}$ is a unit vector,

$$\frac{\nabla \mathbf{F}(\mathbf{x}_0) \mathbf{y}}{|\mathbf{y}|} \sim \frac{\mathbf{F}(\mathbf{x}_0 + L_r \hat{\mathbf{n}}) - \mathbf{F}(\mathbf{x}_0)}{L_r}. \quad (10)$$

Even though L_r is assumed small, the above is $\mathcal{O}(1)$ in L_r because the numerator and denominator are both $\mathcal{O}(L_r)$. For the same \mathbf{y} , an estimate for the true gradient is therefore given by the above expression, with the \mathbf{F} replaced by $\tilde{\mathbf{F}}$. This means that

$$\begin{aligned} \left| \frac{\nabla \mathbf{F}(\mathbf{x}_0) \mathbf{y}}{|\mathbf{y}|} - \frac{\nabla \tilde{\mathbf{F}}(\mathbf{x}_0) \mathbf{y}}{|\mathbf{y}|} \right| &\lesssim \frac{1}{L_r} \left| (\mathbf{F}(\mathbf{x}_0 + L_r \hat{\mathbf{n}}) - \mathbf{F}(\mathbf{x}_0)) - (\tilde{\mathbf{F}}(\mathbf{x}_0 + L_r \hat{\mathbf{n}}) - \tilde{\mathbf{F}}(\mathbf{x}_0)) \right| \\ &\leq \frac{1}{L_r} \left| \mathbf{F}(\mathbf{x}_0 + L_r \hat{\mathbf{n}}) - \tilde{\mathbf{F}}(\mathbf{x}_0 + L_r \hat{\mathbf{n}}) \right| + \frac{1}{L_r} \left| \mathbf{F}(\mathbf{x}_0) - \tilde{\mathbf{F}}(\mathbf{x}_0) \right| \\ &\lesssim \frac{\varepsilon}{L_r} \sqrt{S^2(\mathbf{x}_0 + L_r \hat{\mathbf{n}})} + \frac{\varepsilon}{L_r} \sqrt{S^2(\mathbf{x}_0)} \\ &\lesssim 2 \sqrt{\frac{v_r S^2(\mathbf{x}_0)}{L_r}}, \end{aligned} \quad (11)$$

since $S^2(\mathbf{x}_0)$ is a good approximation to $S^2(\mathbf{x}_0 + L_r \hat{\mathbf{n}})$, and the small quantity ε in the stochastic differential equation formulation is given here by (7). This gives an estimate for the error within the square-brackets term in the FTLE expression (4). The actual $\tilde{\Lambda}$ FTLE value at \mathbf{x}_0 must include the above error term, and so

$$\begin{aligned} \tilde{\Lambda}(\mathbf{x}_0) &= \frac{1}{|T|} \ln \left[\sup_{|\mathbf{y}| \neq \mathbf{0}} \frac{|\nabla \mathbf{F}(\mathbf{x}_0) \mathbf{y}|}{|\mathbf{y}|} \pm 2 \sqrt{\frac{v_r S^2(\mathbf{x}_0)}{L_r}} \right] \\ &= \frac{1}{|T|} \ln \left[e^{\Lambda(\mathbf{x}_0)|T|} \pm 2 \sqrt{\frac{v_r S^2(\mathbf{x}_0)}{L_r}} \right] \\ &= \frac{1}{|T|} \ln \left[e^{\Lambda(\mathbf{x}_0)|T|} \left(1 \pm \frac{2 \sqrt{v_r S^2(\mathbf{x}_0)}}{e^{\Lambda(\mathbf{x}_0)|T|} \sqrt{L_r}} \right) \right] \\ &= \frac{1}{|T|} \left\{ \ln \left[e^{\Lambda(\mathbf{x}_0)|T|} \right] + \ln \left[\left(1 \pm \frac{2 \sqrt{v_r S^2(\mathbf{x}_0)}}{e^{\Lambda(\mathbf{x}_0)|T|} \sqrt{L_r}} \right) \right] \right\} \\ &\approx \Lambda(\mathbf{x}_0) \pm \frac{2 \sqrt{v_r S^2(\mathbf{x}_0)}}{e^{\Lambda(\mathbf{x}_0)|T|} \sqrt{L_r} |T|}, \end{aligned}$$

where the last step is based on the standard Taylor expansion $\ln(1+b) \approx b$ assuming that b is small. Thus, an *interval* for the FTLE is

$$\tilde{\Lambda}(\mathbf{x}_0) \in [\Lambda(\mathbf{x}_0) - \mathcal{E}(\mathbf{x}_0), \Lambda(\mathbf{x}_0) + \mathcal{E}(\mathbf{x}_0)], \quad (12)$$

in which the *FTLE error* is

$$\mathcal{E}(\mathbf{x}_0) = \frac{2\sqrt{v_r S^2(\mathbf{x}_0)}}{e^{\Lambda(\mathbf{x}_0)|T|}\sqrt{L_r}|T|}. \quad (13)$$

As required, the error \mathcal{E} has dimensions of inverse time because (see Appendix A) S^2 has dimensions of the square-root of time in the way that it is defined here. An important first step in calculating (13) is in fact the computation of $S^2(\mathbf{x}_0)$, which is detailed in Appendix A. This only requires knowledge of the same data that is used to compute the FTLE field, and is a slightly more elaborate computation (the flow map gradient needs to be computed at each time-step, and then an integration over time needs to be computed). Now, the error \mathcal{E} can be viewed as a field on the space of initial conditions \mathbf{x}_0 , allowing for the error, and an appropriate interval as in (12), to be computed as fields. As will be demonstrated, the *nonhomogeneous* nature of the error field allows for assigning different certainties to the FTLE values across the domain, thereby tempering any conclusions related to heterogeneous transport implications from the FTLE field.

An important observation is that one cannot make the FTLE error smaller by simply increasing the resolution, while keeping the velocity uncertainty fixed but nonzero. Having spatially higher-resolved data in which the uncertainty at the data points remains high does *not* improve conclusions; it actually decreases certainty of the calculated FTLE value. This is due to the necessity of computing the flow map gradient in which the location uncertainty $\sim \sqrt{L_r v_r}$ while the spatial scale which divides this goes as $\sim L_r$. It is the ratio v_r/L_r which must be made small, if possible. (If v_r were zero, then the analysis needs adjustment, because then the stochasticity is ignored and errors are simply caused by the $\mathcal{O}(L_r)$ interpolation error in determining $\nabla \mathbf{F}$. This will lead to the error decreasing with L_r . However, this is *not* the situation which is examined in this paper, which assumes nonzero velocity uncertainties.)

The fact that improving the spatial resolution does *not* increase the certainty in the FTLE field will be demonstrated in the next section, using stochastic simulations. Put another way, the theoretical error expression gives the unexpected implication that it is the smallness of the ratio v_r/L_r , rather than making v_r or L_r small by themselves, that leads to a diminished error. Since this ratio has dimensions of inverse time, it is suggested that the FTLE field has a higher degree of certainty if

$$\frac{L_r}{v_r} \gg T_r, \quad (14)$$

in which T_r is the time-resolution of the data. Blindly reducing L_r without taking this into account is likely to *decrease* the accuracy of an FTLE field.

The conclusions here are based on using a first-order method (10) to approximate the gradient of \mathbf{F} . Using higher-order methods will result in higher powers of L_r in the denominator of (13); for example, an n th order scheme will generate $L_r^{n+1/2}$. Consequently, $\mathcal{E} \sim v_r^{1/2} L_r^{-n-1/2}$, and the basic conclusions (error decays with v_r but increases with L_r) will continue to hold. Interestingly, higher-order schemes will result in *worse* behavior as L_r is decreased.

3. Validations and implementations. In this Section, several examples are examined numerically. The first example demonstrates the potential for spurious conclusions if treating the FTLE field as exactly known, and shows that taking




FIGURE 1. Double-gyre computations with $L_r = 0.005$ and $v_r = 2 \times 10^{-5}$: (a) FTLE field Λ , and (b) the error field \mathcal{E} .

the theoretical error into account avoids this. Probability distributions generated from stochastic simulations with 100,000 realizations validate the correctness of the error bounds. The second example examines in more depth the spatial variation of the FTLE field, in particular the spatial correlation of the FTLE error field and statistically computed FTLE ‘spread’ indicators based on stochastic simulations. Furthermore, the dependence of these statistical measures on v_r and L_r is examined, validating the direct and inverse square-root relationships implication from the theoretical error (13). The third example shows the computability of the FTLE error field for oceanographic velocity data.

3.1. Double-gyre. For these FTLE computations, the double-gyre flow [?] with $\mathbf{x} = (x_1, x_2)$, and velocity field given by

$$\mathbf{u}(x_1, x_2, t) = \begin{pmatrix} -\pi A \sin(\pi f(x_1, t)) \cos(\pi x_2) \\ \pi A, \cos(\pi f(x_1, t)) \sin(\pi x_2) \frac{\partial f}{\partial x_1}(x_1, t) \end{pmatrix},$$

in which

$$f(x_1, t) = \varepsilon_{dg} \sin(\omega t) x_1^2 + (1 - 2\varepsilon_{dg} \sin(\omega t))$$

is used. Highly popular in Lagrangian coherent structure analysis (see citations in [?], the double-gyre is defined on the invariant rectangle $[0, 2] \times [0, 1]$. In keeping with the original simulations [?], the parameter values $A = 0.1$, $\omega = 2\pi/10$, $\varepsilon_{dg} = 0.1$, $t_0 = 0$ and $T = 15$ will be used throughout. The same spatial resolution, L_r , will be used in the x - and y -directions.

First, let $L_r = 0.005$ and take $v_r = 2 \times 10^{-5}$. Since the velocity scale in the problem can be considered order-unity, this corresponds to the velocity field being known to extraordinary accuracy (a highly unlikely scenario in realistic situations). Fig. 1 shows the result of simulations for this case, with (a) displaying the computed FTLE field, and (b) the error field. Nearby errors, being less than 0.02, indicate that the blue ridge (at value ~ 0.35), remains sufficiently distinct from adjacent regions (with values ~ 0.2 or lower) even when the uncertainty in the FTLE values is taken into account. With the aid of the error field, one can therefore conclude that the ridge is genuine, and not an artefact of the resolution and velocity uncertainties. The two almost elliptical regions in (b) centered near $(0.5, 0.6)$ and $(1.5, 0.4)$ with approximate radius 0.1, are likely to be genuinely coherent regions




FIGURE 2. As in Fig. 1, but with $L_r = 0.05$ and $v_r = 0.002$.

of small stretching, since the uncertainties are small. On the other hand, there is slightly less confidence in the two regions (‘two-pronged’ lightest yellow patches) of comparably small FTLE values in (a), near $(0.4, 0.1)$ and $(1.85, 0.85)$, because the errors from (b) are about 0.08. On the whole, this highly-resolved FTLE field and its error field indicate a very high certainty in the results.

Any realistic data set is unlikely to possess the level of resolution above, with lengthscale 0.5% of the system lengthscale. Thus, the same calculations are shown in Fig. 2, with $L_r = 0.05$ (a ten-fold reduction in resolution, yet still 5% of the system lengthscale) and $v_r = 0.002$ (a Eulerian velocity uncertainty remaining generous at 0.2%). The values of the FTLE field in (a) have not changed that much in comparison to the higher resolution of Fig 1, but the spatial variation is instructive. Firstly, the ‘smudges’ in the FTLE field in (a) (despite the smoothness of the velocity field) highlights that understanding errors is paramount when dealing with realistic data. In Fig. 2(a), while the actual ridge is ill-defined, the tiny region of low Λ near $(1.05, 0.1)$ seems inconsistent; this is where Fig. 1(a) identifies the well-defined ridge associated with the stable manifold whose ‘beginning’ is just about here. Now Fig. 2(b) has correctly identified that there is an uncertainty as high as 1.4, which is significantly more than the computed FTLE values. Thus, having the uncertainty theory helps analyse the accuracy of the FTLE calculations. There are other regions which are identified by (b) as having high uncertainty (about 0.3); these are near the ‘two-pronged’ regions.

To investigate in detail the adequacy of the error field of Fig. 2, stochastic simulations are now performed using 100,000 realizations of the stochastic differential equation (6). The value $\varepsilon = \sqrt{L_r v_r}$ is used, with $L_r = 0.05$ and $v_r = 0.002$. The numerics are performed using the Euler-Maruyama algorithm [?] with $\Delta t = 0.01$. To obtain the most realistic conditions, within each of the 100,000 simulations, each initial condition is assumed to be subject to *independently chosen* Brownian motions (as opposed to choosing one initially, and then applying it to all initial conditions as is the understanding in the ‘random dynamical systems viewpoint [?]). Since nearby trajectories evolve with independent uncertainties, this results in flow map gradients having high variability. For each realization, having computed the final location of all the initial conditions, the flow map gradient is computed, and thereby the FTLE field determined for each realization. Fig. 3 shows the results

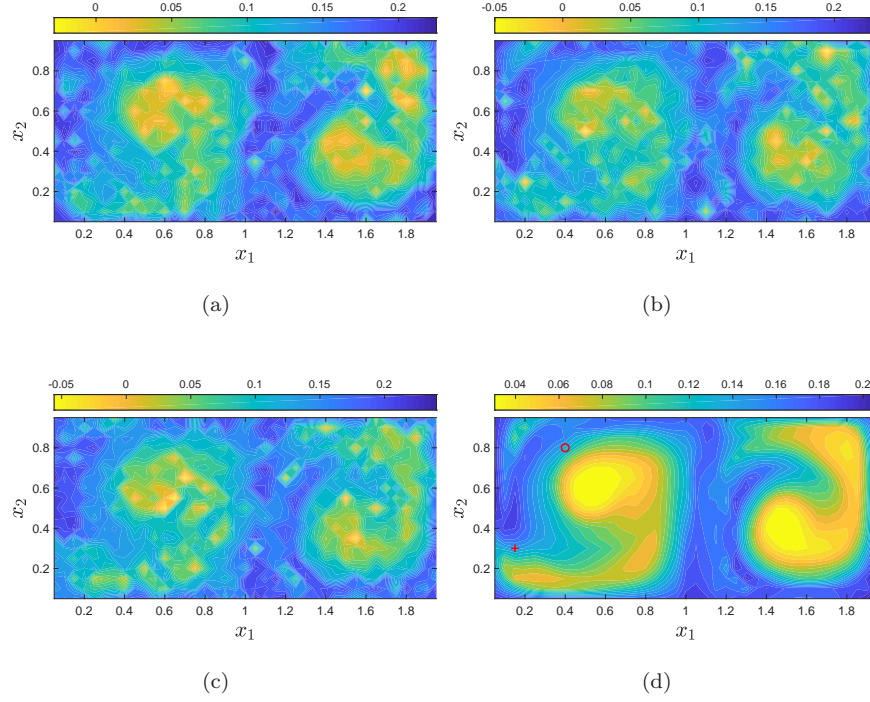


FIGURE 3. (a,b,c) Three (of the 100,000) sample FTLE fields, and (d) the mean FTLE field, computed when using stochastic simulation of (6) for the double-gyre system with $L_r = 0.05$ and $v_r = 0.002$.




FIGURE 4. (a) The standard deviation of the FTLE fields computed from the 100,000 stochastic simulations, and (b) the theoretical error field (the same as Fig. 2(b), but scaled to elucidate variation at values of FTLE comparable to (a)).




FIGURE 5. The probability density of the FTLE value from the stochastic simulations of the double-gyre at two selected points indicated in Figs. 4(a) and 3(d): (a) maximum standard-deviation point (at red ‘x’), and (b) minimum standard deviation point (at red circle).

of computing the FTLE field for three of the 100,000 realizations. Some common features seem to be indistinctly seen.

The mean (average) of the 100,000 FTLE fields displayed in Fig. 3(d). The mean field captures many of the features of the highly-resolved Fig. 1(a), with respect to the heterogeneity and magnitude. However, smaller-scale details are not obtained, in particular, the two-pronged valleys and (more importantly) the *sharp* FTLE ridge. The FTLE ridge *is* approximately obtained as a softer ridge. The lack of sharpness is of course understandable, because it is obtained from averaging stochastic simulations. Using a smaller ε value would sharpen it; however, since $\varepsilon = \sqrt{L_r v_r}$, this would physically mean having better spatial resolution, or more certainty in the velocity field. This is not always controllable in a realistic system, and so obtaining a picture such as Fig. 3(d) is probably the best one can expect. It is clear that performing stochastic simulations, thereby allowing for uncertainties to be present but then averaging them out, provides a significantly better approach than simply computing the FTLE by pretending that there are no uncertainties (compare Figs. 3(d) and 2(a)).

A computational proxy for the uncertainty in the FTLE field would be the standard deviation (obtained at each point by computing the standard deviation of the 100,000 FTLE values from the samples). This standard-deviation field computed from the stochastic simulations is shown in Fig. 4(a). The red ‘x’ is the point corresponding to the maximum value, and the red circle to the minimum value, of the standard deviation. These two points are also similarly labelled in Fig. 3(d). The comparison of Fig. 4(a) should be to the theoretical error as shown in Fig. 2(b). Unfortunately, the preternaturally large value near (1.05, 0.1) renders the smaller variations in the field to not be observable. Thus, the same error field of Fig. 2(b) is displayed in Fig. 4(b), but with the scaling changed to be comparable to the standard-deviation in Fig. 4(a). The patchiness of the theoretical error field is inevitable, given the relatively low resolution here. The positioning and size of the smaller error regions (yellow) in the two subfigures of Fig. 4—the first from




FIGURE 6. The probability density of the FTLE value from the stochastic simulations of the double-gyre with $L_r = 0.05$ and $v_r = 0.002$; the red star is the value from the standard FTLE computation: (a) at $(1.05, 0.1)$ and (b) at $(1.15, 0.1)$.

stochastic simulations and the second from the theoretical error— are consistent. Therefore, the theoretical error is once again important.

However, the subfigures are different in many aspects. The reason is anticipated to be the fact that the distributions of the FTLE values are *not* symmetrical in general. To understand this further, the point at which the standard deviation is maximum (red ‘x’ in Figs. 4(a) and 3(d)) is identified, as is the point at which the standard deviation is the least (red circles in the same figures). It should be mentioned that the location of these points in the FTLE field in Fig. 3(d) has no obvious relationship to ‘standard’ points of interest (ridges, flat areas) in the field. The empirical probability distribution of the FTLE values at these points, from the 100,000 stochastic simulations, is displayed in Fig. 5. Panel (a) corresponds to the maximum standard deviation point, and as expected, there is a large spread in the FTLE values. The red star in the figure corresponds to the deterministically computed FTLE value, and the arrows emanating in the two directions show the uncertainty interval of the FTLE value from the theoretical expression. There is a strange structure in the FTLE distribution, with a narrow but very high spike near 0.2, but a solid peak near 0.14. The theoretical error bounds very nicely encapsulate the extent of the FTLE distribution, but do not give any insights into the nature of this distribution. The FTLE distribution at the (red circle) point of minimum standard deviation, shown in (b) has a very pronounced but asymmetric peak. The uncertainty interval—which is symmetric by construction— captures the larger spread (towards the left), thereby overestimating the uncertainty interval towards the right. In general though, the theoretical uncertainty interval is strongly validated by the stochastic simulations here.

Are the stochastic simulations consistent with the error bound near the point $(1.05, 0.1)$, near where the FTLE field produced an unexpectedly shallow trough? The empirical probability distribution for the FTLE field at this point is illustrated in Fig. 6(a), with the computed FTLE value once again as a red star, and the interval indicated by the red line. In this case, the theoretical error is enormous (specifically, 1.3131), and so the theoretical FTLE interval extends far beyond the entire range of the stochastic simulations. The stochastic simulations return an average FTLE value of around 0.21; a large discrepancy with the deterministic FTLE computation yielding -0.04 . That is, large stretching is to be expected, in contrast to shrinking as predicted from the deterministic FTLE. The theoretical




FIGURE 7. Wobbly Duffing from from $t_0 = 4$ with time-of-flow $T = 4$, with $L_r = 0.2$ and $v_r = 0.002$: (a) FTLE field, and (b) FTLE error field.

error bound remains correct, but is a large overestimate in this case. Reasons for why the FTLE computation returns such a ‘bad’ value near here at his low resolution, but is not too unexpected in the remainder of the domain is not obvious. Nevertheless, the theoretical error bound correctly identifies this as a highly suspect value (as highlighted in Fig. 2(b)), demonstrating the usefulness of the uncertainty procedure. The corresponding picture at the nearby point (1.15, 0.1) in Fig. 2(b) displays how the FTLE distributions can take on complicated shapes, but once again the uncertainty bounds work well.

An observation related to the asymmetry of the FTLE distributions is in order. The theoretical errors were initially obtained for the *stretching field*, whose logarithm needed to be taken to compute the FTLE. If one assumes that the errors are symmetric in the stretching field, this would imply an *asymmetry* in the FTLE because of this logarithm. Indeed, the Taylor approximation $\ln(1 \pm b) \sim \pm b$ used in the derivation of the error bound ignores this asymmetry, because b is assumed small. The logarithm would tend to give a longer tail in the negative direction, with sharper descent in the positive direction, for the distribution. This is observed in Figs. 5 and 6, and in many other calculations that were performed.

3.2. Wobbly Duffing. To generalize the previous example, an unsteady system in which the flow is not time-periodic, bounded or area-preserving is next considered to demonstrate that these issues are not impediments to the calculations. This example however is one in which the stable and unstable manifolds are known *exactly* if considered as an infinite-time flow. In finite-time situations, defining these entities becomes ambiguous [?, ?]. The unsteady Eulerian velocity field is [?]

$$\mathbf{u}(x_1, x_2, t) = \frac{1}{h^2 + m^2} \begin{pmatrix} h\dot{h} + m\dot{m} + 2mh & h^2 - m^2 + \dot{m}h - m\dot{h} \\ h^2 - m^2 - \dot{m}h + m\dot{h} & h\dot{h} + m\dot{m} - 2mh \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - \left(\frac{hx_1 - mx_2}{h^2 + m^2} \right)^3 \begin{pmatrix} m \\ h \end{pmatrix},$$

where the choice

$$h(t) = 1 + a \tanh t \cos(\omega t) \quad \text{and} \quad m(t) = -a \tanh t \sin(\omega t)$$

is made, and the overdot denotes the time-derivative. (Different h and m functions are permitted; the following information still holds within appropriate parameter regimes [?].) As explained in [?], the velocity field is obtained by converting the classical Duffing system [?, ?] to a frame of reference which rotates in a ‘wobbly’ fashion. The ‘figure-8’ structure of the heteroclinic manifolds (surrounding two coherent eddies) of the classical steady system [?, ?] lies on the wobbly curve

$$[m(t)x_1 + h(t)x_2]^2 = [h(t)x_1 - m(t)x_2]^2 \left[1 - \frac{1}{2} \left(\frac{h(t)x_1 - m(t)x_2}{h(t)^2 + m(t)^2} \right)^2 \right] \quad (15)$$

for any choice of $a \in (-1, 1)$ or ω (and indeed any consistently taken choice of the functions h and m), as time t progresses [?]. However, restricting to finite time means that one can only expect to see these structures in some approximate way from the FTLE field. The parameters used throughout this Section are $a = 0.4$, $\omega = 8$, $t_0 = 4$ and $T = 4$, and the domain at time t_0 is always the rectangle $(x_1, x_2) \in [-2.5, 2.5] \times [2.0, 2.0]$.

With the choice $L_r = 0.2$ and $v_r = 0.002$, the computed FTLE field and its error are shown in Fig. 7. The available spatial resolution L_r is used in evaluating (15), and is overlaid in red. There is a blue ridge emanating from near $(0, 0)$ in the FTLE field, which straddles the stable manifold. The FTLE error plot in (c) indicates that errors on and near the ridge are around 0.05, and since the ridge exhibits a sharp drop in FTLE value of around 0.4 as one goes across it, the identification of this ridge in (a) is justifiable. If v_r is changed (not shown), the only difference that occurs is that the scaling of (b) changes; increasing the velocity uncertainty can therefore change this conclusion.

The FTLE field also identifies low stretching regions, which are expected to be related to low mixing. In Fig. 7(a), there are two triangular yellow regions (the valleys) which are such low regions. With our prior knowledge of the existence of the (red) manifolds cutting these, the indications are that where the manifolds cross these regions, there must be some level of stretching (since these regions will eventually get towards the moving hyperbolic trajectory located at the point of intersection of the heteroclinic manifolds). However, because the flow considered is for finite time, it is not clear whether $T = 4$ is sufficient to obtain a high stretching rate in this region. On the other hand, having the FTLE field identify this as a *low* stretching region may be suspect under conditions of velocity uncertainty, given our prior knowledge. Access to the error plot in (b), however, identifies exactly these regions (darkest blue regions) as being places where the FTLE field could have an error of ± 0.2 , and hence the FTLE values may *not* be that small. An intuitive understanding of the reason for this error is that exponential stretching is only experienced in a small neighborhood of the hyperbolic trajectory [?]. The FTLE uncertainty is likely capturing the fact that the uncertainty in the location of evolved particles is on the boundaries of such a neighborhood. In contrast, for regions of low FTLE which are well within the manifold structure, the error is not as large, and these are identifiable as being more likely to have low stretching (and consequently more coherent).

How well does the theoretical error stand up to stochastic simulations? To understand this further, stochastic simulations were performed as in Section 3.1, and using 1000 realizations. If the theoretical error is a good indicator of the stochastic simulations, then there must be a good correlation between it and measures for the error obtained from the stochastic simulations. For each data point in the spatial




FIGURE 8. Scatter plots of the theoretical error with (left) the standard deviation, and (right) the range, obtained from 1000 stochastic simulations. Each point (red circle) in the plot corresponds to a point in the spatial domain at time t_0 , and the blue line is a linear fit. The values chosen are (top) $L_r = 0.2$ and $v_r = 0.002$, and (bottom) $L_r = 0.05$ and $v_r = 0.00002$.

domain, the standard deviation and the range (from the 1000 simulations) was computed. These are two different ways of computing an ‘error’ or ‘spread’ from the stochastic simulations. However, associated with every data point, there is available a theoretical error (illustrated for example in Fig. 7(b)). Thus, it is possible to view the correlation—across all the spatial data points—between each of the two error measures from the stochastic simulations and the theoretical error. These scatter plots are shown in Fig. 8 for two choices of the parameters: (top) $L_r = 0.2$ and $v_r = 0.002$, and (bottom) $L_r = 0.05$ and $v_r = 0.00002$. The linear correlations are ~ 0.79 for the upper plots, and ~ 0.94 for the lower ones. This means that under some choices of the parameters L_r and v_r , there is *excellent* spatial correlation between the theoretical, and stochastically-obtained, FTLE. (In most cases examined the correlations were greater than 0.78, but in certain parameter regimes, they were well above 0.94.) In other words, the FTLE error field developed here offers a persuasive measure of the nonhomogeneous distribution of errors in any computed FTLE field.

A deeper analysis of how the errors change with the velocity uncertainty v_r is undertaken in Fig. 9. With $L_r = 0.2$ kept fixed, and a choice of v_r , 1000 stochastic

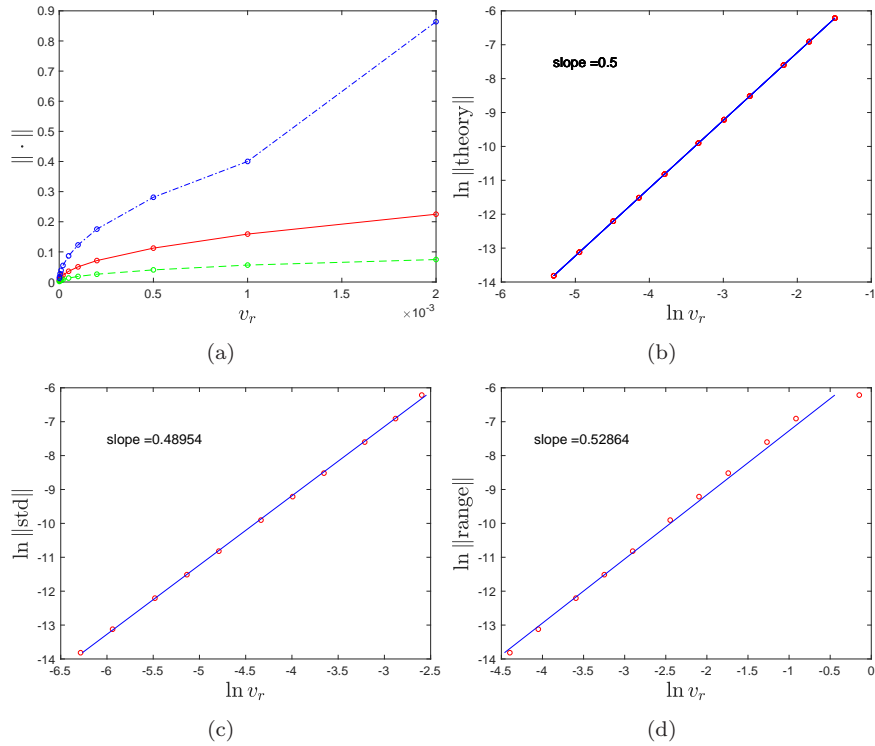


FIGURE 9. The v_r dependence analyzed for the theoretical and stochastically-determined errors in the FTLE field, for the wobbly Duffing system with $L_r = 0.2$: (a) theoretical (red-solid), standard deviation (green-dashed) and range (blue-dot-dashed) norms, and (b-d) investigations on the dependence of v_r on the FTLE spreading measures $||\text{theory}||$, $||\text{std}||$ and $||\text{range}||$ respectively.

simulations were performed for a range of values of v_r . Each point in the spatial domain has a standard deviation of the FTLE field which can be computed based on this; a *standard deviation norm* across the domain is obtained by taking the maximum of these standard deviation values. Similarly, each point has a range of FTLE based on the stochastic simulations, from which a *range norm* is constructed by taking the maximum across all spatial points. These shall be denoted by $||\text{std}||$ and $||\text{range}||$ respectively, while the *theoretical norm* is $||\text{theory}|| = \sup_{\mathbf{x}_0} \mathcal{E}(\mathbf{x}_0)$. This procedure was followed for a range of v_r . The variation of the three norms as v_r is changed in displayed in Fig. 9(a). The theoretical error norm lies between the standard deviation and the range norms for all values of v_r . Indeed, its value is about double the standard deviation of the FTLE field computed from the stochastic simulations. The second inequality used in (11) in estimating the error is likely to be a highly conservative one, and the indications from the current computational work is that perhaps a better empirical estimate of the standard deviation of the FTLE field would be given by $\mathcal{E}(\mathbf{x}_0)/2$ in general. This observation is also consistent with the double-gyre simulations shown in Fig. 4. In any case, providing a single number

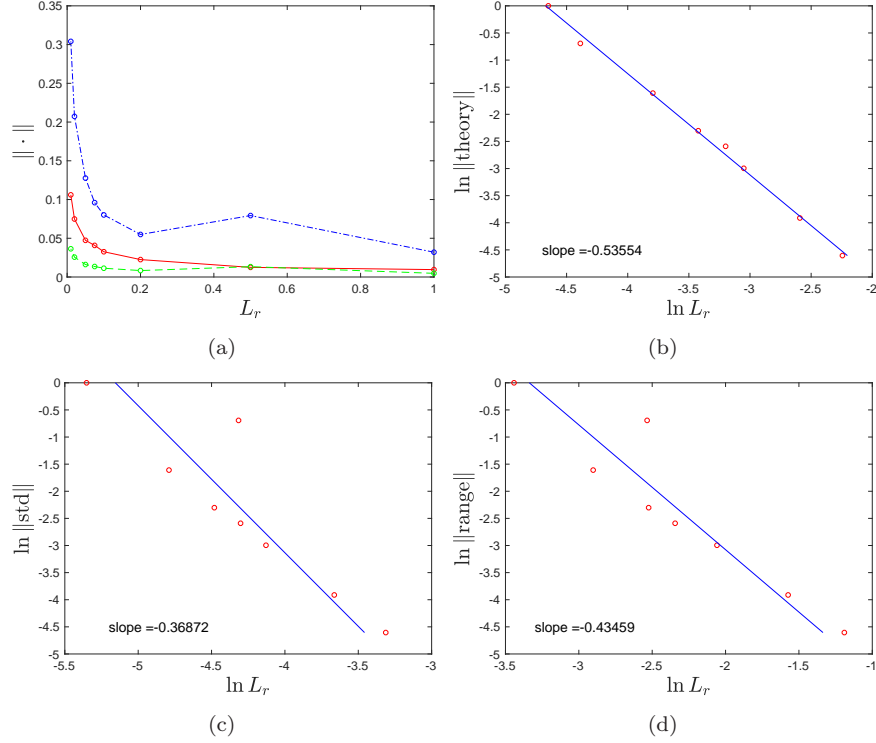


FIGURE 10. The L_r dependence analyzed for the theoretical and stochastically-determined errors in the FTLE field, for the wobbly Duffing system with $v_r = 0.00002$: (a) theoretical (red-solid), standard deviation (green-dashed) and range (blue-dot-dashed) norms, and (b-d) investigations on the dependence of v_r on the FTLE spreading measures $\|\text{theory}\|$, $\|\text{std}\|$ and $\|\text{range}\|$ respectively.

for the ‘spread’ of a probability distributions when it can have forms such as shown in Figs. 5(a) and 6(b) is ambiguous, and \mathcal{E} or $\mathcal{E}/2$ has been shown to be effective.

The theoretical error $\|\text{theory}\|$, based on the expression (13), is expected to depend on v_r in the form $\sqrt{v_r} = v_r^{1/2}$. For the values of v_r considered here, this is indeed the case (Fig. 9(b)). One way of evaluating the theory is to see whether stochastic measures for the FTLE’s spread, notably $\|\text{std}\|$ and $\|\text{range}\|$, also exhibit this dependence. Figs. 9(c) and (d) give strong evidence that both measures have consistent dependence with the theoretical bound, thereby adding credence to the theory.

The theoretical error (13) exploded the naive expectation that the FTLE would become *more* accurate by decreasing L_r . To investigate this, exactly the same type of analysis as for Fig. 9 is performed for changing L_r , and shown in Fig. 10. While (a) suggests that the theoretical error is about double the standard deviation computed from the stochastic simulations (as previously), (b) demonstrates that the theoretical error decreases with L_r roughly as $1/\sqrt{L_r}$. Since this was computed from the theoretical expression (13), this is no surprise. However, the lower panels are based on *stochastic simulations*, computed by ignoring anything to do with

(13). Both the standard deviation and range norms decrease with L_r at roughly the same exponent, indicating that higher resolution (with other parameters kept fixed) lowers the FTLE error. This—initially surprising—result is perfectly consistent with the theory developed here.

3.3. Oceanographic FTLEs. Next, *realistic* data from an unsteady flow is analyzed from this viewpoint. Unlike in the previous cases, the absence of a known truth impedes the validation of the results. However, it is shown that the error in the FTLE field can be computed, thereby allowing any results from the field to be treated within an uncertainty range, as would be scientifically appropriate. The data used is that available from the European Commission’s Copernicus Marine Environment Monitoring Service, comprising ocean surface geostrophic velocities in the Gulf Stream extension region based on sea-surface height observations. It has a resolution of 0.5 degrees (in both latitude and longitude), and the reported daily geostrophic velocity data is computationally not exactly divergence-free. The intention of this Section is to demonstrate the *computability* of the FTLE error bound in such realistic situations to complement the FTLE field, rather than to reach any particular scientific conclusions. Future work in using this tool is to be anticipated.

Fig. 11 demonstrates some examples of the calculations, with the axes being the longitude and latitude. The left panels are the FTLE fields (in units of day^{-1}) under the stated time-ranges, and the right panels are the corresponding FTLE error fields with the choice $v_r = 0.001 \text{ deg/day}$ (the typical velocity scale is in the data $\sim 1 \text{ deg/day}$). The white regions are the north American continent. The darker areas—appearing as curves in some cases—in the FTLE fields (left pictures) are regions of large stretching, while the yellow patches are sometimes construed as eddy-like structures on the flanges of the Gulf Stream because they have low stretching. While the FTLE field demonstrates some smoothness and features, the fact that the uncertainty field has a ‘patterned’ nature is a signal: since the data has resolution of 0.5 degrees, regions which are away from the gridpoints must have uncertainties associated with them. The low error bands (in yellow) in the error plots do indeed conform to the locations of gridpoints. The FTLE field contours—obtained as they are by automatic smooth interpolation by the software used (Matlab in this case)—hides this fact.

The general observation from these calculations (and others not shown) is that the FTLE fields determined from this data must be treated with much caution. The spatial resolution appears to be a significant stumbling block in inferring smoothly defined ridges and regions of low stretching from the left panels in Fig. 11, *purely based on this data*.

4. Concluding remarks. FTLE fields are often computed in realistic flows, and used to reach conclusions regarding stretching and transport of fluid. The Eulerian data that is used in such endeavors can never be known accurately on the full spatial domain because of resolution and measurement/modelling uncertainties. Until now, there has not been a systematic way to evaluate the accuracy of FTLE fields. This article does so in the context of two-dimensional finite-time (potentially non-area-preserving) flows, by computing an uncertainty interval (akin to providing error-bars at each spatial location) for the FTLE field. While the theoretical development in Section 2 works in three dimensions as well, at present, there only exists a methodology for computing the required S^2 -field in two dimensions [?].




FIGURE 11. FTLE fields (left) and FTLE error fields (right) for several situations based on the Copernicus data: (top) from 30 January 2016 to 29 February 2016 (middle) 25 February 2011 to 26 April 2011, and (bottom) 25 February 2011 to 13 September 2011.

To validate the results, it was necessary to examine systems in which detailed information was known. The impact, in particular, of identifying FTLE ridges and valleys was viewed through this lens. Using the FTLE uncertainty was shown to be necessary to avoid reaching unjustified conclusions. In realistic data—in which the

ground truth is unknown—the FTLE uncertainties developed here should be used to mollify any conclusions reached from a standard computation of the FTLE field.

While the results presented here have mostly assumed that the time-of-flow T were positive, the results are equally valid for $T < 0$. This corresponds to the computation of backwards-time FTLE fields. The error expression (13) is therefore also valid for $T < 0$ (and the required expressions for S^2 as given in Appendix A have also been written to be correct even for $T < 0$).

The investigations here indicate that performing many stochastic simulations of (6) and then averaging them will provide a better estimate for the FTLE field than simply computing it using (1), i.e., as if it were deterministic. This supports the work of Guo et al [?] who provided several ways of stochastically averaging the FTLE field. Doing so, however, is computationally expensive, since a multitude of stochastic simulations needs to be performed for each and every point in the spatial domain. Another slight issue is that, even if FTLE ridges are relevant and exist in the flow, averaged stochastic simulations will make them less sharp, and hence more difficult to identify. (This is as it should be: uncertainties *do* mean that the ridge is less certain [?].) Having a theoretical error field—computable directly from the deterministic data, and incorporating the spatial resolution L_r and Eulerian velocity uncertainty scale v_r —is a more desirable approach. The stochastic simulations performed here, and their comparison to the theoretical FTLE error, provide good evidence to support this approach. The spatial heterogeneity of the error field is particularly important in reaching conclusions based on the features observed in the FTLE field, and this is captured well by the theoretical error field. Thus, it is hoped that this paper will help generate more informed conclusions from FTLE fields computed in the future.

An important caution, implied by the theoretical error (13) and subsequently demonstrated in stochastic simulations, is that simply reducing the spatial resolution does *not* make a computed FTLE field more accurate. The dependence on the spatial resolution L_r and the Eulerian velocity uncertainty v_r is more subtle; the error reduces as $\sqrt{v_r/L_r}$. (If ignoring the velocity uncertainty, of course, the leading-order dependence on L_r is as expected; the error decreases with L_r .) The intuitive reason for the inverse dependence with L_r is that in computing the stretching (3), the numerator possesses an uncertainty of $\mathcal{O}(\sqrt{L_r})$ which interacts with the $\mathcal{O}(L_r)$ in the denominator. There is no evidence in the literature that this fact has as yet been realized, despite the multitudinous computations of FTLE fields across many disciplines. A simple sanity check on whether the FTLE field is believable is given in (14): the spatial resolution should be *significantly larger* than the velocity uncertainty times the temporal resolution. The ‘best’ resolution for usage must be chosen based on an assessment of the velocity uncertainty v_r as well, and needs to be sufficiently large in comparison with $v_r T_r$, where T_r is the temporal spacing.

Appendix A. Stochastic sensitivity and its computation. The theory developed in [?] is associated with a more general system than (6): it additionally includes a dimensional spatiotemporally-dependent diffusion matrix. This has been set to the identity, and its dimensions transferred to ε , in explaining the results for the special case of (6) in this Appendix.

Suppose that an initial condition \mathbf{x}_0 at time t_0 progresses to a point \mathbf{X} at time $t_0 + T$ according to the deterministic flow (1). For this deterministic flow, let $\mathbf{F}_{t_1}^{t_2}$ be the flow map from a time t_1 to t_2 , each of which is in the interval $(t_0, t_0 + T)$

(this notation allows for t_2 to be less than t_1 ; this would then be flow in backwards time according to (1)). Thus $\mathbf{X} = \mathbf{F}_{t_0}^{t_0+T}(\mathbf{x}_0)$ and $\mathbf{x}_0 = \mathbf{F}_{t_0+T}^{t_0}(\mathbf{X})$.

For a trajectory \mathbf{x}_t evolving according to the *stochastic* flow (6) with initial condition \mathbf{x}_0 , define the random two-dimensional deviation vector

$$\mathbf{z} = \frac{\mathbf{x}_{t_0+T} - \mathbf{x}(t)}{\varepsilon} = \frac{\mathbf{x}_{t_0+T} - \mathbf{F}_{t_0}^{t_0+T}(\mathbf{x}_0)}{\varepsilon},$$

Define a general vector $\hat{\mathbf{n}}$ pointing in a direction associated with the polar angle θ from the point \mathbf{X} , and the orthogonal rotation matrix J , by

$$\hat{\mathbf{n}} = \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \quad \text{and} \quad J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

Let \mathbb{E} and \mathbb{V} be the expectation and variance operators, with respect to many realizations of the stochastic dynamics (6). Then, the relevant results from [?] are that

$$\lim_{\varepsilon \rightarrow 0} \mathbb{E}[\mathbf{z}] = \mathbf{0} \quad \text{and} \quad \lim_{\varepsilon \rightarrow 0} \sup_{\theta} \sqrt{\mathbb{V}[\mathbf{z} \cdot \hat{\mathbf{n}}]} = \sqrt{S^2(\mathbf{x}_0)}, \quad (16)$$

where the stochastic sensitivity field S^2 is computed by the following procedure. Construct the 2×2 matrix

$$H(\mathbf{X}, t) := \exp \left[\int_t^{t_0+T} [\nabla \cdot \mathbf{u}] \left(\mathbf{F}_{t_0+T}^\xi(\mathbf{X}), \xi \right) d\xi \right] J \nabla \mathbf{F}_{t_0+T}^t(\mathbf{X}),$$

and define in terms of its components

$$\begin{aligned} P(\mathbf{X}) &= \left| \frac{1}{2} \sum_{i=1}^2 \sum_{j=1}^2 \int_{t_0}^{t_0+T} H_{ij}^2(\mathbf{X}, t) dt \right|, \\ L(\mathbf{X}) &= \frac{1}{2} \int_{t_0}^{t_0+T} \left[\sum_{i=1}^2 H_{i2}^2(\mathbf{X}, t) - \sum_{i=1}^2 H_{i1}^2(\mathbf{X}, t) \right] dt, \\ M(\mathbf{X}) &= \int_{t_0}^{t_0+T} \sum_{i=1}^2 [H_{i1}(\mathbf{X}, t) H_{i2}(\mathbf{X}, t)] dt, \text{ and} \\ N(\mathbf{X}) &= \sqrt{L^2(\mathbf{X}) + M^2(\mathbf{X})}. \end{aligned}$$

Each of these quantities can be computed for a given \mathbf{X} easily; the complication (in comparison to a standard FTLE computation) is that the flow gradient needs to be evaluated at all instances in time as opposed to only at the final instance. This is necessary in computing H , and then its elements need to be appropriately integrated over time to generate P , L , M and N . Having computed these, the stochastic sensitivity $S^2(\mathbf{x}_0)$ on the t_0 -space is now [?]

$$S^2(\mathbf{x}_0) = P(\mathbf{X}) + N(\mathbf{X}). \quad (17)$$

Since S^2 at \mathbf{x}_0 is given in terms of its deterministic image \mathbf{X} at time $t_0 + T$, there are several ways in which S^2 can be computed as a field over the space of initial conditions \mathbf{x}_0 . One is to take a grid of particles \mathbf{x}_0 at the time-slice t_0 , follow deterministic advection and find \mathbf{X} , and then use deterministic advection backwards from \mathbf{X} while keeping track of nearby trajectories, hence computing $H(\mathbf{X}, t)$ at each time and then performing the necessary time integrals. In doing so—and indeed in all the other methods to be described—it is necessary to interpolate available velocity values to the inevitable non-grid locations that the particles venture. The

second method is to take a uniform grid of \mathbf{X} in the time-slice $t_0 + T$ and perform the backwards advection as described above; this will generate information on a *nonuniform* grid of \mathbf{x}_0 in the time-slice t_0 . Then, one can use an operation such as Matlab's `gridfit` algorithm to extend to a uniform grid to enable visualization. A third method is to take the points \mathbf{X} at time $t_0 + T$ generated from forward advection, supplement this nonuniform set with their nearest neighbors (2^n points at distance L_r in each coordinate direction when doing the computations in \mathbb{R}^n). By advecting this *nonuniform* collection backwards in time, it is possible to keep track of the various gradients required. The second approach was followed in the analytical examples described here, while the third was used for the oceanographic data set because it was more robust in preventing particles landing on land (at which the oceanic velocities are undefined). There is clear scope of determining more efficient and robust ways of computing S^2 .

It is noted that the uncertainty lengthscale $\varepsilon\sqrt{S^2(\mathbf{x}_0)}$ (as presented in (9)) arises because $\varepsilon\mathbf{z}$ gives the deviation in location. By projecting this over all possible directions $\hat{\mathbf{n}}$ and then maximizing the variance over all such directions, and then taking the square-root (to convert a variance to a standard-deviation) yields the result, valid to leading-order in the small quantity ε .

Acknowledgments. The author is grateful to Irina Rypina for providing access to the Copernicus oceanographic data. Support from the Australian Research Council (grants FT130100484 and DP200101764) and the Australian Department of Education and Training (Endeavour Research Leadership Award) are also gratefully acknowledged.

E-mail address: sanjeeva.balasuriya@adelaide.edu.au