

Hyperbolic neighborhoods as organizers of finite-time exponential stretching

Sanjeeva Balasuriya^{1†}, Rahul Kalampattel², and Nicholas T. Ouellette^{3‡}

¹School of Mathematical Sciences, University of Adelaide, Adelaide, SA 5005, Australia

²School of Mechanical Engineering, University of Adelaide, Adelaide, SA 5005, Australia

³Department of Civil and Environmental Engineering, Stanford University, Stanford, CA 94305, USA

(Received xx; revised xx; accepted xx)

Accepted for publication in *Journal of Fluid Mechanics*, 26 September 2016

Hyperbolic points and their unsteady generalization—hyperbolic trajectories—drive the exponential stretching that is the hallmark of nonlinear and chaotic flow. In infinite-time steady or periodic flows, the stable and unstable manifolds attached to each hyperbolic trajectory mark fluid elements that asymptote either towards or away from the hyperbolic trajectory, and which will therefore eventually experience exponential stretching. But typical experimental and observational velocity data is unsteady and available only over a finite time interval, and in such situations hyperbolic trajectories will move around in the flow, and may lose their hyperbolicity at times. Here we introduce a way to determine their region of influence, which we term a hyperbolic neighborhood, that marks the portion of the domain that is instantaneously dominated by the hyperbolic trajectory. We establish, using both theoretical arguments and empirical verification from model and experimental data, that the hyperbolic neighborhoods profoundly impact the Lagrangian stretching experienced by fluid elements. In particular, we show that fluid elements traversing a flow experience exponential boosts in stretching while within these time-varying regions, that greater residence time within hyperbolic neighborhoods is directly correlated to larger Finite-Time Lyapunov Exponent (FTLE) values, and that FTLE diagnostics are reliable only when the hyperbolic neighborhoods have a geometrical structure that is ‘regular’ in a specific sense.

1. Introduction

Lagrangian particle trajectories in steady flows are strongly influenced by the presence of hyperbolic (saddle-like) stagnation points, and their attached stable (resp. unstable) manifolds consisting of particles that approach each stagnation point asymptotically in forward (resp. backward) time (Guckenheimer & Holmes 1983; Wiggins 1992). The same is true in idealized unsteady flows, but in this case the hyperbolic stagnation points generalize to *hyperbolic trajectories*, to which are also attached stable and unstable manifolds, although now all these entities move with time. Stable (resp. unstable) manifolds are flow separators in forward (resp. backward) time, demarcating, for example, boundaries between different coherent structures in flows (Rom-Kedar *et al.* 1990; Pierrehumbert 1991; Samelson 1992; Peacock & Haller 2013). Therefore, transport can be enhanced by

† Email address for correspondence: sanjeevabalasuriya@yahoo.com

‡ Email address for correspondence: nto@stanford.edu

breaking these barriers in various ways (Rom-Kedar *et al.* 1990; Wiggins 1992; Stroock *et al.* 2002; Grigoriev 2005; Mezić *et al.* 2010; Balasuriya & Finn 2012; Balasuriya 2010, 2015*b,a*).

Because of its deep connections to mixing and transport, there has been significant effort devoted to developing metrics for locating and characterizing hyperbolicity in fluid flows. Ideally, one desires finite regions with hyperbolic character, as these regions can be considered to be the most strongly mixing parts of the flow. Early work focused on the intuitive (although somewhat fraught (Haller 2005)) similarity between hyperbolicity and strain, and ellipticity and rotation, and used the local velocity gradient as an indicator for hyperbolic or elliptic structures. This approach led to, for example, the so-called Q -criterion, (Hunt *et al.* 1988), the Δ -criterion (Chong *et al.* 1990), and the λ_2 -criterion (Jeong & Hussain 1995), all of which consider (in slightly different forms) the local balance between strain and rotation. In two-dimensional flows, the analogous quantity is often known as the Okubo–Weiss criterion (Okubo 1970; Weiss 1991); and we note that higher-order corrections to the Okubo–Weiss criterion that involve the acceleration in the flow have also been developed (Basdevant & Philipovitch 1994; Hua & Klein 1998). All of these measures, however, have flaws as indicators of true hyperbolic and elliptic regions in unsteady flows (Haller 2005). These methods will all, for example, misidentify a pure straining field in a rotating frame of reference as a pure vortex.

Doing a better job of locating real hyperbolic and elliptic structures in relation to their transport implications for nearby fluid particles necessitates a Lagrangian framework, which casts the evolution of the flow into a form more explicitly compatible with dynamical systems theory. More importantly, Lagrangian analysis incorporates the unsteadiness of the flow field in an intrinsic and natural way, so that the structures found are guaranteed to be relevant for the evolution of the flow. A range of methods have been and continue to be developed to locate Lagrangian structures, including topological analyses based on braids (Allshouse & Thiffeault 2012; Budišić & Thiffeault 2015), sets whose boundaries retain their curvature (Ma & Bollt 2014), ergodic partitions and entropy (Budišić & Mezić 2012; Froyland & Padberg-Gehle 2012), clustering (Huntley *et al.* 2015; Froyland & Padberg-Gehle 2015), Perron–Frobenius transfer operator methods that help determine regions that ‘hold together’ (Dellnitz & Junge 2002; Froyland *et al.* 2010; Froyland & Padberg 2009), analysis of its adjoint, the Koopman operator (Mezić 2013), and a Lagrangian generalization of the Okubo–Weiss criterion via so-called mesochronic analysis (Mezić *et al.* 2010).

The most commonly used method for determining the most straining regions in a Lagrangian sense, however, is the Finite-Time Lyapunov Exponent (FTLE) and related techniques (Shadden *et al.* 2005; Allshouse & Peacock 2015; Huntley *et al.* 2015; He *et al.* 2016; Lee *et al.* 2007; Waugh *et al.* 2007; Johnson & Meneveau 2015; Balasuriya 2015*a*; Samelson 2013). Over a given time period, this method computes the stretching experienced by a fluid element placed at each location, and by plotting this stretching field, one can identify maximum regions. Typically, these maxima occur along ridges, which are then extracted and interpreted as the boundaries between coherent structures. The maximum stretching property is consonant with the similar property exhibited by stable and unstable manifolds (Guckenheimer & Holmes 1983) in infinite-time flows. For steady flows, these objects are regions in space that are attracted to or repelled from hyperbolic (saddle-like) stagnation points. In unsteady flows, they are time-varying regions that are attracted to or repelled from hyperbolic trajectories that move with time (Balasuriya 2012; Balasuriya & Padberg-Gehle 2014). Rather than FTLEs, Haller (2011, 2015) chooses to use a *different* property of stable and unstable manifolds in identifying what have become known as hyperbolic Lagrangian Coherent Structures (LCSs): namely

that they are co-dimension-one entities to which there is maximal attraction or repulsion over a finite time period of the flow. These structures are curves in two dimensions, or surfaces in three dimensions (Blazevski & Haller 2014). It should be noted that the term LCS is often (and confusingly) used in the literature to refer to *all* of the above characterizations, although it is known that they do not in general coincide (Haller 2015; Karrasch 2015).

In this article, we seek to locate and track regions in space in which exponential stretching of fluid elements occurs. However, a fundamental difficulty in analyzing *finite-time* unsteady flows (i.e., in any realistic data set obtained from observations or experiments), is that the growth of *any* smooth function can be bounded by a growth of the form $Ae^{\lambda t}$, where λ is a rate constant: for *any* given λ , one can simply choose A large enough. This property means that the stretching rate, λ , is ambiguous for finite-time flows. Methods for defining hyperbolic trajectories (in infinite time) usually rely on such exponential decay estimates (Coppel 1978; Balasuriya 2012), which are related to the stretching of nearby fluid elements; but how can this notion be extended to finite-time situations? There are continuing attempts to resolve this issue in the time domain, for example by setting $A = 1$ (Doan *et al.* 2012; Karrasch 2013); but such ad hoc choices are unsatisfying. Moreover, there can be situations in which there is (intuitively) exponential stretching over some times, and no stretching whatsoever over others. In other words, there are *time intervals* that might be associated with a particular form of exponential stretching. Revealing, or defining, this behavior in the time domain is therefore fraught with difficulties. We offer an alternative viewpoint.

Regardless of the precise definition we take, much of our intuition about the function these structures play in the flow dynamics comes from the analogy with stable and unstable manifolds (Haller 2015; Balasuriya 2016*b*). A long-term hope is that fluid transport optimization and control based on stable and unstable manifolds (Balasuriya & Finn 2012; Balasuriya & Padberg-Gehle 2014) might be adaptable to finite-time flows (Ouellette *et al.* 2016). But, although inspired by properties associated with stable and unstable manifolds, hyperbolic structures in finite-time unsteady flows possess quite different characteristics. In the first place, hyperbolic structures, however they are defined, may be ephemeral. Moreover, during their existence, they will drive exponential stretching only for fluid elements that are *sufficiently close*. In contrast to infinite-time flows, fluid elements near a repelling LCS or a forward-time FTLE ridge (the analogs of a stable manifold in each of the characterizations) in finite-time flows may not come close enough to the associated hyperbolic trajectory to feel its effects before the hyperbolic trajectory disappears. Therefore, what we are really interested in is the region of influence, as a function of time, of each hyperbolic trajectory during its lifetime.

Here, we introduce a new way of thinking about this zone of influence, arising directly from the methods of nonlinear dynamics. Hyperbolic points or trajectories are so called because the *linearized* flow around them produces trajectories that are locally hyperbolæ. We therefore define the *hyperbolic neighborhood* around one of these locations to be the region where the linearized flow is a good approximation to the full flow. By construction, the fluid elements in this neighborhood feel the influence of the nearby hyperbolic structure so strongly that they will experience immediate exponential stretching, with minimal effects from the rest of the flow field. Once elements leave this region, they will no longer be so strongly influenced. Thus, elements traversing the flow domain will experience ‘exponential boosts’ to their stretching as they pass through hyperbolic neighborhoods. The irregular spatial distribution of hyperbolic neighborhoods, their time evolution, and their appearance and disappearance therefore provide a template for the evolution of stretching in the flow field.

Below, in §2, we motivate and give a precise definition of these hyperbolic neighborhoods, and explain the role they play in organizing the stretching of fluid elements. Connections to FTLE and LCS methods are also discussed, since these techniques are of great current interest in fluid dynamics ranging from micro- to geophysical scales (Huntley *et al.* 2015; He *et al.* 2016; Ouellette *et al.* 2016; Samelson 2013; Johnson & Meneveau 2015; Balasuriya 2015a; Haller *et al.* 2016). In §3, we explicitly calculate the hyperbolic neighborhoods for several analytical examples, including the Duffing equation, the steady Taylor–Green flow, and the oscillating Taylor–Green flow. We demonstrate that residence time within hyperbolic neighborhoods is closely connected to the FTLE fields, and that fluid elements entering these neighborhoods markedly increase their stretching, and reduce it once they leave. We offer insights as to how the geometry of these neighborhoods impacts interpretations that one might reach by examining FTLE fields, and indeed show that the degeneracy of the neighborhoods can be held accountable for some instances in which FTLE fields are known to provide unexpected results. Then, to show that these neighborhoods are indeed computable from real data, we present measurements from a turbulent quasi-two-dimensional laboratory flow in §4. Finally, in §5, we summarize our results, discuss the current limits of our analysis, and outline what we see as the most promising paths forward.

2. Defining Hyperbolic Neighborhoods

2.1. Steady Flows: Exponential Boosts in Stretching

In developing our methodology for unsteady flows, it helps to first build intuition for a steady flow given by

$$\dot{x}_i = u_i(\mathbf{x}), \quad (2.1)$$

where the dot denotes differentiation along a Lagrangian trajectory. Stagnation points $\bar{\mathbf{x}}$ satisfy $u_i(\bar{\mathbf{x}}) = 0$. If $\frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t))$ has eigenvalues that are not imaginary, and with at least one with a positive real part and another with a negative real part, then $\bar{\mathbf{x}}$ is a hyperbolic point (Guckenheimer & Holmes 1983; Ouellette & Gollub 2008). It has attached to it stable and unstable manifolds with dimensions given by the number of real-part-negative and positive eigenvalues respectively (Guckenheimer & Holmes 1983; Balasuriya 2012). These emanate in the directions spanned by the corresponding eigenvectors (Guckenheimer & Holmes 1983), and their global structure organizes the dynamics of the flow (Perry & Chong 1987; Rom-Kedar *et al.* 1990). The eigenvalues and vectors give us this information because if the flow is linearized about the point $\bar{\mathbf{x}}$ by setting $\mathbf{x}(t) = \bar{\mathbf{x}} + \mathbf{y}(t)$, then Taylor expansions give

$$\dot{y}_i = y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}) + \mathcal{O}(|\mathbf{y}|^2), \quad (2.2)$$

and neglecting higher-order term tells us that the velocity is locally linear in \mathbf{y} . If $\mathbf{y}(0)$ were chosen in an eigenvector direction corresponding to an eigenvalue λ of $\frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}})$ (i.e., $\nabla \mathbf{u}(\bar{\mathbf{x}})$), then it is clear from equation (2.2) that $\mathbf{y}(t) = \mathbf{y}(0)e^{\lambda t}$ within the region in which the higher-order terms can be legitimately discarded. A situation associated with this reasoning is shown in figure 1, corresponding to a two-dimensional flow with one positive and one negative eigenvalue. The stable manifold W^s of $\bar{\mathbf{x}}$ comes in along the eigenvector corresponding to the negative eigenvalue, and the unstable manifold W^u is similarly associated with the positive eigenvalue. The neighborhood H is a conceptual region in which the linear approximation (2.2) is deemed valid. Once fluid elements are within H , they will stretch exponentially, with a rate given by the positive eigenvalue,

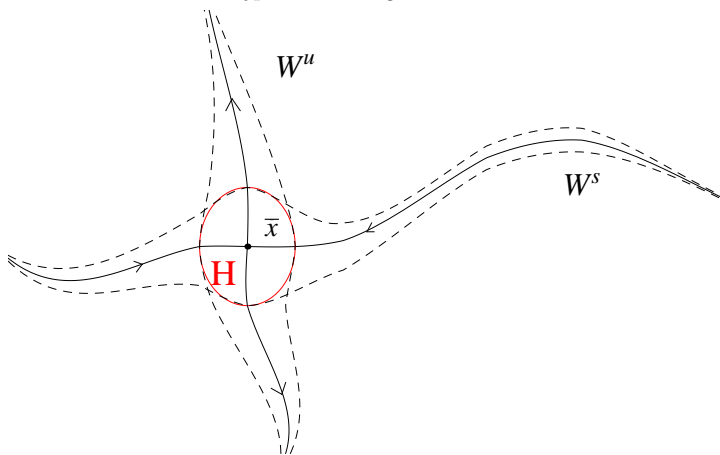


FIGURE 1. Hyperbolic neighborhood H (red) for a steady situation: fluid lying within H experiences exponential stretching, while fluid elements within the dashed envelopes around the stable W^s and unstable W^u manifolds of the stagnation point \bar{x} will stretch once inside H .

since they are pulled away by the unstable manifold. This will happen for fluid blobs that traverse the stable manifold as well, since except for the (one-dimensional) part of the blob lying exactly on the stable manifold, the parts of the fluid element on its two sides will get pulled apart exponentially due to the unstable manifold. Now, since the linear approximation only works while in H , there is no expectation of exponential stretching for instances when fluid elements are outside H . However, fluid lying within the dashed envelope surrounding W^s will eventually be within H , and *then* will exponentially stretch. Similarly, fluid in the dashed envelope around W^u will have experienced exponential stretching at some time in the past.

The region H also has an underappreciated connection to both FTLE ridges and LCSs. Suppose the forward-time FTLE was computed for the situation pictured in figure 1. If this was done for a *short* time period T , then parts of the stable manifold close to \bar{x} will emerge as large values of the FTLE field, since fluid elements there are within the immediate exponential stretching regime. However, if one computes for a larger time, then large FTLE values will also arise for parts of the stable manifold that are further away. In performing large T calculations, another important factor is that it is *only* trajectories that are *exactly* on the stable manifold that will reside within H forever after entering H . These will therefore continue to incur exponential stretching forever, in contrast to other trajectories that will leave H . Thus, if performing a forward-time FTLE calculation with large T , the stable manifold will emerge as having anomalously large values of the FTLE and will be visible as a ridge. The sensitivity of FTLE ridges to the time T over which the computation is performed is therefore directly related to whether that time would put fluid elements into H or not, since it is only inside H that exponential stretching occurs. Similarly, consider the definition of a repelling hyperbolic LCS of Haller (2011, 2015). Exponential repulsion from the stable manifold will occur only once particles are within H , and thus once again the method displays an intimate connection between the finite time over which the calculation is performed and whether this time would put certain particles into H in its duration or not. In other words, the sensitivity of both FTLEs and hyperbolic LCSs to the finite time considered is linked to H . We will explore the FTLE connection in more detail in §3.3.

More realistically, the global flow will *not* be linear, and will contain several hyperbolic stagnation points. In this generic situation, fluid traversing the flow domain would have

‘exponential boosts’ to its stretching each time it went through a neighborhood H of a hyperbolic point. These neighborhoods are therefore crucial to understanding the Lagrangian evolution of stretching.

2.2. Unsteady Flows: Hyperbolicity and Its Relationship to FTLEs/LCSs

When the flow is unsteady, we must generalize from hyperbolic fixed points (since in general there may no longer be any true fixed points) to hyperbolic trajectories. These objects are special solutions $\bar{\mathbf{x}}(t)$ of the Lagrangian equations of motion

$$\dot{x}_i = u_i(\mathbf{x}, t), \quad (2.3)$$

for which the variational equation

$$\dot{y}_i = y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t) \quad (2.4)$$

possesses an exponential dichotomy (Coppel 1978; Balasuriya 2012). Intuitively, this condition requires that there be one direction along which one can choose a material point a distance $\mathbf{y}(t)$ away from $\bar{\mathbf{x}}(t)$ such that $\mathbf{y}(t)$ decays exponentially quickly in forward time, and another direction along which $\mathbf{y}(t)$ decays exponentially in backward time. These conditions again imply the local existence of a stable and unstable manifold attached to the hyperbolic trajectory $\bar{\mathbf{x}}(t)$, although unlike in the steady case these manifolds evolve in time. The basic picture is something like that of figure 1, but the point $\bar{\mathbf{x}}$ is moving with time, as are the local directions of emanation of the stable/unstable manifolds and the manifolds themselves. Given these differences, it is clear that simply looking for instantaneous stagnation points need not capture hyperbolic stretching; nor do eigenvalues and eigenvectors provide information about the local behavior (Balasuriya 2016b).

In practice, *finite-time hyperbolic trajectories* tend to be determined in different ways depending on the particular method being used, and there is no universally agreed-upon definition for them. If using FTLEs, they would be intersections of the forward- and backward-time FTLE ridges (Voth *et al.* 2002). We however note that this can be an ambiguous exercise, since defining exactly what makes a feature a ‘ridge’ is unclear. For example, how sharp must it be? Moreover, there is no guarantee that if one tracks an intersection point of ridges with time, it will be a trajectory of the flow (Haller 2015). If using hyperbolic LCSs, they would be intersections of repelling and attracting LCSs. These characterizations mean that hyperbolic trajectories for finite-time flows may exist only for some times, and not others; although the trajectory of a given fluid element must be continuous in incompressible flow, it need not remain hyperbolic over its entire history. In particular, we can no longer say that a fluid element on the stable manifold will eventually experience exponential stretching; the time it would take for that fluid element to arrive near the hyperbolic trajectory may be longer than the hyperbolic trajectory exists. Fluid elements that are close to the hyperbolic trajectory will, however, experience the exponential stretching that is the signature of hyperbolicity. The dynamics of those fluid elements are therefore dominated by the hyperbolic trajectory, and determine its zone of influence in the fluid—which we shall refer to as its *hyperbolic neighborhood* and label by $H(t)$. This neighborhood is likely to contain segments of the analogs of stable and unstable manifolds of the hyperbolic trajectory, as fluid elements on the manifolds and near to the hyperbolic trajectory will certainly feel their effects; but it will also likely contain fluid elements *not* on the manifolds but still close to the hyperbolic trajectory, in the region of the flow field where the streamlines are close to hyperbolæ. Thus, the

hyperbolic neighborhood is distinct from a simple estimate of the manifolds, due to its finite-time nature.

If the flow is steady, the *directionality* of the stretching incurred in forward time while within a hyperbolic neighborhood is indicated by the direction of the unstable manifold in figure 1. In other words, this is the direction of the eigenspace spanned by the eigenvectors associated with eigenvalues with positive real part of the Jacobian matrix evaluated at the hyperbolic point $\bar{\mathbf{x}}$. In backward time, the stretching occurs in the direction of the eigenspace associated with eigenvalues with negative real part. This directionality becomes less easy to define for unsteady flows since, just as the hyperbolic trajectory $\bar{\mathbf{x}}$ is moving with time, so are the directions of emanation of the stable/unstable manifolds. Formally, Balasuriya (2016*b*) shows that these time-varying directions can be characterized in terms of Oseledets spaces or through exponential dichotomy projections, but these are difficult to use in most practical settings. Trajectories traversing $H(t)$ will experience exponential stretching instantaneously in these directions, but since they are changing with time, the accumulated stretching while within $H(t)$ cannot be encapsulated as being in a particular direction. Nevertheless, over a finite time interval, exponential separation between adjacent trajectories is to be expected while both are within $H(t)$.

Before providing a definition as to how to determine $H(t)$, we make the important observation that the variational equation (2.4) has a central role in determining either FTLEs or hyperbolic LCSs. If the finite time under consideration is from a time t_0 to a time $t_0 + T$, for example, let $\Phi(\mathbf{x}, t_0, T)$ be the flow map that gives the location at time $t_0 + T$ of the fluid element that was at \mathbf{x} at time t_0 . It can be shown (Haller 2011) that its gradient

$$Z_{ij}(T) = \frac{\partial \Phi_i}{\partial x_j}(\mathbf{x}, t_0, T) \quad (2.5)$$

is a fundamental matrix solution to the variational equation (2.4), in the sense that

$$\frac{d}{dT} Z_{ij}(T) = Z_{kj}(T) \frac{\partial}{\partial x_k} u_i(\mathbf{x}(T), T) \quad ; \quad \mathbf{x}(t_0) = \mathbf{x}. \quad (2.6)$$

The right Cauchy–Green strain tensor is formed from the contraction of the gradient of the flow map with itself, so that

$$C_{ij}(\mathbf{x}, t_0, T) = \frac{\partial \Phi_k}{\partial x_i} \frac{\partial \Phi_k}{\partial x_j}. \quad (2.7)$$

If the largest eigenvalue of the Cauchy–Green tensor is λ_{\max} , then the FTLE σ is defined by (Shadden *et al.* 2005)

$$\sigma(\mathbf{x}, t_0, T) = \frac{1}{T} \ln \left(\sqrt{\lambda_{\max}(\mathbf{x}, t_0, T)} \right). \quad (2.8)$$

We will in Section 3.1 examine the FTLE in some more detail within the context of computation. The definition of hyperbolic LCSs (Haller 2011, 2015) *also* depends specifically on eigenvalues and eigenvectors of the Cauchy–Green tensor, or equivalently, on the singular values and vectors of Z , the fundamental matrix solution of the variational equation. The crucial observation is that both FTLEs and hyperbolic LCSs therefore depend on a solution to the variational equation (2.4), *which is only valid if the higher-order nonlinear terms are discarded*. In other words, conclusions arising from these definitions are only reasonable *in a small enough neighborhood* traveling along with the trajectory that is being examined.

2.3. Hyperbolic Neighborhoods: Dominance of Linearity

How, then, does one determine a hyperbolic neighborhood? The key to our method is the requirement that it is precisely the region in space around the hyperbolic trajectory, where the linearized flow near the hyperbolic trajectory gives a good approximation to the actual flow. To see how this condition helps us to estimate $H(t)$, let us recall how equation (2.4) is obtained. First, we let $\mathbf{x}(t) = \bar{\mathbf{x}}(t) + \mathbf{y}(t)$, where $\mathbf{y}(t)$ is the deviation away from the hyperbolic trajectory $\bar{\mathbf{x}}(t)$. Inserting into equation (2.3), we obtain

$$\dot{\hat{x}}_i + \dot{y}_i = u_i(\bar{\mathbf{x}} + \mathbf{y}, t) = u_i(\bar{\mathbf{x}}, t) + y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}, t) + \mathcal{O}(|\mathbf{y}|^2), \quad (2.9)$$

where the second equality follows from Taylor-expanding the velocity field. Now, if we assume that the variation $\mathbf{y}(t)$ is very small at time t so that the second fluid element is very close to the hyperbolic trajectory, we can keep only terms up to $\mathcal{O}(|\mathbf{y}|)$, thereby recovering equation (2.4). This derivation, however, also allows us to bound the error introduced by keeping only the first-order terms. Since $\bar{\mathbf{x}}(t)$ is itself a Lagrangian trajectory, we can also write

$$\dot{y}_i = u_i(\bar{\mathbf{x}} + \mathbf{y}, t) - \dot{\hat{x}}_i = u_i(\bar{\mathbf{x}} + \mathbf{y}, t) - u_i(\bar{\mathbf{x}}, t), \quad (2.10)$$

where we have made no approximations. The error in the variational approximation is then the difference between this equation and the variational equation. Now, above we described the hyperbolic neighborhood $H(t)$ as being the region around $\bar{\mathbf{x}}(t)$ in which its hyperbolic character is preserved so that the flow stretches exponentially. We can now make this notion more precise by defining $H(t)$ to be the region around $\bar{\mathbf{x}}(t)$ in which the replacement of equation (2.10) by equation (2.4) is valid—in other words, the region in which the approximation error is small when compared with the retained terms. This condition can be quantified by

$$\left| y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t) \right| \gg \left| u_i(\bar{\mathbf{x}}(t) + \mathbf{y}, t) - u_i(\bar{\mathbf{x}}(t), t) - y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t) \right|. \quad (2.11)$$

This condition requires that the estimate of the velocity difference between points $\bar{\mathbf{x}} + \mathbf{y}$ and $\bar{\mathbf{x}}$ by the velocity gradient at $\bar{\mathbf{x}}$ alone is very good—or, in other words, that the flow is very close to the linear flow at the hyperbolic trajectory. As long as the velocity gradient at $\bar{\mathbf{x}}(t)$ has no purely imaginary eigenvalues, this condition will be achievable for some non-empty set.

DEFINITION 1 (HYPERBOLIC NEIGHBORHOOD (HN)). *The hyperbolic neighborhood $H(t)$ of a hyperbolic trajectory $\bar{\mathbf{x}}(t)$ is the set*

$$H(t) := \left\{ \bar{\mathbf{x}}(t) + \mathbf{y} : \left| y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t) \right| > E \left| u_i(\bar{\mathbf{x}}(t) + \mathbf{y}, t) - u_i(\bar{\mathbf{x}}(t), t) - y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t) \right| \right\}, \quad (2.12)$$

which depends on a constant $E \geq 1$.

We have introduced a constant $E \geq 1$ that explicitly allows us to specify how accurate we require the linear approximation to be. In general, as E is increased, the size of $H(t)$ will decrease (though there are exceptions, which we shall address in §3.3)—and thus for very large values of E , $H(t)$ may be unmeasurable from experiments or observations where the spatial resolution of the velocity field is always finite. We therefore expect that some tuning of E will likely be necessary in these cases so as to draw meaningful conclusions. We will show via analytical examples that making a choice of E is essentially

equivalent to making a choice of a cut-off exponential stretching rate, and that the results obtained by this process are not hugely sensitive to E .

We remark that our definition of hyperbolic neighborhoods (HNs) is quite different from the concept of ‘LCS cores’ introduced by Olascoaga & Haller (2012) that attempt to delimit the most important segments of the attracting/repelling hyperbolic LCSs in two dimensional flows. Having identified a hyperbolic trajectory and its attached attracting LCS curve, they follow the curve until the strain rate towards it reverses direction, and clip the attracting LCS at that point. In computing the strain rate, they use quantities derived from the *linearized* flow, essentially the variational equation at each point on the LCS. By doing this, they are able to clip an attracting and a repelling LCS near a hyperbolic trajectory, with the understanding that the flow topology nearby is like hyperbolæ. This process identifies two intersecting curve segments, and not a *region*. Moreover, the method of Olascoaga & Haller (2012) does not seek to check when a linear approximation is legitimate. Indeed, the process uses strictly linearized flow on the structures. In contrast, our hyperbolic neighborhoods directly investigate the concept of linearity, and thereby exponential stretching.

We have already mentioned that exponential behavior cannot be strictly verified in finite-time flows, since any function can be bounded by an exponential over a finite time interval. We have avoided this difficulty by directly searching for *linearity* in the flow field, which drives ‘instantaneous exponential stretching’ in an intuitive sense. While this latter concept cannot be defined clearly, the former can be quantified, and leads to our definition of hyperbolic neighborhoods.

In thinking in terms of hyperbolic neighborhoods, we have also highlighted some aspects of FTLEs/LCSs that are often hidden. These quantities are usually calculated at a time t_0 , with computations being done until a time $t_0 + T$, for some fixed T . Thus, T might be thought of as the time period over which the relevant structures are important according to whatever definition one is using. In test examples such as the double gyre (Shadden *et al.* 2005), hyperbolicity lasts over any time period chosen, and thus the choice of T (apart from not being too large or small) usually does not affect outputs significantly. The transitory nature of hyperbolic structures in realistic flows associated with finite-time data, however, makes the choice of T crucial. Referring back to figure 1, consider fluid elements chosen near W^s far to the right at time t_0 , but imagine that this is a finite-time, unsteady situation. If the time T is such that (i) the hyperbolicity of $\bar{\mathbf{x}}$ disappears during this time, (ii) these elements have not entered H by time $t_0 + T$, or (iii) the fluid elements passed quickly through H and are well away from H by time $t_0 + T$, then the influence of the hyperbolic neighborhood may be not noticed at all. These considerations can make the methodology being used highly sensitive to T . This sensitivity is embraced by our concept of hyperbolic neighborhoods, since they are able to identify, at each time t between t_0 and $t_0 + T$, the parts of space that experience hyperbolicity, thereby enabling a broader picture of how to instantaneously segment the flow domain into regions in which the flow behaves ‘differently.’ It must also be emphasized that our method works even for compressible flows with general time-dependence, and in any dimension, assuming the hyperbolic trajectories can be determined.

Our definition of hyperbolic neighborhoods is also objective (Truesdell & Noll 2004; Haller 2015; Peacock *et al.* 2015), in that they are invariant under affine coordinate transformations of the form

$$z_i = Q_{ij}(t)x_j + b_i(t), \quad (2.13)$$

where $\mathbf{b}(t)$ is an arbitrary vector and $Q_{ij}(t)$ is an orthogonal tensor. The invariance occurs since the t in (2.12) appears only as a parameter (there are no t -derivatives, for

example). The hyperbolic trajectory location would be at $\bar{z}_i(t) = Q_{ij}(t)\bar{x}_j(t) + b_i(t)$ in the transformed frame, and if \mathbf{w} is the deviation from this, we have $w_i = Q_{ij}(t)y_j$ as the relationship between the local deviations in the two reference frames. Thus, if $\mathbf{U}(\mathbf{z}, t)$ is the velocity field as represented in the \mathbf{z} -frame, the linear term in the \mathbf{z} -frame takes the form

$$w_j \frac{\partial}{\partial z_j} U_i(\bar{\mathbf{z}}(t), t) = (Q_{jk}(t)y_k) \left(\frac{\partial}{\partial x_\ell} u_i(\bar{\mathbf{x}}(t), t) Q_{\ell j}^{-1}(t) \right) = y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t), \quad (2.14)$$

which is invariant. The other terms appearing in (2.12) involve the difference between the velocity at two points, which is also invariant under the transformation, and thus—assuming that an objective criterion has been used in identifying the hyperbolic trajectories—hyperbolic neighborhoods are objective.

We note that we have outlined here a method for determining the hyperbolic neighborhood associated with a particular hyperbolic trajectory—but not a method for finding the hyperbolic trajectory itself. In general, finding such objects can be difficult, but there are also analytical (Balasuriya 2011) and numerical methods (Dellnitz & Junge 2002; Balasuriya & Padberg-Gehle 2014) as alternatives to the curve intersection ideas associated with FTLEs and LCSs. A tractable method must be chosen for general experimental or observational data, as we outline in §4.2.

3. Analytical Examples

To build intuition for what these hyperbolic neighborhoods mean, we consider here several analytical examples where the hyperbolic neighborhoods can be explicitly determined. We will demonstrate the fact that trajectories entering an HN display boosts to their instantaneous exponential stretching rate, which diminishes as they leave the HN. The residence time within HNs will be established to be correlated to the long-term stretching rates as identifiable via FTLEs. The connection between the factor E in the definition of HNs and a ‘threshold stretching rate’ will be qualitatively and quantitatively established through these examples. Furthermore, we will show why some instances in which FTLEs provide unexpected results (Haller 2011) can be attributed directly to the geometry of the HNs.

3.1. Steady Taylor–Green Flow

The well-known steady Taylor–Green flow is defined by the streamfunction

$$\psi(\mathbf{x}) = \cos x_1 \cos x_2, \quad (3.1)$$

which gives the velocity field as

$$\begin{aligned} \dot{x}_1 &= -\cos x_1 \sin x_2 \\ \dot{x}_2 &= \sin x_1 \cos x_2. \end{aligned} \quad (3.2)$$

Since this flow is periodic in both directions, it has an infinite number of critical points (Ouellette & Gollub 2007). Let us focus on the critical point $\bar{\mathbf{x}} = (\pi/2, \pi/2)$, which is hyperbolic; and note that, as this is a steady flow, this point is also a hyperbolic trajectory. For this flow, the velocity gradient is given by

$$\nabla \mathbf{u} = \begin{pmatrix} \sin x_1 \sin x_2 & -\cos x_1 \cos x_2 \\ \cos x_1 \cos x_2 & -\sin x_1 \sin x_2 \end{pmatrix}, \quad (3.3)$$

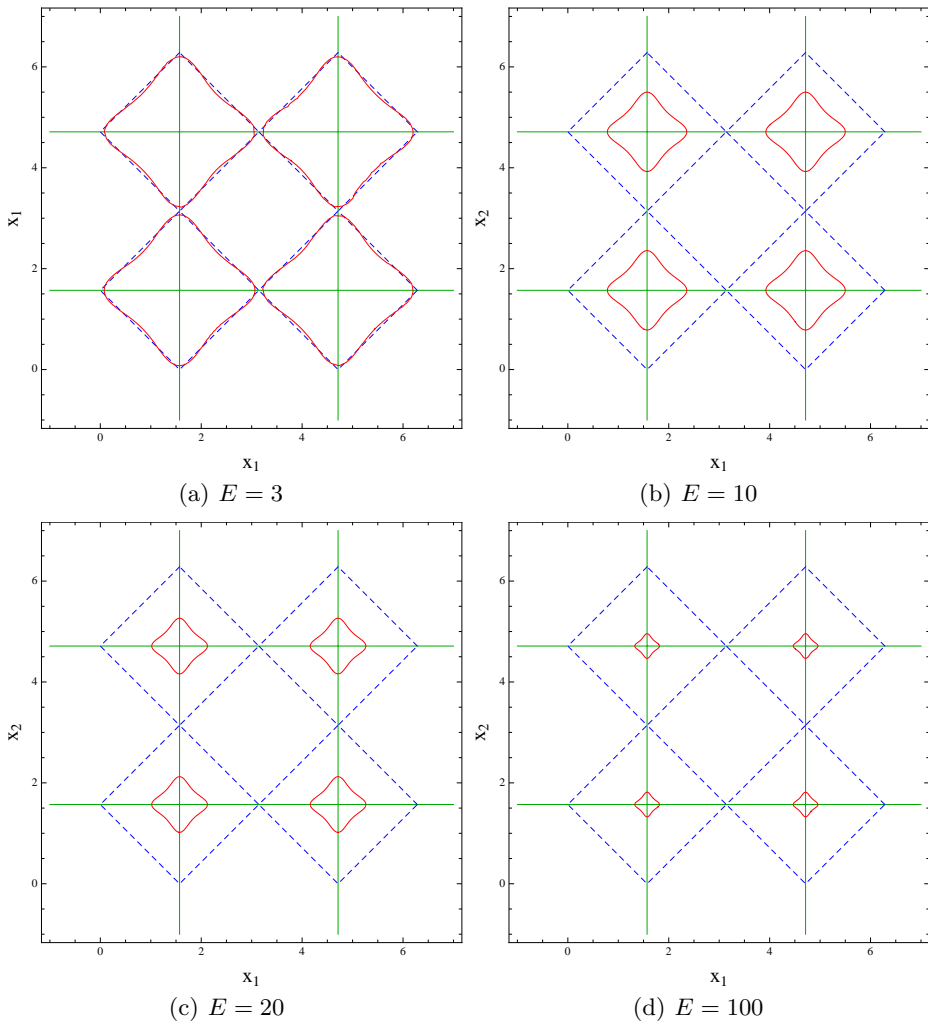


FIGURE 2. Hyperbolic neighborhoods (red solid lines) for the steady Taylor–Green flow, as defined by equation (3.6), for several values of E . The stable/unstable manifold connections between the hyperbolic points are shown with green horizontal and vertical lines. For comparison, the blue dashed lines show the regions that would be identified as hyperbolic based on the instantaneous Okubo–Weiss criterion (equation (3.8)).

and so we have

$$y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}, t) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ -y_2 \end{pmatrix}. \quad (3.4)$$

Therefore,

$$\begin{aligned} \mathbf{u}(\bar{\mathbf{x}} + \mathbf{y}, t) - \mathbf{u}(\bar{\mathbf{x}}, t) - \nabla \mathbf{u}(\bar{\mathbf{x}}, t) \cdot \mathbf{y} &= \begin{pmatrix} -\cos\left(\frac{\pi}{2} + y_1\right) \sin\left(\frac{\pi}{2} + y_2\right) \\ \sin\left(\frac{\pi}{2} + y_1\right) \cos\left(\frac{\pi}{2} + y_2\right) \end{pmatrix} - \begin{pmatrix} y_1 \\ -y_2 \end{pmatrix} \\ &= \begin{pmatrix} \sin y_1 \cos y_2 - y_1 \\ -\cos y_1 \sin y_2 + y_2 \end{pmatrix}. \end{aligned} \quad (3.5)$$

The hyperbolic neighborhood associated with a hyperbolic point $\bar{\mathbf{x}} = (\bar{x}_1, \bar{x}_2)$ is

therefore

$$H = \left\{ (\bar{x}_1 + y_1, \bar{x}_2 + y_2) : y_1^2 + y_2^2 > E^2 \left[(\sin y_1 \cos y_2 - y_1)^2 + (-\cos y_1 \sin y_2 + y_2)^2 \right] \right\}, \quad (3.6)$$

where for convenience we have squared both sides of the inequality in equation (2.12). This neighborhood is shown in figure 2, along with the hyperbolic neighborhoods associated with the hyperbolic points at $(\pi/2, 5\pi/2)$, $(5\pi/2, \pi/2)$, and $(5\pi/2, 5\pi/2)$. For comparison, we also show the regions that would be identified as hyperbolic based on the Okubo–Weiss criterion (Okubo 1970; Weiss 1991); that is, regions that satisfy

$$\det \nabla \mathbf{u} < 0, \quad (3.7)$$

or

$$-\sin^2 x_1 \sin^2 x_2 + \cos^2 x_1 \cos^2 x_2 < 0 \quad (3.8)$$

in this case. For this flow, the HNs with $E = 3$ are quite similar (although not identical) to those regions predicted to be hyperbolic based on the Okubo–Weiss criterion. This similarity is reasonable, since for steady flows the Okubo–Weiss criterion is fairly effective at indicating hyperbolic regions; objections to it are primarily concerned with the relative time scales over which the velocity gradient and the vorticity gradient evolve (Rivera *et al.* 2014). As E is increased, the HNs become smaller. We can think of E as being our measure for how closely we require linearity.

An interesting observation from the nature of the HNs in figure 2 is that they exhibit a weak elongation in the direction of the stable and unstable manifolds; they ‘stick out’ slightly in those directions. The intuition for this observation is that the exponential stretching is *exact* in the direction of the stable/unstable manifolds local to the hyperbolic point, and identifiable with the eigenvalue of the linearized velocity gradient matrix. The direction of the manifold is in fact the corresponding eigenvector direction. Slightly off a stable/unstable manifold, while the behavior is certainly dominated by the eigenvalue associated with this nearby manifold, there is also a contribution to the stretching arising from the other eigenvalue, which has the opposite sign. This, then, diminishes the stretching in comparison to being exactly on the manifold. Thus, the region of dominance of linearity tends to push out along the stable and unstable manifold directions in this case. As we shall see later, there are instances in which this generic behavior does *not* occur, in what we shall identify to be *degenerate* HNs.

Next, we verify our claim that the exponential stretching rate within HNs is anomalously larger than outside HNs. To do so, consider a pair of points α^1 and β^1 at an initial time t_0 , located close to $(\pi/2, \pi)$, i.e., near the stable manifold of $(\pi/2, \pi/2)$. Let $\Phi_\alpha^1(t)$ and $\Phi_\beta^1(t)$ be trajectories of (3.2), as a function of time, that satisfy $\Phi_\alpha^1(t_0) = \alpha^1$ and $\Phi_\beta^1(t_0) = \beta^1$. In this vein, imagine *many* pairs of trajectories $\Phi_\alpha^i(t)$ and $\Phi_\beta^i(t)$ that satisfy $\Phi_\alpha^i(t_0) = \alpha^i$ and $\Phi_\beta^i(t_0) = \beta^i$. If $\lambda^i(t)$ is the instantaneous measure of the exponential rate of stretching of the distance between the trajectories $\Phi_\alpha^i(t)$ and $\Phi_\beta^i(t)$ at time t , we can think of approximating it via

$$e^{\lambda^i(t)\Delta t} = \frac{\left| \Phi_\alpha^i(t + \Delta t) - \Phi_\beta^i(t + \Delta t) \right|}{\left| \Phi_\alpha^i(t) - \Phi_\beta^i(t) \right|} \quad (3.9)$$

over a short time interval $[t, t + \Delta t]$. Thus, we can computationally recover the instantana-

neous exponential stretching rate for a pair of trajectories starting at α^i and β^i as

$$\lambda^i(t) = \frac{1}{\Delta t} \ln \frac{|\Phi_\alpha^i(t + \Delta t) - \Phi_\beta^i(t + \Delta t)|}{|\Phi_\alpha^i(t) - \Phi_\beta^i(t)|}. \quad (3.10)$$

We will momentarily use this measure—a pairwise instantaneous stretching rate—to compute stretching rates for the Taylor–Green flow. First, though, we establish a connection between it and FTLEs. Imagine that we start at time t_0 with a pair of points α^i and β^i , and track the trajectories Φ_α^i and Φ_β^i numerically until time $t_0 + T$, thereby obtaining information at discrete times $t_0 + j\Delta t$, $j = 0, 1, 2, \dots, n$, where $n\Delta t = T$. Thus, the average *accumulated* instantaneous stretching over the full finite time period is

$$\begin{aligned} \Lambda^i(t_0, t_0 + T) &:= \frac{1}{n} \sum_{j=0}^{n-1} \frac{1}{\Delta t} \ln \frac{|\Phi_\alpha^i(t_j + \Delta t) - \Phi_\beta^i(t_j + \Delta t)|}{|\Phi_\alpha^i(t_j) - \Phi_\beta^i(t_j)|} \\ &= \frac{1}{n\Delta t} \ln \frac{|\Phi_\alpha^i(t_0 + T) - \Phi_\beta^i(t_0 + T)|}{|\Phi_\alpha^i(t_0) - \Phi_\beta^i(t_0)|} \\ &= \frac{1}{T} \ln \frac{|\Phi_\alpha^i(t_0 + T) - \Phi_\beta^i(t_0 + T)|}{|\Phi_\alpha^i(t_0) - \Phi_\beta^i(t_0)|}. \end{aligned} \quad (3.11)$$

If thinking of α^i as a fixed initial location at time t_0 , the forward FTLE field at α^i , given in terms of the Cauchy–Green tensor by (2.8), has the equivalent form (Shadden *et al.* 2005; Haller 2015; Balasuriya 2012)

$$\sigma(\alpha^i, t_0, T) = \frac{1}{T} \ln \sup_{\beta^i} \frac{|\Phi_\alpha^i(t_0 + T) - \Phi_\beta^i(t_0 + T)|}{|\Phi_\alpha^i(t_0) - \Phi_\beta^i(t_0)|}, \quad (3.12)$$

where the supremum over β^i indicates that one takes initial locations β^i in all possible directions from α^i , but then chooses the direction such that the ratio of expansion above is maximized. Indeed, Samelson (2013) argues that the FTLE defined in (3.12) might be more naturally interpreted in the fluid dynamics context as a Finite-Time Lagrangian Strain; to avoid confusion as to exactly what ‘strain’ may mean, we offer the additional modification that perhaps it should be thought of as a Finite-Time Lagrangian Stretching. In any case, the connection between the average accumulated stretching (3.11) and the FTLE (3.12) is clear: these are proportional to one another if β^i is chosen infinitesimally closely in the direction in which the FTLE at α^i is maximized. Therefore, the instantaneous stretching rate definition (3.10) we will use for numerical evaluation is strongly connected to the FTLE; loosely speaking, its accumulated effects lead to the FTLE over a finite-time interval. (Though the FTLE does have the added ‘supremum over all directions’ as part of its definition, whereas (3.10) is a pairwise stretching measure.)

With these connections understood, we can now return to the Taylor–Green flow and evaluate how the stretching between trajectories is impacted by being within HNs. We seed a collection of fluid particles (blue xs) as shown in the top panels of figure 3, which will pass into the HN (shown in red) during our time of integration. For each adjacent pair of initial particle locations, we show in the bottom panels how the instantaneous stretching rate (3.10) varies. Each color corresponds to a pair of trajectories. We deem that a pair of trajectories is within the HN centered at $(\pi/2, \pi/2)$ if their midpoint is




FIGURE 3. The flow of trajectories in comparison with the HN (top panels) with the initial seedings shown by x s and the final location by o s, and the computed instantaneous stretching rate (3.10) for the Taylor–Green flow (bottom panels) at different E values, with the range over which a particular colored trajectory is within the HN indicated by the lines at the top.

within the HN at that time. The range of time over which a particular pair of trajectories is within the HN is indicated by the line of identical color towards the top of the figures. All trajectories initially have small values of λ , since they are outside the HN. As they approach the HN, the value of λ increases appreciably, thereby demonstrating that strong exponential stretching is occurring. If $E = 5$, the HN is chosen to be fairly large, and if $E = 10$, it is smaller (and shown in figure 2(b)). When E is larger, a smaller number of trajectory pairs are indicated as being within an HN, and moreover are within an HN for a shorter time than when E is smaller. Thus, figure 3 validates very clearly our claim that trajectories entering HNs experience a boost in their exponential stretching, and lose it as they leave.

The time of residence—as indicated by the lengths of the lines at the top of figure 3—is clearly linked to when λ is large in the corresponding curve below. In other words, if the stretching curves were cut off below a certain threshold, it would appear that the time-ranges at the top of the figure would be obtained. We investigate this quantitatively in figure 4. For a given E value, we determine the stretching threshold λ_o such that if keeping only time values for which the stretching rates are above λ_o in figure 3, we recover ‘close to’ the time range within the HN as reflected by the time-range lines at the top of figure 3. We define this closeness as follows. Let $\text{HN}^i(t_j)$ be the indicator functions as to whether trajectory pair i is within the HN at time t_j , i.e., $\text{HN}^i(t_j) = 1$ if the midpoint




FIGURE 4. (a) The stretching rates of figure 3(b), but with the range over which λ is larger than a threshold value λ_o indicated by the lines at the top of the figure. (b) λ_o values as functions of E plotted on logarithmic axes for the Taylor–Green flow, chosen by minimizing (3.13), with identical parameter values as for figure 3 being used.

is within HN at time t_j and 0 if not. Let $I^i(t_j)$ be an indicator function as to whether $\lambda(t_j) > \lambda_o$ at time t_j . Then, we define the discrepancy between these by

$$\epsilon(\lambda_o) = \frac{\Delta t}{mT} \sum_{j=0}^{n-1} \sum_{i=1}^m |\text{HN}^i(t_j) - I^i(t_j)|. \quad (3.13)$$

where there are m trajectory pairs considered. Note that this is essentially an L^1 -norm, suitably averaged over time and number of pairs of trajectories. By sweeping over choices of λ_o , we determine the value of λ_o that minimizes $\epsilon(\lambda_o)$. In figure 4(a), we show the computed range functions $I^i(t)$ at the top of the figure for $E = 10$, for a direct comparison with figure 3(b) (in which the range functions at the top of the figure correspond to $\text{HN}^i(t)$). Thus, choosing an HN with $E = 10$ is very similar to choosing a threshold stretching rate of $\lambda_o = 0.813$, which is shown by the dashed blue line in figure 4(a). In other words, this is the choice which makes the ‘being within HN’ and ‘being above the stretching threshold’ conditions similar. In this way, we can find the appropriate threshold stretching λ_o corresponding to different E values, and present this information in figure 4(b). As E gets larger the variation of λ_o slows; this corresponds to the statement that as one approaches the hyperbolic trajectory arbitrarily closely, the stretching rate is effectively the stretching rate of the hyperbolic trajectory. Theoretically we would expect this value to approach 1, which is the unstable eigenvalue associated with the stagnation point $(\pi/2, \pi/2)$. Obtaining this numerically in the limit $E \rightarrow \infty$ requires continually refining the spatial resolution, since as $E \rightarrow \infty$, the size of the HN itself approaches zero, but must remain resolved for these computations to be performed. Figure 4(b) therefore enables a *quantitative* evaluation of how the factor E in the definition of the HN is related to how the exponential stretching within the HN is close to that of its core hyperbolic trajectory.

3.2. Duffing Equation

As a second example, consider the steady, unforced Duffing equation (Wiggins 1992; Balasuriya 2011)

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= x_1 - x_1^3 \end{aligned} \quad (3.14)$$




FIGURE 5. Hyperbolic neighborhoods (boundaries shown with dashed curves) for the Duffing equation, as computed from equation (3.18), for $E = 8$ (red), 20 (blue), and 40 (green). The stable and unstable manifolds of the origin are shown by the thick black curves, and the thin curves show how two different trajectories, starting at locations α and β , traverse the HNs.

has a hyperbolic fixed point at the origin, to which is attached one-dimensional stable and unstable manifolds (see figure 5). Since this is again a steady (autonomous) system, the corresponding hyperbolic trajectory is simply $(0, 0, t)$ for all t ; moreover, the hyperbolic neighborhood $H(t)$ will also be the same for all t . The velocity gradient is easily calculated to be

$$\nabla \mathbf{u} = \begin{pmatrix} 0 & 1 \\ 1 - 3x_1^2 & 0 \end{pmatrix}, \quad (3.15)$$

from which we obtain

$$y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}, t) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} y_2 \\ y_1 \end{pmatrix}. \quad (3.16)$$

This system also has $\mathbf{u}(\bar{\mathbf{x}}(t), t) = 0$ for all t , and so

$$\mathbf{u}(\bar{\mathbf{x}} + \mathbf{y}, t) - \mathbf{u}(\bar{\mathbf{x}}, t) - \nabla \mathbf{u}(\bar{\mathbf{x}}, t) \cdot \mathbf{y} = \mathbf{u}(\mathbf{y}, t) - \begin{pmatrix} y_2 \\ y_1 \end{pmatrix} = \begin{pmatrix} 0 \\ -y_1^3 \end{pmatrix}. \quad (3.17)$$

Using the definition in equation (2.12), $H(t)$ will therefore be the set of points where

$$H = \left\{ (\bar{x}_1 + y_1, \bar{x}_2 + y_2) : \sqrt{y_1^2 + y_2^2} > E |y_1^3| \right\} = \left\{ (x_1, x_2) : \sqrt{x_1^2 + x_2^2} > E |x_1^3| \right\}, \quad (3.18)$$

since $(\bar{x}_1, \bar{x}_2) = (0, 0)$. The HNs that satisfy this inequality are shown by dashed curves in figure 5. The neighborhoods have a waist-like structure with a pinch along the x_1 -axis. They are constrained along the x_1 direction, but unbounded in the x_2 direction. This behavior can be explained by examining the equations of motion (equation (3.14)): these equations are *always* linear in x_2 , and thus the linear approximation is always good in that direction. On the other hand, as x_1 becomes larger and larger, the linear approximation to the term $x_1 - x_1^3$ clearly becomes poor. Physically, the shape of the hyperbolic neighborhood can be understood by noting that this dynamical system also contains two elliptic fixed points at $(-1, 0)$ and $(1, 0)$. These elliptic points effectively shield fluid elements that are far away from the origin in the x_1 direction from the influence of the hyperbolic point. No such structures exist along the x_2 -axis, however, and so the influence of the hyperbolic point is felt everywhere in that direction.

We first establish as before that an increase in stretching occurs when trajectories enter




FIGURE 6. The instantaneous stretching rate (3.10) for the Duffing flow at different E values, with the range over which a particular colored trajectory is within the HN indicated by the lines at the top.

the HN. We choose many initial trajectory pairs near $(-1, 1)$, fairly close to the stable manifold. The evolution of the instantaneous stretching rate (3.10) is shown in figure 6. The upper lines indicate whether the trajectory pair identified by a particular color lies within the HN at the origin, for each of the choices (a) $E = 5$ and (b) $E = 100$. As for the Taylor–Green flow, there is a clear indication of the stretching rate increasing when close to the hyperbolic point. Using E as a cut-off for identifying when a trajectory pair is within the HN (as used in obtaining the time ranges at the top of the figures) is once again essentially equivalent to choosing a cut-off value of the instantaneous stretching rate λ . Notice that we have chosen E values that are highly disparate in this instance, and yet the results are not too different.

We next examine the impact of the nature of the HNs on stretching rates. The HN here is ‘partially degenerate’ in that it does not consist of a well-defined finite neighborhood of the hyperbolic point as shown in figure 1, but rather is unbounded in the x_2 -direction. (We will define a ‘degenerate HN’ in the next section.) The asymmetry of this situation has a strong impact on the stretching rates. Consider the two initial locations α and β , chosen symmetrically on the two sides of the stable manifold as shown in figure 5. If taking $E = 40$, the forward trajectory Φ_α through α goes through a significant portion of the HN, but the trajectory Φ_β through β spends much less time within the HN. Thus, we would expect a pair of trajectories that are close to Φ_α to have larger stretching rates than a pair close to Φ_β , a fact that would not be obvious without the notion of HNs. We investigate this further in figure 7. The left panels display the residence times within the HN, where we have chosen $E = 20$, while the right panels are the forward FTLE field. We have examined the flow from the initial time $t = 0$ to times $t = 1$ (top), $t = 3$ (middle) and $t = 6$ (bottom). Firstly, we note the strong correlation between the residence times and the forward FTLE fields, providing evidence that the accumulated stretching (as captured by the FTLE field) has an intimate link to how much time was spent in the HN. The longer the time spent within the HN, the longer the impact of the exponential stretching.

We note that there is also a fundamental difference in computing the residence time and the FTLE. For the residence time, we choose a point, iterate it in forward time, and catalog the time spent within the HN. For an FTLE calculation, on the other hand, behaviour of *nearby* points is also necessary to compute the FTLE value at one point. The residence time measure is therefore an intrinsic property of each trajectory without




FIGURE 7. The residence times in the HN with $E = 20$ (left), compared with the FTLE field (right) for the Duffing equation computed over different times from $t = 0$ to $t = 1$ (top), to $t = 3$ (middle) and to $t = 6$ (bottom).

regard to nearby motion; the strong correlation to nearby deformations as illustrated in figure 7 is particularly interesting. Of course, we cannot claim that the residence time is better or easier to use, since it requires knowledge of the location of the hyperbolic trajectory and its accompanying neighborhood.

As the final time value is increased, both the residence time and the FTLE plots display a sharpening ridge that, of course, approaches the stable manifold shown in figure 5. For smaller times, this ridge is less well-defined, and is moreover tilted away from the true stable manifold (for which $x_2 = x_1 \sqrt{1 - x_1^2}/2$ is the true expression in the vicinity of the origin, and which is well-approximated by the ridges in the lower panels of figure 7). The




FIGURE 8. Hyperbolic neighborhoods (boundaries shown with dashed curves) for the example of Haller (2011) as given in (3.19), as computed from equation (3.20), for $E = 6$ (red), 20 (blue), and 200 (green). The stable and unstable manifolds of the origin are shown by the thick black curves.

reason for this tilting is clear from the asymmetry of the HNs as shown in figure 5. The HNs are symmetric about the x_2 -axis, and trajectories closer to this axis spend longer times within the HNs. Thus, they accumulate more stretching; i.e., Φ_α from figure 5 will have both larger residence times and larger stretching than Φ_β . Therefore, both the residence time and the FTLE plots will be tilted towards the x_2 -axis for short times, as indicated by the numerics in figure 7. Moreover, we would expect the FTLE field at α to be larger than at β for this reason, and this is clearly visible in figure 7(f). There is a clear patch of very low FTLE values appearing to the lower-left of FTLE ridge near $(-0.4, 0.4)$, in contrast with the upper-right region.

The top panel in figure 7 is particularly instructive in interpreting FTLE fields. If figure 7(b) were considered, we might imagine that the stable manifold is tilted closer to the x_2 -axis than it really is. In this case we know exactly why this is so; it is because the HN—which is of infinite extent—is tilted thus, and trajectories closer to the x_2 axis will experience exponential stretching more. This is true for short enough times; once trajectories leave the HN, there is no longer any contribution to exponential stretching. Thus, the nature of the HNs provides additional insight into how the finite time value used in FTLE calculations affects the final results.

3.3. Degenerate HNs

The HNs of the Duffing system did not look like the generic bounded neighborhood shown in figure 1, in that they were unbounded. This has an influence on the FTLEs, which are a particular realization of the accumulated effects of the stretching rate (3.10). In some cases in which FTLEs give ‘unexpected’ results, we can show that this is *because* the HNs do not look like figure 1. We examine two examples in this section, in which the HNs are *degenerate* (in a way that we will define), and identify how the degenerate nature of the HNs affect the FTLEs in particular ways. First, consider the example due to Haller (2011), given by

$$\begin{aligned}\dot{x}_1 &= x_1 \\ \dot{x}_2 &= -x_2 - x_2^3.\end{aligned}\tag{3.19}$$

This system has a hyperbolic stagnation point at the origin, with its stable manifold along the x_2 -axis and its unstable manifold along the x_1 -axis. Thus, the initial expectation

might be that the forward-time FTLE field would reveal a ridge along the x_2 -axis and that the backward-time FTLE field would identify a ridge along the x_1 -axis. However, Haller (2011) shows, both numerically and analytically, that the forward-time FTLE field is *constant* everywhere (with no ridges), and the backward-time FTLE field has a *trough* (rather than a ridge) along the x_1 -axis. This example is offered by Haller (2011) to show why FTLEs can fail in identifying prominent Lagrangian structures, and here we are able to use the concept of HNs to provide insight as to *why* FTLEs do not work.

The calculations for the HN are very similar to that for the Duffing system, and we get

$$H == \left\{ (x_1, x_2) : \sqrt{x_1^2 + x_2^2} > E |x_2^3| \right\}. \quad (3.20)$$

These HNs are simply a horizontal/vertical flip of the Duffing ones, and are shown in figure 8. Our first observation is that these extend infinitely in the x_1 -direction, thereby displaying a feature similar to figure 5. However, there is an additional characteristic: whichever E we choose, the HN *always* contains the *full* unstable manifold. Put another way, unlike for the Taylor-Green flow (figure 2) or the Duffing equation (figure 5), each HN boundary *fails to cross all manifold branches emanating from the hyperbolic trajectory*. This shall be our definition for an HN to be *degenerate*.

Consider trajectories coming in from a positive x_2 value in figure 8. Whether coming in along the stable manifold *or not*, once they enter an HN, they will remain in it forever, because the HNs form horizontal bands around the unstable manifold, extending infinitely in x_1 . Thus, the trajectories will all accumulate the stretching associated with being inside an HN in the same way. The end result will be that the forward-time FTLEs will *not* be able to distinguish the stable manifold as a ridge. In the nondegenerate HN situation of figure 1, in contrast, it is *only* trajectories coming in along the stable manifold that will remain forever in the HN. The fact that all other trajectories eventually exit means that when computing forward-time FTLEs, they will eventually stop accumulating stretching, thereby resulting in a smaller FTLE value than for trajectories along the stable manifold. This explains why the forward-time FTLE field is constant and fails to identify the stable manifold as a ridge.

The backward-time FTLE failed in identifying the unstable manifold in a rather more spectacular way as shown by Haller (2011): the x_1 -axis emerges as a *trough* of the backward-time FTLE field rather than a *ridge*. That is, it has a *smaller* accumulation of stretching than nearby points, which appears contrary to the intuition emerging from the generic picture of figure 1. This failure is again because of the degenerate nature of the HN, but the reason is more subtle than for the forward-time FTLE. The standard expectation from a *generic* HN as shown in figure 1 would be that in the limit $E \rightarrow \infty$, the HN would shrink towards the hyperbolic point in a smooth way, retaining connections to the stable and unstable manifolds. This is indeed illustrated for the Taylor-Green flow in figure 2, and the generic picture is that the HN is pushed outwards along the stable and unstable manifold directions. However, the situation of figure 8 is different, because of the ‘waistlike’ structure of the HNs. These pinch *in* exactly along the stable manifold as E gets larger. In this degenerate situation the part of the stable manifold remaining within the HN as $E \rightarrow \infty$ is dramatically less than for nearby trajectories because of the inward pinching. (Indeed, it is easy to see that only a length of $2/\sqrt{E}$ of the stable manifold is within an HN with general E value.) The pinching means that trajectories coming in very close to the x_1 -axis encounter the HN boundary having experienced far less of an excursion in the x_2 -direction than trajectories coming in slightly further off the x_1 -axis. It is these x_2 -excursions that contribute to stretching in backward-time (since

the stable manifold is in the x_2 -direction), and therefore trajectories along the x_1 -axis incur *less* stretching than adjacent ones. Thus, a trough rather than a ridge appears in the backward-time FTLE field.

The subtlety of this situation is further evidenced by the fact that the backward-FTLE trough is a *shallow* one, as shown by Haller (2011). In contrast, ridges identified in a nondegenerate situation like that shown in figure 1 must be *sharp*, since trajectories on the stable manifold remain in the HN for infinite times, in contrast with a nearby trajectory, however close. The accumulated stretching will therefore be substantially higher. To see this, we would need to ensure that the time interval over which FTLEs are computed is tuned properly, and sufficient resolution is available. Exactly how sharp is ‘sharp’ is, of course, difficult to define. In computational experiments where sharpening is observed as the ridge gets refined, as in set of figures 7(b), (d) and (f), we can have confidence that this sharp ridge has an abruptly larger rate of exponential stretching than areas very near to it.

A second degenerate HN situation deserves mention. Consider the flow

$$\dot{x}_i = B_{ij}x_j \quad (3.21)$$

in which the $n \times n$ constant coefficient matrix B has one positive and one negative eigenvalue. For *any* trajectory of (3.21), the variational equation is then identical to (3.21) itself. In other words, a linearization is *exactly* correct, since there are no nonlinear terms. This reveals that all trajectories of (3.21) are hyperbolic trajectories. Moreover, the HN associated with any such hyperbolic trajectory is \mathbb{R}^n in its entirety, since linearity is exactly correct everywhere. The degeneracy of this HN is because no HN boundary can be defined for an HN occupying *all* space. Thus, the concepts of entering and leaving HNs and of incurring exponential stretching while only being inside an HN are irrelevant. FTLE fields will therefore have no ridges (Haller 2011; Balasuriya 2016a).

Through the above degenerate examples and also the Duffing equation, we have shown that understanding the geometrical structure of HNs is crucial to whether the usage of FTLEs provides the intuitive information that one might expect from examining figure 1. In the Duffing equation, the HNs are only ‘partially degenerate’ in that while being unbounded, the HN boundaries *did* intersect all branches of the stable and unstable manifolds emanating from the hyperbolic point. In this case, subject to standard issues regarding the integration time and spatial resolution, FTLEs will manage to identify stable and unstable manifold structures, as has been done for the Duffing equation (Wiggins 1992; Balasuriya 2016a). If the HNs are genuinely degenerate, as has been argued above, FTLEs display behavior that is at odds with the understanding provided by figure 1. Thus, HNs and their degeneracy is an additional tool for testing the capability of FTLEs (Allshouse & Peacock 2015, cf.) for identifying finite-time analogs of stable and unstable manifolds.

3.4. Oscillating Taylor–Green Flow

Our next analytical example is of an unsteady, chaotic flow, initially introduced for studying transport in Rayleigh–Bénard convection by Solomon & Gollub (1988), and subsequently popular in chaotic advection studies (Babiano *et al.* 2000; Nishikawa *et al.* 2001; Torney & Neufeld 2007; Khurana *et al.* 2011; Khurana & Ouellette 2012). It is obtained by adding a simple lateral oscillation along one of the coordinate axes to the Taylor–Green flow, so that the streamfunction is given by

$$\psi(\mathbf{x}, t) = \sin(x_1 + A \sin \Omega t) \sin x_2 \quad (3.22)$$




FIGURE 9. Hyperbolic neighborhoods for the oscillating Taylor–Green flow, for parameters $\Omega = 2\pi$, and $A = 0.3$, with the locations of the corresponding hyperbolic trajectory shown by the interior dots. Neighborhoods are calculated (a) using a fixed $E = 30$ at different times ($t = 0$ (red), $t = 0.15$ (blue), and $t = 0.8$ (green)); and (b) at a fixed $t = 0$ but with different values of E ($E = 10$ (red), $E = 50$ (blue) and $E = 200$ (green)).

and the velocity field is

$$\begin{aligned}\dot{x}_1 &= \sin(x_1 + A \sin \Omega t) \cos x_2 \\ \dot{x}_2 &= -\cos(x_1 + A \sin \Omega t) \sin x_2.\end{aligned}\quad (3.23)$$

When $A = 0$, this oscillating flow reduces to the steady Taylor–Green case, although, for convenience here, shifted such that the origin is a hyperbolic point. This has an unstable manifold in the x_1 -direction and a stable manifold in the x_2 -direction. For small A , there remains a hyperbolic trajectory near the origin, although now it moves with time. Its location can be determined using Theorem 2.10 of Balasuriya (2011) to be

$$\bar{\mathbf{x}}(t) = \begin{pmatrix} \frac{-A[\Omega \cos \Omega t + \sin \Omega t]}{1 + \Omega^2} \\ 0 \end{pmatrix} + \mathcal{O}(A^2) = \begin{pmatrix} \bar{x}_1(t) \\ 0 \end{pmatrix} + \mathcal{O}(A^2).\quad (3.24)$$

The velocity gradient for this flow is given by

$$\nabla \mathbf{u} = \begin{pmatrix} \cos(x_1 + A \sin \Omega t) \cos x_2 & -\sin(x_1 + A \sin \Omega t) \sin x_2 \\ \sin(x_1 + A \sin \Omega t) \sin x_2 & -\cos(x_1 + A \sin \Omega t) \cos x_2 \end{pmatrix},\quad (3.25)$$

which when evaluated at $\bar{\mathbf{x}}$ becomes

$$\nabla \mathbf{u}(\bar{\mathbf{x}}(t), t) = \begin{pmatrix} \cos(\bar{x}_1(t) + A \sin \Omega t) & 0 \\ 0 & -\cos(\bar{x}_1(t) + A \sin \Omega t) \end{pmatrix}\quad (3.26)$$

so that

$$y_j \frac{\partial}{\partial x_j} u_i(\bar{\mathbf{x}}(t), t) = \cos(\bar{x}_1(t) + A \sin \Omega t) \begin{pmatrix} y_1 \\ -y_2 \end{pmatrix}.\quad (3.27)$$

The velocity difference between $\bar{\mathbf{x}}(t) + \mathbf{y}$ and $\bar{\mathbf{x}}(t)$ is given by

$$\mathbf{u}(\bar{\mathbf{x}}(t) + \mathbf{y}, t) - \mathbf{u}(\bar{\mathbf{x}}(t), t) = \begin{pmatrix} \sin(\bar{x}_1(t) + y_1 + A \sin \Omega t) \cos y_2 \\ -\cos(\bar{x}_1(t) + y_1 + A \sin \Omega t) \sin y_2 \end{pmatrix} - \begin{pmatrix} \sin(\bar{x}_1(t) + A \sin \Omega t) \\ 0 \end{pmatrix}. \quad (3.28)$$

Inserting these quantities into equation (2.12) and using various trigonometric identities, $H(t)$ can then be defined as the set of points $(\bar{x}_1(t) + y_1, y_2)$ satisfying

$$\begin{aligned} \cos^2(\bar{x}_1(t) + A \sin \Omega t) (y_1^2 + y_2^2) &> E^2 \{ [\cos(\bar{x}_1(t) + A \sin \Omega t) (y_2 - \cos y_1 \sin y_2) \\ &+ \sin(\bar{x}_1(t) + A \sin \Omega t) \sin y_1 \sin y_2]^2 \\ &+ [\sin(\bar{x}_1(t) + A \sin \Omega t) (\cos y_1 \cos y_2 - 1) \\ &+ \cos(\bar{x}_1(t) + A \sin \Omega t) (\sin y_1 \cos y_2 - y_1)]^2 \}. \end{aligned} \quad (3.29)$$

Figure 9(a) shows the HNs as calculated from equation (3.29) for three different times. Notice that in this unsteady situation, the HN is *not* symmetric with respect to the location of the hyperbolic trajectory, unlike in the steady cases discussed above. As the hyperbolic trajectory moves back and forth, it drags the neighborhood in which exponential attraction/repulsion reigns back and forth with it, but *because it is moving in the x_1 -direction*, this region is not symmetric with respect to x_1 . Also, the HN appears to move much more significantly than does the hyperbolic trajectory to which it is pinned. In contrast, there is symmetry with respect to the x_2 -direction because there is no hyperbolic trajectory movement in that direction. Therefore, the concept of the HN depends strongly on the Lagrangian nature of the hyperbolic trajectory, but at each instant in time identifies regions that are subject to exponential stretching or contraction as a result of this trajectory. The effect of varying E —which identifies the relative importance of linear effects over nonlinear ones in the neighborhood—is shown in figure 9(b). As E gets very large, the neighborhoods get smaller, but retain their basic structure. Importantly, they are significant patches even when the neighborhood is defined by the linear terms being 200 times greater than the nonlinear ones. Thus, exponential stretching operates over much more than an infinitesimal region near the hyperbolic trajectory. The asymmetry with respect to the hyperbolic trajectory becomes more pronounced at larger E , highlighting the fact that its instantaneous region of influence cannot be naively assumed to be symmetric (as the conceptual figure 1 might suggest), but rather depends strongly on the Lagrangian motion of the hyperbolic trajectory. Adjacent particles entering a hyperbolic neighborhood will experience an exponential boost to their separation while within it.

In figure 10, the relationship between the HN and the stable (green) and unstable (red) manifolds is pictured. It is easy to see that the unstable manifold of the hyperbolic point remains on the line $x_2 = 0$ even for $A \neq 0$, since $x_2 = 0$ is invariant for the unsteady flow (3.23). The time-varying stable manifold is more difficult to determine. We use Theorem 2.7 from Balasuriya (2011) to find its expression correct to $\mathcal{O}(A)$ at each time t , and display this as the dashed green curve. It is indeed possible to prove using Theorem 2.2 of Balasuriya (2016b) that, to $\mathcal{O}(A)$, the stable manifold continues to point vertically upwards from the hyperbolic trajectory, but when progressing away from the hyperbolic point it has a time-varying curvature. Figure 10 indicates that the HNs have no degeneracy, and so stretching accumulation will occur in the generic way. We note that there is a slight discrepancy between where the stable manifold intersects the




FIGURE 10. The HNs and hyperbolic trajectory locations (blue) of the oscillating Taylor-Green flow (3.23) near the origin at two different times, using $A = 0.2$, $\Omega = 2\pi$ and $E = 50$, with the accompanying stable (dashed green) and unstable (dashed red) manifolds also shown.

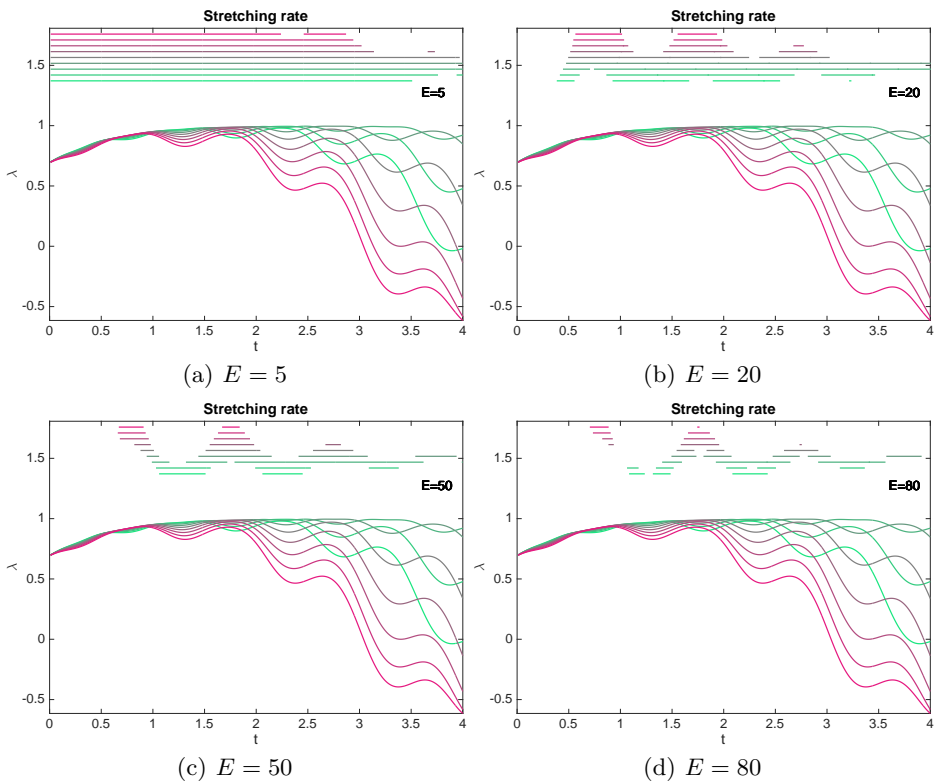


FIGURE 11. The stretching measure (3.10) computed for the oscillating Taylor-Green flow (3.23) with $A = 0.2$ and $\Omega = 2\pi$, for pairs of trajectories seeded near $(0, 0.8)$. The corresponding time in which a trajectory is inside an HN (defined according to different choices of E as shown in each panel) is indicated by the identically colored line towards the top of each figure.

HN and where the HN ‘sticks out.’ This is because we are comparing a Lagrangian entity (the stable manifold) with an instantaneous entity (the HN) that *impacts* the Lagrangian motion. The instantaneous HN is continually trying to adjust itself to keep us with the time variation, and there must therefore be an effective time-lag.

Next, we numerically verify that trajectories experience an exponential boost when entering an HN. We use the stretching measure introduced in (3.10), and seed pairs of trajectories near $(0, 0.8)$, close to the stable manifold of the origin. Figure 11 shows the

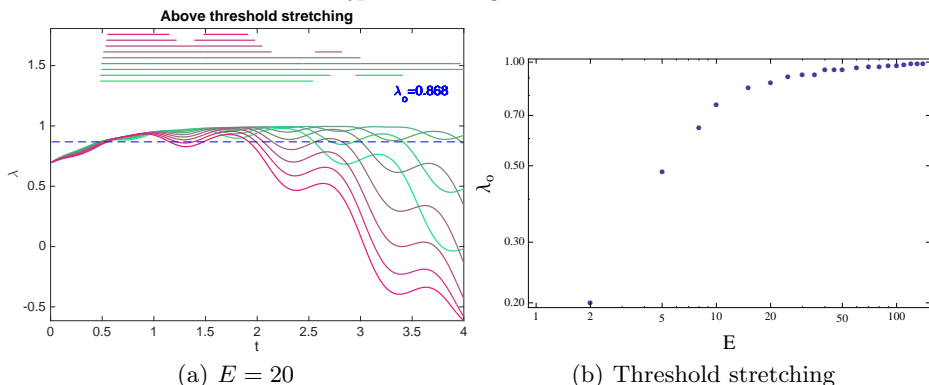


FIGURE 12. (a) The stretching rates of figure 11(b), but with the range over which λ is larger than a threshold value λ_o indicated by the lines at the top of the figure. (b) λ_o values as functions of E plotted on logarithmic axes for the oscillating Taylor–Green flow, chosen by minimizing (3.13), with identical parameter values as for figure 11 being used.

stretching, with each pair of trajectories having been assigned a particular color. The domain in which the pair is deemed to be within the HN is shown by the straight lines at the top of the figures. We show the computations for four different values of E . The stretching is clearly larger when within an HN. The fact that the HNs are oscillating with time here, unlike in the steady Taylor–Green situation, means that there are instances when trajectories leave an HN to be back inside soon afterwards, which is shown by the interrupted nature of the time-domain lines at the top of figure 11. As expected, the stretching is larger when inside an HN, and its increase when entering and decrease when leaving an HN is clearly visible.

In figure 12, we present the relationship between E and an optimum threshold λ_o as defined in §3.1. Figure 12(a), with $\lambda_o = 0.868$, is to be compared with figure 11(b), with $E = 20$. The time ranges at the top of the figures obtained using these values are almost identical. The computation of the optimum λ_o for each E value, shown in figure 12(b), in all cases is associated with an error (3.13) of approximately 10^{-4} , indicating that excellent optimality is obtained. Once again, the threshold λ_o appears to approach a constant value, and for $E \gtrsim 20$, about 90% of this stretching rate is achieved.

To verify that the accumulated stretching, measured via forward-time FTLEs, is related to the residence time within the HNs, we provide figure 13. The choice $E = 10$ has been made for the HN located near the origin at time 0, and in this nondegenerate HN situation, we find that the residence time is well correlated to the residence time within the HN. Specifically, we note the qualitative similarity between the residence time and FTLE plots (how far the ridges extend, the region in which they are more peaked, where they get fatter, etc). The horizontal ridge appearing in figure 13(f) occurs because of the stable manifold connected to the hyperbolic trajectories located near $(\pm\pi, 0)$, which is also responsible for the larger off-ridge FTLE values appearing alongside it. The time of integration is such that particles in the chosen domain approach $(\pm\pi, 0)$ in this case, in contrast to shorter times of integration (though a hint of this issue appears in (d)). These features in (f) are not picked up in (e) because only the HN near $(0, 0)$ has been used. The fact that this flow is genuinely unsteady has not impeded the residence time in an HN as being an excellent proxy for accumulated stretching.




FIGURE 13. The residence times in the HN with $E = 10$ (left), compared with the FTLE field (right) for the oscillating Taylor-Green flow with $A = 0.2$ and $\Omega = 2\pi$, computed over different times from $t = 0$ to $t = 1$ (top), to $t = 3$ (middle) and to $t = 7$ (bottom).

4. Experimental Test

Analytical examples are very useful for developing intuition about what HNs mean in situations where we know a great deal about the underlying dynamical system. The main purpose of these ideas, though, is their application in real experimental or observational data. Thus, it is important to show that the entities we define are actually computable in such a system. To that end, we here demonstrate the calculation of HNs using data from a turbulent quasi-two-dimensional laboratory flow. The insights we obtained from the analytical examples will be useful in this computation.

4.1. Apparatus and Data Acquisition

To generate nearly two-dimensional flow in the lab, we used an electromagnetically driven thin-layer flow cell we have described in detail elsewhere (Kelley & Ouellette 2011*a*; Liao & Ouellette 2013). The working fluid is contained in a cell of lateral dimension $86 \times 86 \text{ cm}^2$, supported by a glass substrate coated in a hydrophobic wax. The fluid itself is a thin layer (roughly 5 mm deep) of an electrolyte (16% NaCl by mass in deionized water), and so conducts electricity. We use a pair of copper bar electrodes mounted at opposite ends of the flow cell to drive d.c. electric currents of up to a few amperes through the fluid. By itself, the current does not generate flow. Underneath the glass substrate, however, lies a square grid of permanent disc magnets with their dipoles pointed vertically; the combination of the lateral current density and the vertical magnetic field applies a Lorentz body force to the fluid, which in turns generates flow. And aside from weak fringing fields from the magnets, this force lies in the plane, and so drives the fluid in a nearly two-dimensional way. The magnets themselves measure 12.7 mm in diameter, and are placed with a center-to-center spacing of 25.4 mm. For the experiments analyzed here, they were arranged in a checkerboard fashion with the north and south poles alternating in each direction.

To measure the flow field, we use particle tracking velocimetry (PTV). We seed the fluid with small (51 μm diameter) fluorescent polystyrene microspheres that absorb in the blue and emit in the green. Because these particles are lighter than the electrolyte (with a specific gravity of 1.05), they float to its surface. On the surface, though, they would experience a long-range attractive force due to surface tension. Thus, to remove this effect, we float a second layer, also 5 mm deep, of fresh water on top of the electrolyte. This miscible interface has no bulk surface tension, and so removes the particle-particle attraction, thereby rendering them faithful flow tracers (Ouellette *et al.* 2008). We illuminate the particles with blue LEDs and record their motion at a rate of 60 frames per second with an IDT MotionPro M5 camera fitted with a Schneider APO-Xenoplan lens. The camera has a resolution of 2320×1728 pixels, with which we image the central $31.7 \times 23.6 \text{ cm}^2$ of the experiment (approximately 12×9 magnet spacings). For the data we show here, the Reynolds number based on the in-plane velocity and the spacing between the magnets is $\text{Re} = 185$, high enough to make the flow weakly turbulent (so that it shows a net inverse energy flux from small to large scales (Kelley & Ouellette 2011*b*), as is expected in two-dimensional turbulence) but not so large as to generate significant three-dimensional motion (Kelley & Ouellette 2011*a*).

Images from the camera are processed to generate particle trajectories using a multi-frame, predictive Lagrangian particle tracking algorithm (Ouellette *et al.* 2006). Accurate particle velocities are then obtained by convolving the trajectories with a differentiating and smoothing kernel (Mordant *et al.* 2004). Since we track roughly 35,000 particles per frame, the instantaneous velocity measurements are spatially dense enough that we can interpolate between them to create highly resolved Eulerian velocity fields. To condition these fields further and ensure that they are as two-dimensional as possible, we project them onto a numerically computed basis of streamfunction eigenmodes (Kelley & Ouellette 2011*a*). Subsequently, to generate Lagrangian statistics, we integrate the equations of motion for fluid elements through these velocity fields using second-order Runge–Kutta integration (Ouellette *et al.* 2008) to generate “virtual” trajectories with any desired initial conditions.

4.2. Extraction of Hyperbolic Trajectories

Finding hyperbolic trajectories in an arbitrary data set can be a challenging task. The mere *definition* of hyperbolic trajectories in finite-time situations is fraught with difficulties, as argued for example in Balasuriya (2016b). Nevertheless, there is a preponderance of methods which can be adapted to find *their own definition* of hyperbolic trajectories, as discussed in the Introduction. The ambiguity of exactly what a hyperbolic trajectory is (as reflected by the various ‘definitions’ in use, which lead to *different* identifications in general), reveals an advantage of using the idea of HNs. It seems that most methods, applied with judicious interpretation and testing to a particular time-varying velocity data set, have the potential for zeroing in on *essentially* the same hyperbolic trajectory. However, its exact positioning will inevitably be slightly different for the different methods. Having the concept of an HN, which locates a *region* as opposed to a *point*, mitigates this uncertainty in the knowledge of the *exact* location.

Most of these methods (e.g., hyperbolic LCSs, Perron-Frobenius operator methods) have absolutely *no* direct evaluation of *exponential stretching*, which is what we are after. Given that this is our specific focus, we will use the method for hyperbolic trajectory identification that is explicitly designed for exponential stretching: FTLEs (Shadden *et al.* 2005; Allshouse & Peacock 2015; Huntley *et al.* 2015; He *et al.* 2016; Lee *et al.* 2007; Waugh *et al.* 2007; Johnson & Meneveau 2015; Balasuriya 2015a; Samelson 2013; Voth *et al.* 2002) as defined in equation (2.8). Thus, we will seek intersections between (sharp enough) ridges of the forward- and backward-time FTLE fields to determine our hyperbolic trajectories (Voth *et al.* 2002). While this process *a priori* does not guarantee that the identified hyperbolic trajectory is a genuine *trajectory* of the flow (Haller 2015, see also figure 7(b,d,e)), we have shown that as long as the HNs remain nondegenerate, then this process conforms to the intuitive picture of figure 1, and thus intersections of stable and unstable manifolds are expected to be obtained with increasing accuracy. Checking that HNs are nondegenerate can therefore be thought of as a test for validating that FTLE ridges do indeed provide the ‘correct’ understanding of how the exponential behavior affects fluid elements. We will *a posteriori* check whether the nondegeneracy property is satisfied.

In figure 14, we show the forward-time FTLE field for fixed t_0 and T for our experimental flow. The field tends to organize itself into a pattern of sharp ridges. These ridges, which identify the regions of the flow field that experience the strongest stretching in the time range from t_0 to $t_0 + T$, mark the analogs of stable (when $T > 0$) and unstable (when $T < 0$; that is, in backwards time) manifolds. The intersections of these forward-time and backward-time ridges are then likely to mark the instantaneous positions of hyperbolic trajectories. Intersections may also be heteroclinic trajectories, but these secondary intersections can be made to be unusual in practice by investigating the behavior of ridges as T decreases, and verifying that FTLE ridges shorten towards the hyperbolic trajectory.

To extract the line-like ridges from this field data, so as to identify their forward- and backward-time intersections, we use an algorithm based on the techniques described by Senatore & Ross (2011). Points that lie on or near the ridges (ridge points), shown in green in figure 14, are initially found by calculating the Hessian of the FTLE field, and identifying regions where the minimum eigenvalue, and the dot product of the corresponding eigenvector with the gradient of the FTLE field, lie within relevant tolerances. A dynamical sharpening scheme is then used to integrate the ridge points along streamlines of the field, towards the crest of the nearest ridge. Integration is stopped when the required termination condition, derived from Mathur *et al.* (2007), is reached;




FIGURE 14. A forward-time FTLE field (blue) computed from the experimental data for a region measuring approximately 12×9 magnet spacings. Ridges of this field, shown in green were extracted using the procedure described in the text. The red points are the result of applying the dynamical sharpening scheme.

this process results in a set of points that identify the ridges. Shown in red in figure 14, these are the FTLE analogs to the unstable manifolds, and each curve must go across a hyperbolic trajectory location, as yet unknown. A similar calculation (not shown), enables the extraction of the backward-time FTLE ridges at this same instance in time. Finally, once the ridges are found in forward and backward time, it is straightforward to identify their intersection points. We treat these points as the instantaneous locations of hyperbolic trajectories.

4.3. Results

Once hyperbolic trajectories have been identified, we can apply equation (2.12) to determine the hyperbolic neighborhood $H(t)$ associated with each one at each time instant. An example is shown in figure 15, where we have used the same time instant as shown in figure 14 and have shown all the hyperbolic trajectories identified by our algorithm (which may only be a subset of the true ensemble of hyperbolic trajectories present). The instantaneous hyperbolic neighborhoods for $E = 2$ are shown by the black contours; the resolution of the velocity field is not sufficient to resolve HNs computed for significantly larger values of E consistently well. Unlike in the analytical examples discussed in §3, the HNs here are much more irregularly shaped, primarily because the flow field itself is spatially disordered. Even so, they are clearly nondegenerate, and share many features with the simple analytical examples. The HNs tend to be elongated along the stable and unstable manifolds associated with each hyperbolic trajectory, although they also extend a finite distance in other directions. It is also interesting to note that the areas of the various neighborhoods are not at all the same; instead, some of the hyperbolic trajectories have an influence over a much larger region of space than others. This observation has significant ramifications. It is certainly reasonable to think that not every hyperbolic trajectory (along with its associated manifolds) should play an equal role in the flow dynamics; it is possible that the area of the associated hyperbolic




FIGURE 15. Several hyperbolic neighborhoods superimposed on the stretching field (that is, the field of Cauchy–Green eigenvalues) from one snapshot of experimental data. Blue colors show the forward-time stretching (that is, the stable manifolds), and red colors show the backward-time stretching (that is, the unstable manifolds). The instantaneous locations of some hyperbolic trajectories are shown with black dots, and the corresponding hyperbolic neighborhoods (calculated using $E = 2$) are shown with black contours. The black scale bar shows the spacing between magnets, and the Cauchy–Green tensor was computed over a time T of approximately the correlation time of the velocity field.

neighborhood may provide a way to quantify just how important a given hyperbolic trajectory is. If E were changed uniformly across the domain, the *relative* sizes of the regions of influence remains unaffected. Thus, the determination of which hyperbolic trajectories are more influential over their surroundings can be made by simply choosing any reasonable E , for example $E = 2$ in figure 15.

There are regions in which HNs overlap, as in the top left of figure 15. Trajectories in such regions experience instantaneous stretching in the unstable manifold directions associated with *each* HN to which they belong. The total stretching will be a superposition of all of the directional stretchings associated with each HN the trajectory is in.

As time evolves, the HNs in this flow evolve as well, moving about and changing their shape as their associated hyperbolic trajectory moves. As they are not material areas, their size is not conserved, and they may grow or shrink with time. Far from being a failing of our technique, this shrinking contains useful information: once a hyperbolic trajectory ceases to be hyperbolic, we would expect its zone of influence to be small. This is reflected in the definition of HNs, as given in (2.12). If the linear terms vanish in comparison with the nonlinear term $u_i(\bar{\mathbf{x}}(t) + \mathbf{y}, t) - u_i(\bar{\mathbf{x}}(t), t)$, then the definition of the HN amounts to finding regions satisfying $0 > E$ where $E \geq 1$. That is, the HN disappears.

Finally, we examine the influence of the HNs on the stretching, as we did for our analytical examples in §3. The situation here is, however, somewhat different, both for




FIGURE 16. The mean stretching rate λ in the experimental flow, averaged over all the HNs shown in figure 15, as described in the text. Times are normalized by the large eddy turnover time T_L . The solid, red curve shows the stretching rate for pairs initially in the HNs, while the dashed, black line shows the stretching rate for pairs seeded uniformly in the domain.

practical and physical reasons. Unlike in our examples the HNs in this flow exist only for finite times and move about irregularly; thus, it is difficult to seed pairs of Lagrangian points that we know will eventually both enter an HN. Additionally, pair separation behaves differently in turbulent flows like this one and in chaotic flows like our analytical examples: trajectories separate according to a power law rather than exponentially (Sawford 2001; Bourgoin *et al.* 2006). In two-dimensional turbulence, however, there is expected to be a transient exponential range for initially nearby particle pairs as long as their separation remains in the range of scales in which the enstrophy cascade operates (Rivera & Ecke 2005). Thus, applying eq. 3.10 is still reasonable for initially close pairs.

We therefore computed the mean stretching rate λ , as defined in eq. 3.10, averaged over both an ensemble of Lagrangian pairs initially in the HNs shown in figure 15 (computed with $E = 2$) and an ensemble seeded uniformly in the flow domain. Data for these two cases are shown in figure 16. The pairs used for this computation were initially separated by 2% of the magnet spacing L_m , and times are normalized by the eddy turnover time T_L of eddies at the magnet scale, computed as $T_L = L_m/U$, where U is the root-mean-square velocity. Unlike the cases shown in figures 3, 6, and 11, we do not pinpoint when individual pairs enter or leave HNs in figure 16. It is clear, however, that the pairs that start in HNs do indeed experience enhanced stretching relative to a typical pair placed arbitrarily in the flow domain. Thus, even in this complex, highly unsteady flow, HNs still play a demonstrable role in boosting the Lagrangian stretching experienced by the fluid.

5. Discussion and Conclusions

We have described here a new method for estimating the instantaneous zone of influence of a hyperbolic trajectory in arbitrary, unsteady flow by comparing the spatial structure of the velocity field to the linearized estimate around the hyperbolic trajectory. By directly analyzing the linearized velocity, we have avoided ambiguities in attempting to define when a finite-time function exhibits exponential decay. We have established strong connections with FTLEs, and shown how the geometrical structure of this zone of influence can affect conclusions that can be reached through examining FTLE fields. We have verified our claim that fluid elements entering these zones of influence exhibit exponential stretching, which diminishes as they leave the zone. Thus, these zones can be

thought of as traveling regions that impart anomalously large stretching on fluid elements entering them.

The essential assumption underlying this technique is the notion that if a hyperbolic trajectory (or indeed any kind of coherent structure) is to be important in a given region of the flow field, then its characteristics should be dominant to leading order in that region. In addition to being reasonable, we also note that this assumption makes our proposed technique very flexible. Here, we have defined hyperbolic neighborhoods, in that we are characterizing the zone of influence of hyperbolic trajectories; nothing stops us, however, from extending our method to defining, e.g., elliptic or parabolic neighborhoods, as long as the underlying point-like elliptic or parabolic structure is known. We thus anticipate that the ideas presented here may find wider application than just searching for hyperbolic structures, just as the ideas behind the original hyperbolic LCSs were later extended to include elliptic and parabolic structures (Haller 2015).

That being said, there is still room for refinement in the ideas we have presented. As has been argued compellingly, if *coherent transport* is the goal, then structures that evolve in a Lagrangian sense with the flow are most worthy of investigation. While hyperbolic trajectories are Lagrangian, the hyperbolic neighborhoods we have defined do not follow the flow. That is, particles within $H(t)$ at some time do not always remain within $H(t)$ as t evolves. This occurs because $H(t)$ is defined as a moving region of influence, rather than a moving group of particles. Its defining characteristic is that while within $H(t)$, fluid elements will experience instantaneous exponential stretching, as we have demonstrated through, for example, figures 3, 6, and 11. Residence time within $H(t)$ is also directly correlated to the stretching incurred over a given time period, as we have shown in figures 7 and 13. But unlike e.g., the Okubo-Weiss and associated criteria, the hyperbolic neighborhoods are also not fully Eulerian, since they are pinned to the (Lagrangian) hyperbolic trajectory. In future work, we plan to explore ways of extending the Lagrangian character of the core hyperbolic trajectory to its corresponding neighborhood by, for example, looking for advected sets that remain in instantaneous hyperbolic neighborhood for long times.

A second productive path forward we can identify is to understand more deeply the role of the coefficient E . We have in figures 4 and 12 shown that E is linked to choosing a stretching threshold, i.e., how close to the stretching rate experienced at the hyperbolic point the stretching rate must be to be considered in the HN. The geometrical deformation of the HNs in the shrinking limit $E \rightarrow \infty$ has been shown in § 3.3 to impact on FTLEs in nontrivial ways. The value of E can also likely be connected to the spatial resolution used in any FTLE or other Lagrangian advection computation, since HNs will be ‘visible’ in such situations only for small enough E .

Finally, hyperbolic neighborhoods may be a useful way to describe the stretching of transported material within a flow field. For example, consider a patch of oil that is released into an unsteady velocity field. The stretching of the oil elements will be fundamentally impacted by hyperbolic neighborhoods in the sense that whenever some oil elements enter a hyperbolic neighborhood, they will stretch exponentially while within it. The oil patch’s stretching over some finite time interval will therefore be strongly governed by the appearance/disappearance, locations, and time-variations of hyperbolic neighborhoods, in relation to the location of oil particles. Hyperbolic neighborhoods thus give us a new way to think of (an aspect of) mixing, which we will pursue in the future.

The authors acknowledge support through a Terman Faculty Fellowship from Stanford University (N.T.O.), a Summer Research Scholarship from the University of Adelaide

(R.K.), and Future Fellowship Grant FT130100484 from the Australian Research Council (S.B.).

REFERENCES

- ALLSHOUSE, M. & PEACOCK, T. 2015 Refining finite-time Lyapunov exponent ridges and the challenges of classifying them. *Chaos* **25**, 087410.
- ALLSHOUSE, M. R. & THIFFEAULT, J.-L. 2012 Detecting coherent structures using braids. *Physica D* **241**, 95–105.
- BABIANO, A., CARTWRIGHT, J. H. E., PIRO, O. & PROVENZALE, A. 2000 Dynamics of a small neutrally buoyant sphere in a fluid and targeting in Hamiltonian systems. *Phys. Rev. Lett.* **84**, 5764–5767.
- BALASURIYA, S. 2010 Optimal frequency for microfluidic mixing across a fluid interface. *Phys. Rev. Lett.* **105**, 064501.
- BALASURIYA, S. 2011 A tangential displacement theory for locating perturbed saddles and their manifolds. *SIAM J. Appl. Dyn. Sys.* **10**, 1100–1126.
- BALASURIYA, S. 2012 Explicit invariant manifolds and specialised trajectories in a class of unsteady flows. *Phys. Fluids* **24**, 12710.
- BALASURIYA, S. 2015a Dynamical systems techniques for enhancing microfluidic mixing. *J. Micromech. Microeng.* **25**, 094005.
- BALASURIYA, S. 2015b Quantifying transport within a two-cell microdroplet induced by circular and sharp channel bends. *Phys. Fluids* **27**, 052005.
- BALASURIYA, S. 2016a *Flow Barriers and Transport in Unsteady Flows: A Melnikov Approach*. Philadelphia: SIAM Press (in press).
- BALASURIYA, S. 2016b Local stable and unstable manifolds and their control in nonautonomous finite-time flows. *J. Nonlin. Sci.* **26**, 895–927.
- BALASURIYA, S. & FINN, M.D. 2012 Energy constrained transport maximization across a fluid interface. *Phys. Rev. Lett.* **108**, 244503.
- BALASURIYA, S. & PADBERG-GEHLE, K. 2014 Accurate control of hyperbolic trajectories in any dimension. *Phys. Rev. E* **90**, 032903.
- BASDEVANT, C. & PHILIPOVITCH, T. 1994 On the validity of the “Weiss criterion” in two-dimensional turbulence. *Physica D* **73**, 17–30.
- BLAZEWSKI, D. & HALLER, G. 2014 Hyperbolic and elliptic transport barriers in three-dimensional unsteady flows. *Physica D* **273–274**, 46–64.
- BOURGOIN, M., OUELLETTE, N. T., XU, H., BERG, J. & BODENSCHATZ, E. 2006 The role of pair dispersion in turbulent flow. *Science* **311**, 835–838.
- BUDIŠIĆ, M. & MEZIĆ, I. 2012 Geometry of the ergodic quotient reveals coherent structures in flows. *Physica D* **241**, 1255–1269.
- BUDIŠIĆ, M. & THIFFEAULT, J.-L. 2015 Finite-time braiding exponents. *Chaos* **25**, 087407.
- CHONG, M. S., PERRY, A. E. & CANTWELL, B. J. 1990 A general classification of three-dimensional flow fields. *Phys. Fluids A* **2**, 765–777.
- COPPEL, W. A. 1978 *Dichotomies in Stability Theory. Lecture Notes Math.* 629. Berlin: Springer-Verlag.
- DELLNITZ, M. & JUNGE, O. 2002 Set oriented numerical methods for dynamical systems. In *Handbook of Dynamical Systems*, , vol. 2, pp. 221–264. Elsevier.
- DOAN, T.S., KARRASCH, D., NGUYEN, T. & SIEGMUND, S. 2012 A unified approach to finite-time hyperbolicity which extends finite-time Lyapunov exponents. *J. Differential Equations* **252**, 5355–5554.
- FROYLAND, G., LLOYD, S. & SANTITISSADEEKORN, N. 2010 Coherent sets for nonautonomous dynamical systems. *Physica D* **239**, 1527–1541.
- FROYLAND, G. & PADBERG, K. 2009 Almost-invariant sets and invariant manifolds – connecting probabilistic and geometric descriptions of coherent structures in flows. *Physica D* **238**, 1507–1523.
- FROYLAND, G. & PADBERG-GEHLE, K. 2012 Finite-time entropy: a probabilistic approach for measuring nonlinear stretching. *Physica D* **241**, 1612–1628.
- FROYLAND, G. & PADBERG-GEHLE, K. 2015 A rough-and-ready cluster-based approach for

- extracting finite-time coherent sets from sparse and incomplete trajectory data. *Chaos* **25**, 087406.
- GRIGORIEV, R.O. 2005 Chaotic mixing in thermocapillary-driven microdroplets. *Phys. Fluids* **17**, 033601.
- GUCKENHEIMER, J. & HOLMES, P. 1983 *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*. New York: Springer.
- HALLER, G. 2005 An objective definition of a vortex. *J. Fluid Mech.* **525**, 1–26.
- HALLER, G. 2011 A variational theory of hyperbolic Lagrangian coherent structures. *Physica D* **240**, 574–598.
- HALLER, G. 2015 Lagrangian Coherent Structures. *Annu. Rev. Fluid Mech.* **47**, 137–161.
- HALLER, G., HADJIGHASEM, A., FARAZMAND, M. & HUHN, F. 2016 Defining coherent vortices objectively from the vorticity. *J. Fluid Mech.* **795**, 136–173.
- HE, G.-S., PAN, C., FENG, L.-H., GAO, Q. & WANG, J.-J. 2016 Evolution of Lagrangian coherent structures in a cylinder-wake disturbed flat boundary layer. *J. Fluid Mech.* **792**, 274–306.
- HUA, B. L. & KLEIN, P. 1998 An exact criterion for the stirring properties of nearly two-dimensional turbulence. *Physica D* **113**, 98–110.
- HUNT, J. C. R., WRAY, A. & MOIN, P. 1988 Eddies, stream, and convergence zones in turbulent flows. *Tech. Rep.* CTR-S88. Center for Turbulence Research.
- HUNTLEY, H., LIPPHARDT, B., JACOBS, G. & KIRWAN, A. 2015 Clusters, deformation, and dilation: Diagnostics for material accumulation regions. *J. Geophys. Res.* **120**, 6622–6636.
- JEONG, J. & HUSSAIN, F. 1995 On the identification of a vortex. *J. Fluid Mech.* **285**, 69–94.
- JOHNSON, P. & MENEVEAU, C. 2015 Large deviation joint statistics of the finite-time Lyapunov spectrum in isotropic turbulence. *Phys. Fluids* **27**, 085110.
- KARRASCH, D. 2013 Linearization of hyperbolic finite-time processes. *J. Differential Equations* **254**, 256–282.
- KARRASCH, D. 2015 Attracting Lagrangian coherent structures on Riemannian manifolds. *Chaos* **25**, 087411.
- KELLEY, D. H. & OUELLETTE, N. T. 2011a Onset of three-dimensionality in electromagnetically forced thin-layer flows. *Phys. Fluids* **23**, 045103.
- KELLEY, D. H. & OUELLETTE, N. T. 2011b Spatiotemporal persistence of spectral fluxes in two-dimensional weak turbulence. *Phys. Fluids* **23**, 115101.
- KHURANA, N., BLAWZDZIEWICZ, J. & OUELLETTE, N. T. 2011 Reduced transport of swimming particles in chaotic flow due to hydrodynamic trapping. *Phys. Rev. Lett.* **106**, 198104.
- KHURANA, N. & OUELLETTE, N. T. 2012 Interactions between active particles and dynamical structures in chaotic flow. *Phys. Fluids* **24**, 091902.
- LEE, Y.-K., SHIH, C., TABELING, P. & HO, C.-M. 2007 Experimental study and nonlinear dynamics analysis of time-periodic micro chaotic mixers. *J. Fluid Mech.* **575**, 425–448.
- LIAO, Y. & OUELLETTE, N. T. 2013 Spatial structure of spectral transport in two-dimensional flow. *J. Fluid Mech.* **725**, 281–298.
- MA, T. & BOLLT, E. 2014 Differential geometry perspective of shape coherence and curvature evolution by finite-time nonhyperbolic splitting. *SIAM J. Appl. Dyn. Sys.* **13**, 1106–1136.
- MATHUR, M., HALLER, G., PEACOCK, T., RUPPERT-FELSOT, J. E. & SWINNEY, H. L. 2007 Uncovering the Lagrangian skeleton of turbulence. *Phys. Rev. Lett.* **98**, 144502.
- MEZIĆ, I. 2013 Analysis of fluid flows via spectral properties of the Koopman operator. *Annu. Rev. Fluid Mech.* **45**, 357–378.
- MEZIĆ, I., LOIRE, S., FONOBEROV, V. A. & HOGAN, P. 2010 A new mixing diagnostic and Gulf oil spill movement. *Science* **330**, 486–489.
- MORDANT, N., CRAWFORD, A. M. & BODENSCHATZ, E. 2004 Experimental Lagrangian probability density function measurement. *Physica D* **193**, 245–251.
- NISHIKAWA, T., TOROCZKAI, Z. & GREBOGI, C. 2001 Advective coalescence in chaotic flows. *Phys. Rev. Lett.* **87**, 038301.
- OKUBO, A. 1970 Horizontal dispersion of floatable particles in the vicinity of velocity singularities such as convergences. *Deep-Sea Res.* **17**, 445–454.
- OLASCOAGA, M. J. & HALLER, G. 2012 Forecasting sudden changes in environmental pollution patterns. *Proc. Natl. Acad. Sci. USA* **109**, 4738–4743.

- OUELLETTE, N. T. & GOLLUB, J. P. 2007 Curvature fields, topology, and the dynamics of spatiotemporal chaos. *Phys. Rev. Lett.* **99**, 194502.
- OUELLETTE, N. T. & GOLLUB, J. P. 2008 Dynamic topology in spatiotemporal chaos. *Phys. Fluids* **20**, 064104.
- OUELLETTE, N. T., HOGG, C. A. R. & LIAO, Y. 2016 Correlating Lagrangian structures with forcing in two-dimensional flow. *Phys. Fluids* **28**, 015105.
- OUELLETTE, N. T., O'MALLEY, P. J. J. & GOLLUB, J. P. 2008 Transport of finite-sized particles in chaotic flow. *Phys. Rev. Lett.* **101**, 174504.
- OUELLETTE, N. T., XU, H. & BODENSCHATZ, E. 2006 A quantitative study of three-dimensional Lagrangian particle tracking algorithms. *Exp. Fluids* **40**, 301–313.
- PEACOCK, T., FROYLAND, G. & HALLER, G. 2015 Introduction to focus issue: Objective detection of coherent structures. *Chaos* **25**, 087201.
- PEACOCK, T. & HALLER, G. 2013 Lagrangian coherent structures: the hidden skeleton of fluid flows. *Physics Today* **66**, 41.
- PERRY, A. E. & CHONG, M. S. 1987 A description of eddy motions and flow patterns using critical-point concepts. *Annu. Rev. Fluid Mech.* **19**, 125–155.
- PIERREHUMBERT, R. 1991 Chaotic mixing of tracer and vorticity by modulated travelling Rossby waves. *Geophys. Astrophys. Fluid Dyn.* **58**, 285–319.
- RIVERA, M. K., ALUIE, H. & ECKE, R. E. 2014 The direct enstrophy cascade of two-dimensional soap film flows. *Phys. Fluids* **26**, 055105.
- RIVERA, M. K. & ECKE, R. E. 2005 Pair dispersion and doubling time statistics in two-dimensional turbulence. *Phys. Rev. Lett.* **95**, 194503.
- ROM-KEDAR, V., LEONARD, A. & WIGGINS, S. 1990 An analytical study of transport, mixing and chaos in an unsteady vortical flow. *J. Fluid Mech.* **214**, 347–394.
- SAMELSON, R. 1992 Fluid exchange across a meandering jet. *J. Phys. Oceanogr.* **22**, 431–440.
- SAMELSON, R. 2013 Lagrangian motion, coherent structures, and lines of persistent material strain. *Annu. Rev. Marine Sci.* **5**, 137–163.
- SAWFORD, B. 2001 Turbulent relative dispersion. *Annu. Rev. Fluid Mech.* **33**, 289–317.
- SENATORE, C. & ROSS, S. D. 2011 Detection and characterization of transport barriers in complex flows via ridge extraction of the finite time Lyapunov exponent field. *Int. J. Numer. Meth. Engng.* **86**, 1163–1174.
- SHADDEN, S. C., LEKIEN, F. & MARSDEN, J. E. 2005 Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D* **212**, 271–304.
- SOLOMON, T. H. & GOLLUB, J. P. 1988 Chaotic particle transport in time-dependent Rayleigh-Bénard convection. *Phys. Rev. A* **38**, 6280–6286.
- STROOCK, A., DERTINGER, S., ADJARI, A., MEZIĆ, I., STONE, H. & WHITESIDES, G. 2002 Chaotic mixer for microchannels. *Science* **295**, 647–651.
- TORNEY, C. & NEUFELD, Z. 2007 Transport and aggregation of self-propelled particles in fluid flows. *Phys. Rev. Lett.* **99**, 078101.
- TRUESDELL, C. & NOLL, W. 2004 *The Non-Linear Field Theories of Mechanics*. New York: Springer.
- VOTH, G. A., HALLER, G. & GOLLUB, J. P. 2002 Experimental measurements of stretching fields in fluid mixing. *Phys. Rev. Lett.* **88**, 254501.
- WAUGH, D., ABRAHAM, E. & BOWEN, M. 2007 Spatial variations of stirring in the surface ocean: a case study of the Tasman Sea. *J. Phys. Oceanogr.* **36**, 526–542.
- WEISS, J. 1991 The dynamics of enstrophy transfer in two-dimensional hydrodynamics. *Physica D* **48**, 273–294.
- WIGGINS, S. 1992 *Chaotic Transport in Dynamical Systems*. New York: Springer.