Lagrangian Coherent Structures: Generalizing Stable and Unstable Manifolds to Non-Autonomous Dynamical Systems

Stuart Kent
Program in Applied Mathematics
University of Arizona

Spring 2008

Abstract

In this paper, Lagrangian Coherent Structures are presented as the extensions of stable and unstable manifolds to non-autonomous dynamical systems. Theoretical and numerical illustrations of Lagrangian Coherent Structures are given, and the results used to analyze some simple systems, both autonomous and non-autonomous. Finally, the application of Lagrangian Coherent Structures to empirical fluid flows is discussed.

1 Dynamical Systems: Background

The theory of dynamical systems combines quantitative and qualitative analysis of systems of differential equations, and was first expounded by Poincaré [1] in 1890 during his work on celestial mechanics. Many other dynamical systems also arise as models of real-world behavior, including oscillating pendulums, population dynamics and fluid flows. For 'simple' systems[1] there are many analytical concepts such as fixed points, periodic orbits, manifolds and attractors with which to analyze both localized and global features of the system. However, many of these features do not appear in systems with aperiodic time dependence. This paper discusses the use of Lagrangian Coherent Structures in non-autonomous dynamical systems as analogs of stable and unstable manifolds in autonomous systems.

1.1 Basic Notation

A *dynamical system* is a set of equations that describes the evolution of a system in time. In general, a dynamical system is of the form

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), t) \tag{1}$$
$$\mathbf{x}(t_0) = \mathbf{x}_0 \tag{2}$$

Here t is the independent variable, and the solution $\mathbf{x}(t)$ represents the state of the system at time t. The collection of all possible states of the system is called the *state space*, denoted D. If $\mathbf{f}$ has no explicit time dependence, then the system is called *autonomous*. It is assumed that the reader is familiar with other terms such as *fixed points* (both *hyperbolic* and *non-hyperbolic*), *phase portraits*, and so on.

The evolution of $\mathbf{x}(t)$ under such a system can be interpreted as a flow. Given an initial point $\mathbf{x}_0 \in D$, the solution $\mathbf{x}(t)$ will trace out a *trajectory* (also called an *orbit*) $C_{\mathbf{x}_0} \subset D$ as t increases.

[1] Here meaning: systems with either no explicit time dependence or a periodic time dependence.

By considering the evolution of $\mathbf{x}(t)$ for all initial conditions $\mathbf{x}_0$, we see that for sufficiently smooth $\mathbf{f}$ the trajectories will define a smooth flow everywhere in D. Before the formal definition, let us look at an illustrative example.

1.1.1 Example: Whirlpool

Consider the dynamical system

$$\dot{r} = -r^2 \tag{3}$$
$$\dot{\theta} = 1 \tag{4}$$

Note that $\dot{r}$ and $\dot{\theta}$ only depend explicitly upon the position variables (r and θ) - there is no explicit t dependence. Therefore we can produce a static image of the *velocity field* $\mathbf{v}(r,\theta) = (\dot{r}(r,\theta), \dot{\theta}(r,\theta))$ at all points in $D = \mathbb{R}^2$.

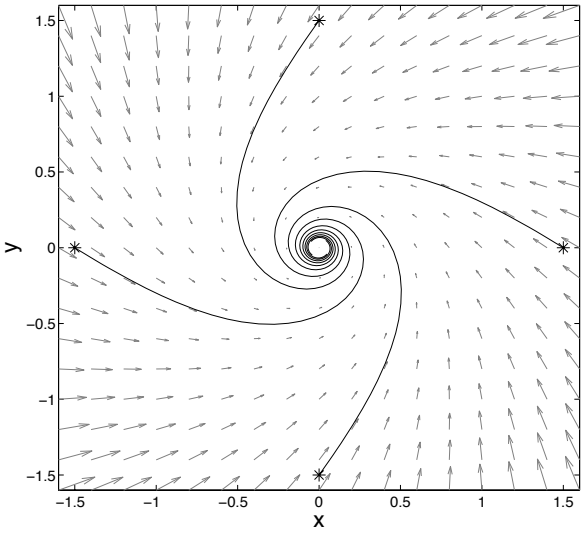

Figure 1: Portion of the velocity field for system (3)-(4).

The solution curves plotted in Fig. 1 behave exactly as if they are tracing out the path of a particle in a fluid flow. I chose this simple autonomous whirlpool example to make the association between trajectories and fluid flow more obvious, but in fact all dynamical systems can be interpreted in this way through the *flow map* defined below.

1.1.2 Flow Maps

Definition: The *flow map* $\phi_{t_0}^t(\mathbf{x})$ with initial time t_0 is given by

$$\phi_{t_0}^t : D \to D \tag{5}$$
$$: \mathbf{x}_0 \mapsto \mathbf{x}(t; t_0, \mathbf{x}_0) \tag{6}$$

$\phi_{t_0}^t$ takes all initial conditions in D and computes their integrated positions at time t. It is referred to as a flow map because it mimics the effect of a fluid flow carrying particles as described above. There are a couple of properties worth noting:

- $\phi_{t_0}^{t_0}(\mathbf{x}) = \mathbf{x}_0$
- $\phi_{t_0}^{t+s}(\mathbf{x}) = \phi_s^{t+s}(\phi_{t_0}^s(\mathbf{x})) = \phi_t^{t+s}(\phi_{t_0}^t(\mathbf{x}))$

This definition will be useful for analyzing trajectory divergence in §2.1.1.

1.2 Stable and Unstable Manifolds in Autonomous 2D Flows

The following time-independent example will aid in motivating the search for global structures in more general time-dependent systems.

1.2.1 Example: Simple Pendulum

A simple undamped planar pendulum of mass m and length l has a single degree of freedom θ (see Fig. 2) that satisfies

$$\ddot{\theta} = -\frac{g}{l}\sin(\theta) \tag{7}$$

This equation can be recast in the form (1)-(2). Setting $\dot{\theta} = \xi$, (7) becomes

$$\dot{\theta} = \xi \tag{8}$$
$$\dot{\xi} = -\frac{g}{l}\sin(\theta) \tag{9}$$

The associated phase plane (with some representative trajectories plotted) is given in Fig. 3.

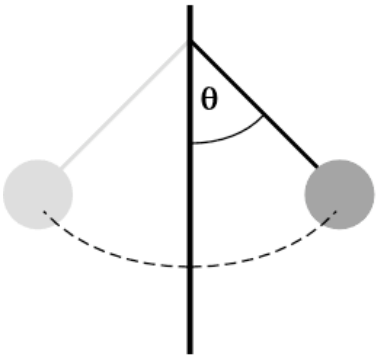

Figure 2: Simple pendulum, $\theta = 0$ defined to be the downward normal.

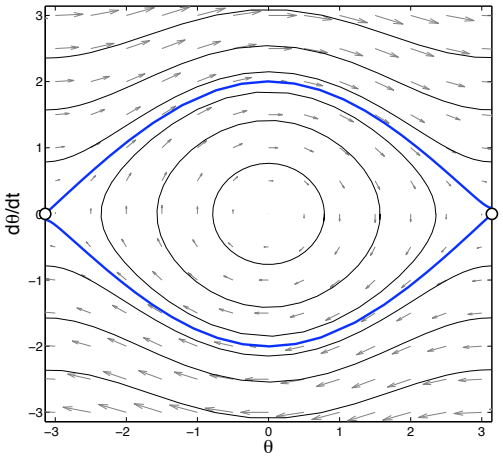

Figure 3: (Flattened) phase plane for the simple pendulum.

To agree with the physical system, the line $(-\pi, \xi)$ is henceforth identified with the line (π, ξ) so that the phase plane actually forms a cylinder.[2] The system now has just two distinct fixed points: $(\theta, \xi) = (0, 0)$ (center) and $(\theta, \xi) = (-\pi, 0)$ (saddle). Two trajectories (in blue) asymptotically approach the fixed point at $(-\pi, 0)$ as $t \to \pm\infty$. Before discussing their significance, it is useful to introduce the concepts of stable and unstable manifolds in autonomous systems.

[2] See [2] for a diagrammatical illustration of this identification.

Definition: The *stable manifold* $W^s(\mathbf{x}^*) \subset D$ of a hyperbolic fixed point $\mathbf{x}^*$ is given by

$$W^s(\mathbf{x}^*) = \left\{ \mathbf{x}_0 \in D : \lim_{t \to \infty} \mathbf{x}(t; \mathbf{x}_0) = \mathbf{x}^* \right\} \tag{10}$$

Definition: The *unstable manifold* $W^u(\mathbf{x}^*) \subset D$ of a hyperbolic fixed point $\mathbf{x}^*$ is given by

$$W^u(\mathbf{x}^*) = \left\{ \mathbf{x}_0 \in D : \lim_{t \to -\infty} \mathbf{x}(t; \mathbf{x}_0) = \mathbf{x}^* \right\} \tag{11}$$

Here the notation $\mathbf{x}(t; \mathbf{x}_0)$ has been used to emphasize that the solution trajectories $\{\mathbf{x}(t) : t \in [0, \infty)\}$ and $\{\mathbf{x}(t) : t \in (-\infty, 0]\}$ depend upon the initial state $\mathbf{x}(0) = \mathbf{x}_0$. Both definitions could equally well have been given in terms of the flow map ϕ.

Identifying the stable and unstable manifolds of saddle points is particularly useful because they divide the phase plane into regions of distinct asymptotic behavior. In the simple pendulum example above, $(-\pi, 0)$ is a saddle point with $W^s((-\pi, 0)) = W^u((-\pi, 0))$ given by the blue trajectories. These curves separate the physical motion of the pendulum into two types: oscillation ($\dot{\theta}$ changes sign in $(-\pi, \pi)$) and spinning ($\dot{\theta}$ is one-signed in $(-\pi, \pi)$). Such curves are also called *separatrices*, and in autonomous systems like this one their position in the phase plane is fixed for all t.

In non-autonomous 2D systems, fixed points are much less likely to exist due to the explicit time dependence of $\mathbf{f}$.[3] Even if a non-autonomous system does have fixed points, many of the analytical tools developed for autonomous systems (such as stable and unstable manifolds) cease to be clearly defined, or even meaningful, in the non-autonomous case. Lagrangian Coherent Structures (LCSs) are a generalization of the separatrices seen in §1.2, and provide one way of uncovering global behavior and trends in more complex non-autonomous flows.

2 Lagrangian Coherent Structures

2.1 Precursor to the Formal Definition

To help identify the non-autonomous analog of stable and unstable manifolds in autonomous systems, consider the system (12)-(13) below.

$$\dot{x} = x - y \tag{12}$$
$$\dot{y} = -(2x + y) \tag{13}$$

This system has a saddle point at $\mathbf{x}^* = (0, 0)$, and as before the stable and unstable manifolds of $\mathbf{x}^*$ have been plotted in blue in Fig. 4. In Fig. 5, the trajectories $\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$ corresponding to the initial states $\mathbf{x}_1(0) = (0.66, 2.0)$ and $\mathbf{x}_2(0) = (0.80, 2.0)$ have been plotted for $t \in [0, 1.8]$. Despite the fact that $\mathbf{x}_1(0)$ and $\mathbf{x}_2(0)$ are relatively close, the trajectories $\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$ diverge as t increases. This is an illustration of the general result that trajectories initially on opposite sides of a stable manifold will diverge as $t \to +\infty$.[4] The extent of this divergence depends upon the structure of the phase portrait: in the simple system above, $\|\mathbf{x}_1(t) - \mathbf{x}_2(t)\| \to \infty$ as $t \to \infty$.

[3] Note that making the 2D system $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), t)$ autonomous by introducing the variable $x_3 = t$ and setting $\dot{x}_3 = 1$ will give a 3D autonomous system that still has no fixed points.

[4] The analagous statement for the unstable manifold is: trajectories initially on opposite sides of an unstable manifold will diverge as $t \to -\infty$.

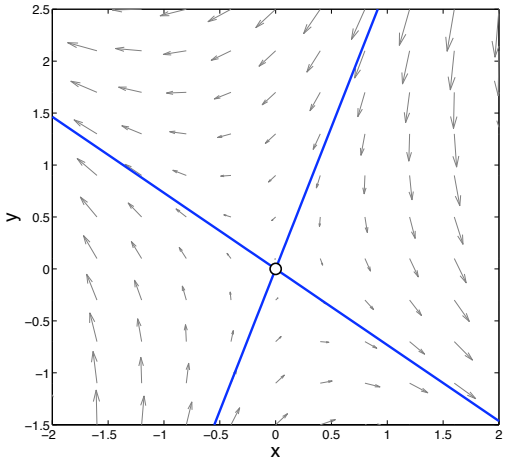

Figure 4: Stable and unstable manifolds for (12)-(13).

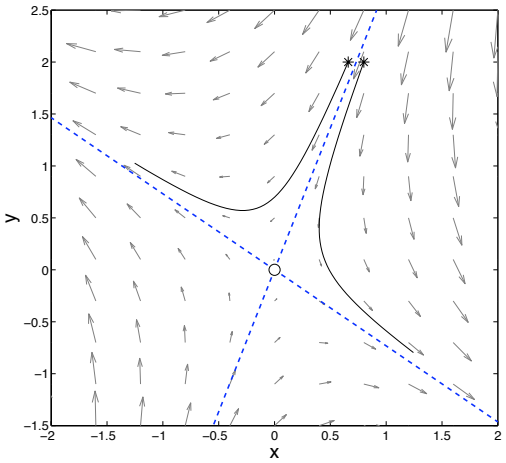

Figure 5: Trajectories starting on opposite sides of the stable manifold diverge as t increases.

The idea of the 'eventual divergence' of trajectories separated by stable or unstable manifolds (as $t \to \pm\infty$) can be used to search for structures with similar properties in non-autonomous flows. In autonomous systems, the *Lyapunov exponent* provides a quantitative measure of how rapidly two close initial points separate as $t \to \infty$. For non-autonomous systems, an asymptotic or instantaneous approximation is of limited use; instead a finite-time measure that depends explictly upon t is required.

2.1.1 Deriving the Finite-Time Lyapunov Exponent (FTLE)

Give an arbitrary point $\mathbf{x}(t_0) \in D$, let $\mathbf{y}(t) = \mathbf{x}(t) + \delta\mathbf{x}(t)$ for all t, where $\|\delta\mathbf{x}(t_0)\|_2$ is assumed to be less than $\eta > 0$ small. The aim of the FTLE is to estimate the maximal growth of $\|\delta\mathbf{x}(t)\|_2$ during the fixed time period $[t_0, t_0 + T]$.

Using the ideas developed in §1.1.2, write

$$\begin{aligned} \delta\mathbf{x}(t_0+T) &= \phi_{t_0}^{t_0+T}(\mathbf{y}) - \phi_{t_0}^{t_0+T}(\mathbf{x}) && (14)\\ &= \phi_{t_0}^{t_0+T}(\mathbf{x}+\delta\mathbf{x}) - \phi_{t_0}^{t_0+T}(\mathbf{x}) && (15)\\ &= \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}\delta\mathbf{x}(t_0) + \mathcal{O}\left(\|\delta\mathbf{x}(t_0)\|_2^2\right) && (16) \end{aligned}$$

Since $\|\delta\mathbf{x}(t_0)\|_2$ is assumed to be small, the $\mathcal{O}\left(\|\delta\mathbf{x}(t_0)\|_2^2\right)$ term is negligible.[5] Thus, taking the L_2 norm of both sides gives

$$\|\delta\mathbf{x}(t_0+T)\|_2 = \left\|\frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}\delta\mathbf{x}(t_0)\right\|_2 \quad (17)$$

So far, nothing has been said about the orientation of the perturbation vector $\delta\mathbf{x}(t)$ in the (x, y) plane. Since we're aiming to quantify the maximal separation of trajectories initially close to $\mathbf{x}$,

[5] In fact, neglecting this term requires that $\|\delta\mathbf{x}(t)\|_2$ remains small for all $t \in (t_0, t_0 + T]$.

the next step is to maximize the expression (17) over all small initial perturbations $\delta\mathbf{x}(t_0)$:

$$\max_{0<\|\delta\mathbf{x}(t_0)\|_2\leq\eta} \|\delta\mathbf{x}(t_0+T)\|_2 \quad = \quad \max_{0<\|\delta\mathbf{x}(t_0)\|_2\leq\eta} \left\| \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}} \delta\mathbf{x}(t_0) \right\|_2 \tag{18}$$

For convenience in the remaining calculations, introduce the notation

$$J = J(\mathbf{x}, t_0, T) = \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}} \tag{19}$$

so that equation (18) becomes

$$\begin{aligned} \max_{0<\|\delta\mathbf{x}(t_0)\|_2\leq\eta} \|\delta\mathbf{x}(t_0+T)\|_2 &= \max_{0<\|\delta\mathbf{x}(t_0)\|_2\leq\eta} \|J\delta\mathbf{x}(t_0)\|_2 & (20)\\ &= \max_{0<\|\delta\mathbf{x}(t_0)\|_2\leq\eta} \left\{ \left(\frac{\|J\delta\mathbf{x}(t_0)\|_2}{\|\delta\mathbf{x}(t_0)\|_2} \right) \|\delta\mathbf{x}(t_0)\|_2 \right\} & (21)\\ &= \eta \|J\|_2 & (22) \end{aligned}$$

Therefore, the maximum divergence of trajectories near a point $\mathbf{x}$ is governed by $\|J(\mathbf{x}, t_0, T)\|_2$. The FTLE incorporates this measure of divergence through the following definition:

Definition: The *Finite-Time Lyapunov Exponent* (FTLE) $\sigma_{t_0}^T(\mathbf{x})$ at a point $\mathbf{x} \in D$ at time t_0 with finite integration time T is given by

$$\sigma_{t_0}^T(\mathbf{x}) = \frac{1}{|T|} \ln \left\| \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}} \right\|_2 \tag{23}$$

It follows that

$$\max_{0<\|\delta\mathbf{x}(t_0)\|_2\leq\eta} \|\delta\mathbf{x}(t_0+T)\|_2 \quad = \quad \eta e^{\sigma_{t_0}^T(\mathbf{x})|T|} \tag{24}$$

For practical purposes note that $\|J\|_2 \equiv \sqrt{\lambda_{\max}(J^*J)}$ is just the *spectral norm* of J. The reader may be wondering why the FTLE has been defined to introduce the 'artificial' exponential in equation (24). Though the maximal growth factor for perturbations is shown in (22) to simply be $\|J\|_2$, the choice of scaling in (23) will better allow the regions where trajectories separate most rapidly to be identified.

For T and t_0 fixed, $\sigma_{t_0}^T(\mathbf{x})$ is a scalar field defined on D, referred to as the FTLE *field*. If we allow t_0 to vary, then $\sigma_{t_0}^T(\mathbf{x})$ becomes a time-varying scalar field that measures the finite-time expansion of nearby trajectories throughout D. This interpretation is a little dense, so let's look at an example of the FTLE field for an autonomous system.

2.1.2 Example: A Numerical FTLE Field

Recall the simple pendulum example from §1.2.1. Fig. 6 shows the FTLE field for (8)-(9), with $T = 20$. Since this particular system is autonomous, the FTLE field varies only with T, and not with t_0. As one would expect (see §2.1), the FTLE is largest near the stable manifolds of the fixed point $(-\pi, 0)$ seen in Fig. 3 - so the FTLE field allows the usual stable and unstable manifolds to be identified in autonomous systems. §4 contains some comments on generating the FTLE field in Fig. 6.

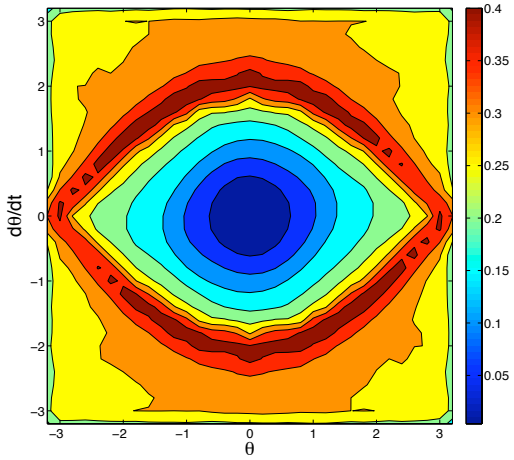

Figure 6: Numerical FTLE field for the simple pendulum (8)-(9); $T = 20$.

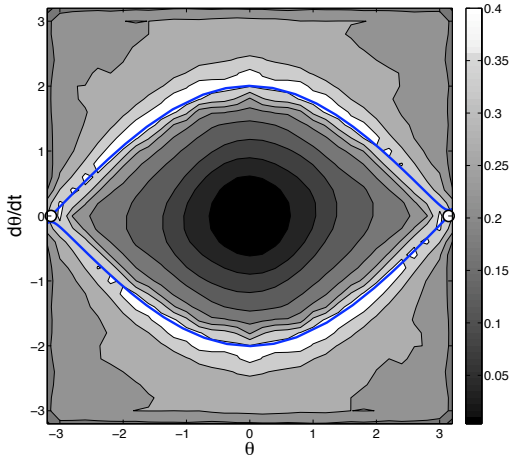

Figure 7: The ridges of high values in the FTLE field correspond to the manifolds of $(-\pi, 0)$ seen in Fig. 3.

This example indicates that the structures we are interested in finding in more general, non-autonomous systems - the analogs of stable and unstable manifolds in autonomous systems - ought to be characterized by ridges of high values in the FTLE field. These ridges, referred to as Lagrangian Coherent Structures, can be defined using the Hessian H of the FTLE field given by

$$H = H(\mathbf{x}, t_0, T) = \frac{d^2 \sigma_{t_0}^T(\mathbf{x})}{d\mathbf{x}^2} \tag{25}$$

Now is a good time to introduce the formal definition of a Lagrangian Coherent Structure as given by [3].

2.2 Formal Definition of a Lagrangian Coherent Structure in Two Dimensions

Definition: A *Lagrangian Coherent Structure* (LCS) of a dynamical system at time t_0 with integration time T is an injective curve $\mathbf{c} = \mathbf{c}(s) \subset D \subset \mathbb{R}^2$ parameterized by $s \in (a, b) \subset \mathbb{R}$ such that:

- $\mathbf{c}'(s)$ and $\nabla \sigma_{t_0}^T(\mathbf{c}(s))$ are parallel $\forall s \in (a, b)$;
- $\hat{\mathbf{n}}^{tr} H \hat{\mathbf{n}} = \min_{\|u\|=1} \mathbf{u}^{tr} H \mathbf{u} < 0$, where $\hat{\mathbf{n}} = \hat{\mathbf{n}}(s)$ is the unit normal vector to $\mathbf{c}(s)$ for all s.

This is a complicated looking definition, but does exactly what we previously described - identifies the ridges of high FTLE values at time t_0. The first condition identifies the direction of $\mathbf{c}$ for all s, and the second ensures that the ridge is a peak (and not a trough) of the FTLE field. Before continuing to look at some examples, there are several points worth noting:

- If $\mathbf{f}$ is a smooth function of $\mathbf{x}$ and t, the FTLE field is smooth and therefore so are any associated LCSs.
- LCSs generally vary with time in non-autonomous systems. We will use the notation $\mathbf{c}(s) = \mathbf{c}(s, t)$ when it is desirable to emphasize this time-dependence.
- A FTLE field may have several LCSs (as in §2.1.2 where two LCSs exist).

- A LCS is not necessarily a level curve of $\sigma_{t_0}^T(\mathbf{x})$.
- A LCS ridge is not necessarily a global maximum.
- The formal definition of a LCS gives no indication as to why the 'Lagrangian' part of 'Lagrangian Coherent Structure' is included. It refers to the Lagrangian perspective of fluid flows because a LCS is approximately advected by the flow map ϕ. This is due to the definition of the FTLE, which incorporates integrated particle trajectories.
- In autonomous systems, LCSs are independent of t_0 and T, and correspond to the stable and unstable manifolds of hyperbolic fixed points (as in §2.1.2).
- The best choice for $|T|$ is problem dependent. Ideally $|T|$ should be chosen large enough that ridges in the FTLE field become clearly defined, but not so large that the linearization applied to (16) is invalidated (this second problem arises when numerically computing the FTLE field with a finite grid spacing). Computation time also plays a role. The sign of T controls whether attracting ($T < 0$) or repelling ($T > 0$) LCSs are found.

As an illustration of finding LCSs analytically, consider the instantaneous analytic FTLE field given by

$$\sigma_{t_0}^T(\mathbf{x}) = \frac{3x^2 - x - 2}{1 + y^2}, \quad x, y \in [-2, 2]. \tag{26}$$

To satisfy the first condition in the definition above, we require that $\mathbf{c}'(s)$ and $\nabla\sigma_{t_0}^T(\mathbf{c}(s))$ are parallel $\forall s \in (a, b)$. Noting that

$$\nabla\sigma_{t_0}^T(\mathbf{x}) = \left(\frac{6x - 1}{1 + y^2}\ ,\ -\frac{2y(3x^2 - x - 2)}{(1 + y^2)^2}\right) \tag{27}$$

it is clear that choosing $\mathbf{c}(s) = \{(s, 0) : s \in [-2, 2]\}$ gives $\nabla\sigma_{t_0}^T(\mathbf{c}(s)) = (6s - 1, 0) \parallel \mathbf{c}(s)$ in $[-2, 2]$ as required. It remains to check where the second condition holds. Along $\mathbf{c}(s)$, the normal vector $\hat{\mathbf{n}}(s) = (0, \pm 1)$ for all $s \in [-2, 2]$. This implies that $\hat{\mathbf{n}}^{tr} H \hat{\mathbf{n}} = \partial_y^2[\sigma_{t_0}^T((s, 0))]$ when H is evaluated at $(s, 0)$. Thus

$$\hat{\mathbf{n}}^{tr} H \hat{\mathbf{n}} = -2(3s^2 - s - 2) \tag{28}$$

The expression above satisfies the second condition of the definition of an LCS just when $(3s^2 - s - 2) > 0$ i.e. when $s \in [-2, -\frac{2}{3}) \cup (1, 2]$. To see this, write $\mathbf{u} = (a, \sqrt{1 - a^2})$. We then want to minimize $\mathbf{u}^{tr} H \mathbf{u}$ over a. Some algebra gives that

$$\mathbf{u}^{tr} H \mathbf{u} = 2a^2(3s^2 - s + 1) - 2(3s^2 - s - 2) \tag{29}$$

Since $(3s^2 - s + 1) > 0$ for $s \in [-2, 2]$, we see that this function attains a negative minimum just when $a = 0$ (i.e. $\mathbf{u} = (0, 1) = \hat{\mathbf{n}}$) and $(3s^2 - s - 2) > 0$. Therefore this field contains two distinct LCSs: $\{(x, 0) : x \in [-2, -\frac{2}{3})\}$ and $\{(x, 0) : x \in (1, 2]\}$. The FTLE field and LCSs are plotted in Fig.s 8 and 9, and as expected the LCSs correspond to the 'high ridges' of the FTLE field.

In most problems the FTLE field is not known analytically, so numerical techniques must be used to estimate the locations of LCSs (as in Fig. 7, where the FTLE field allowed us to approximately locate the stable manifolds).

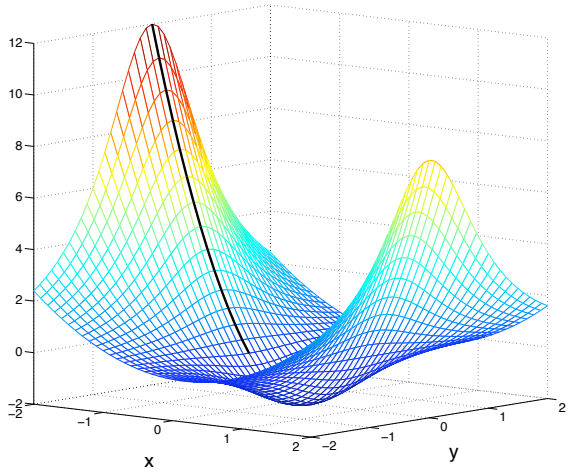

Figure 8: FTLE field (26) with the first LCS $\{(x,0) : x \in [-2, -\frac{2}{3})\}$ marked in black.

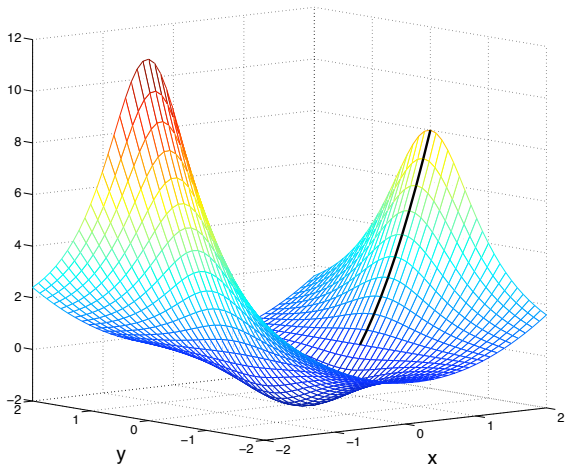

Figure 9: FTLE field (26) with the second LCS $\{(x,0) : x \in (1,2]\}$ marked in black.

2.3 The Flux Across a Lagrangian Coherent Structure

Before seeing a LCS in action in a non-autonomous system, it is interesting to consider the flow close to $\mathbf{c}(s)$. Since a LCS is designed to separate flows exhibiting different behavior, we would hope that the flux across a LCS is at least small, if not zero. This is a reasonable expectation because the FTLE is derived from integrated trajectories, which are themselves material curves. The following argument quantitatively estimates the flux across a LCS, and allows us to find the conditions that guarantee it is small.

Define a function $L(\mathbf{x}, t)$ by

- $|L(\mathbf{x}, t)| = \|\mathbf{x} - \mathbf{x}_q(t)\|$, where $\mathbf{x}_q(t)$ is the closest point on $\mathbf{c}(s;t)$ to $\mathbf{x}$ at time t;
- $L(\mathbf{x}, t)(((\mathbf{x} - \mathbf{x}_q(t)) \times \mathbf{c}_s(s;t)) \cdot \hat{\mathbf{k}}) \geq 0$, where $\hat{k} = (0, 0, 1)$.

This definition gives the signed distance from any $\mathbf{x} \in D$ to the nearest point on the LCS. That is, when traveling along $\mathbf{c}(s)$ in the positive $\mathbf{c}'(s)$ direction, points to the right of the LCS correspond to a positive value of L, while points to the left of the LCS correspond to a negative value of L. The LCS itself is now described by the set $\{\mathbf{x} : L(\mathbf{x}, t) = 0\}$. We are interested in the total flux $\Phi(t)$ across the LCS, which is given precisely by

$$\Phi(t) = \int_a^b \left.\frac{dL}{dt}\right|_{L=0} ds \tag{30}$$

A more explicit form for $\left.\frac{dL}{dt}\right|_{L=0}$ is derived (lengthily) in [3] - it is found that

$$\left.\frac{dL}{dt}\right|_{L=0} = \frac{\|\nabla\sigma\|}{(\hat{\mathbf{n}}^{tr} H \hat{\mathbf{n}})} \left\langle \hat{\mathbf{t}}, \frac{\partial \hat{\mathbf{n}}}{\partial t} - J\hat{\mathbf{n}} \right\rangle + \mathcal{O}\left(\frac{1}{|T|}\right) \tag{31}$$

where $\hat{\mathbf{t}}$ is a unit vector tangent to $\mathbf{c}(s)$, and all terms on the right hand side are evaluated on $\mathbf{c}(s)$. The first term on the right hand side is small for well-defined (sharp) LCSs, and is zero for time-independent flows. In fact, the inner product measures the difference between the local rotation rates of the velocity field and the LCS. To see this, consider each of the terms $\langle \hat{\mathbf{t}}, \frac{\partial \hat{\mathbf{n}}}{\partial t} \rangle$ and $\langle \hat{\mathbf{t}}, J\hat{\mathbf{n}} \rangle$ separately:

- $\langle \hat{\mathbf{t}}, \frac{\partial \hat{\mathbf{n}}}{\partial t} \rangle$: Write $\hat{\mathbf{n}} = (\cos\theta, \sin\theta)$. Then $\hat{\mathbf{t}} = (-\sin\theta, \cos\theta)$. This gives

$$\left\langle \hat{\mathbf{t}}, \frac{\partial \hat{\mathbf{n}}}{\partial t} \right\rangle = (-\sin\theta, \cos\theta)(-\dot{\theta}\sin\theta, \dot{\theta}\cos\theta)^{tr} = \dot{\theta} \tag{32}$$

  Since θ parameterizes the normal vector $\hat{\mathbf{n}}$ to $\mathbf{c}(s)$, $\langle \hat{\mathbf{t}}, \frac{\partial \hat{\mathbf{n}}}{\partial t} \rangle = \dot{\theta}$ measures the local rotation rate of the LCS.

- $\langle \hat{\mathbf{t}}, J\hat{\mathbf{n}} \rangle$: this term represents the component of the linearized velocity field J applied to the normal vector $\hat{\mathbf{n}}$ that lies tangent to the LCS. Therefore, it measures how much the local Eulerian field rotates vectors that are initially perpendicular to the LCS.

So combining the above expressions in the term $\langle \hat{\mathbf{t}}, \frac{\partial \hat{\mathbf{n}}}{\partial t} - J\hat{\mathbf{n}} \rangle$ gives the difference between the local rotation rates of the velocity field and the LCS as claimed above. The second term on the right hand side is small for $|T|$ large, and hence we see that the total flux across the LCS is small for 'sharp' LCSs with $|T|$ large. This is another factor to consider when selecting the integration time T for particular problems.

2.4 Lagrangian Coherent Structures in a Non-Autonomous System

So far, we have verified that LCSs are consistent with the stable and unstable manifolds of autonomous systems. However, the real objective of the definition of a LCS is to allow these concepts to be generalized to non-autonomous flows. We now introduce the forced pendulum governed by the equation

$$\ddot{\theta}(t) = (2.5\cos 5t - 1)\sin\theta(t) \tag{33}$$

The forcing term $2.5\cos(5t)\sin\theta(t)$ corresponds to vertical oscillations of the suspension point of the pendulum, with period $\frac{2\pi}{5}$s and amplitude 2.5m. We will use LCSs to determine the dynamics of this pendulum, and to compare the behavior of this system with the unforced pendulum of §1.2.1.

Fig. 10 shows the FTLE field at $t = 0$ for (33) with $T = 10$. It is immediately clear that the phase space for the forced pendulum is markedly more complicated than that of the unforced pendulum. Since the phase space is 2π-periodic in θ, the FTLE field shows three of the 'eye' shaped outlines similar to the LCSs observed in Fig. 6. As $T > 0$ for this field, trajectories of the forced pendulum diverge away from the LCSs as t increases.

What can be deduced about orbits of this system based on the position of the LCSs? Using our previous analysis of the LCSs for the unforced pendulum, it seems reasonable to expect that the pendulum exhibits normal oscillatory motion for initial conditions 'inside' the central eye. Inside the other eyes the pendulum does not oscillate in quite the same manner, since $\dot{\theta}$ is one-signed in these regions. Instead the pendulum motion will consist of a base spinning rate, plus a periodic perturbation in θ (small compared to the rate of spinning). The final, most interesting region to consider is the space 'in between' the three eyes. Orbits that begin in this region are constrained to remain there for all future times (since the LCSs partion the phase space in this case). However, the FTLE is clearly positive in this region, which indicates that close initial conditions will diverge from one another. This combination of divergent trajectories and a completely bounded subset of the phase space indicates that in this region we expect to observe *mixing*, an important phenomenon in real fluid flows.

Having deduced this behavior, we should now verify it. Fig. 11 is the *time-*$\frac{2\pi}{5}$ *return map* associated with (33). A map of this type allows us to visualize the behavior of a 3-dimensional dynamical

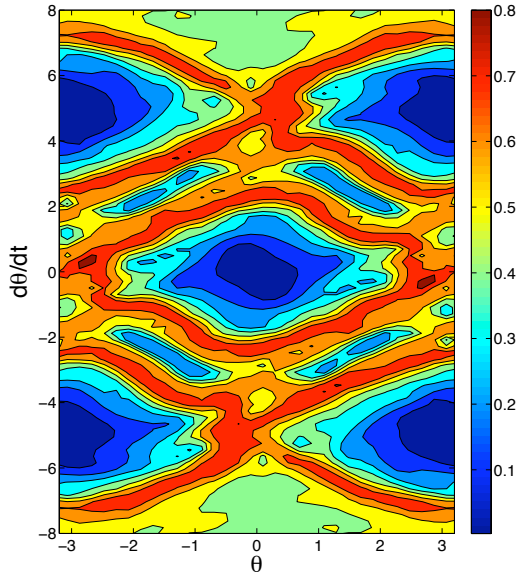

Figure 10: FTLE field for (33). The approximate locations of the LCSs are given by the red curves (high ridges) in the FTLE field.

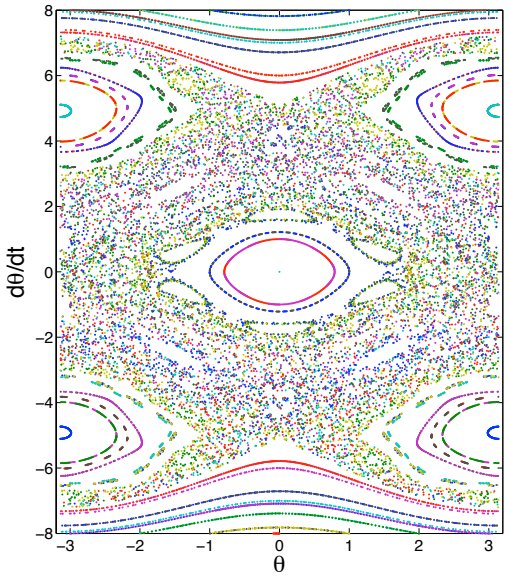

Figure 11: The time-$\frac{2\pi}{5}$ map for (33). 120 initial conditions were iterated 1000 times and plotted in various colors to emphasize trajectory mixing.

system with a 2-dimensional plot. In our case, the dimension we have effectively removed corresponds to time t. Hence, although the phase portrait varies with time, we may use Fig. 10 to analyze the full 3-dimensional flow. Comparing Fig. 11 and Fig. 10, it is clear that the behavior we predicted from the LCSs matches the orbits shown in the return map. In particular, the mixing of trajectories outside the LCSs is illustrated by the intermingling of colors in the return map. Such behavior would not necessarily be evident from studying the Eulerian velocity field alone.

2.5 Lagrangian Coherent Structures and Empirical Data

Since LCSs are essentially defined by the flow map ϕ using a finite integration time T, they are particularly suited to studying real-world fluid flows where asymptotic behavior is unimportant. In the previous sections, analytic autonomous systems have been used to illustrate the main concepts of both the FTLE and LCSs. However most applications of LCSs are drawn from real data, which complicates the calculation of the FTLE field (and hence the discovery of LCSs). For example, to generate Fig. 6, equations (8)-(9) were integrated from $t = t_0$ to $t = T$ numerically, but for a real data set recorded on a finite grid one does not have the luxury of simply integrating points forwards in time. Instead, the data collected on a finite grid is interpolated in order to better approximate the real fluid flow. This technique is used in all the examples presented below.

2.5.1 Ocean Flows and Pollution

The study of ocean transport is one area in which LCSs have proved extremely useful. VHF radar enables very high resolution recording of ocean surface currents, but finding the interesting patterns and trends hidden within such large data sets is not easy. Lekien et al [4] applied the ideas presented in this paper to VHF data from the Florida coast, and identified a LCS in the surface currents (see Fig. 12).

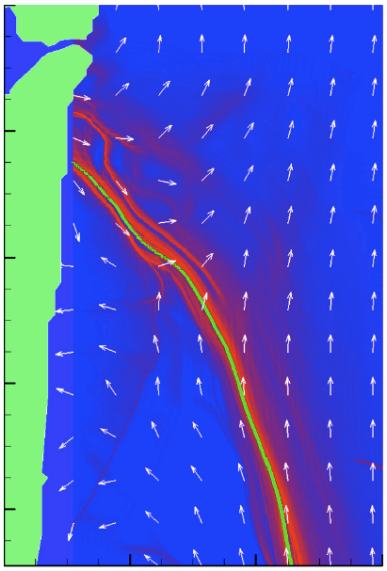

Figure 12: Discrete velocity field (white arrows), FTLE field (blue-low to red-high) and LCS (green) for Florida coast. Image from [3].

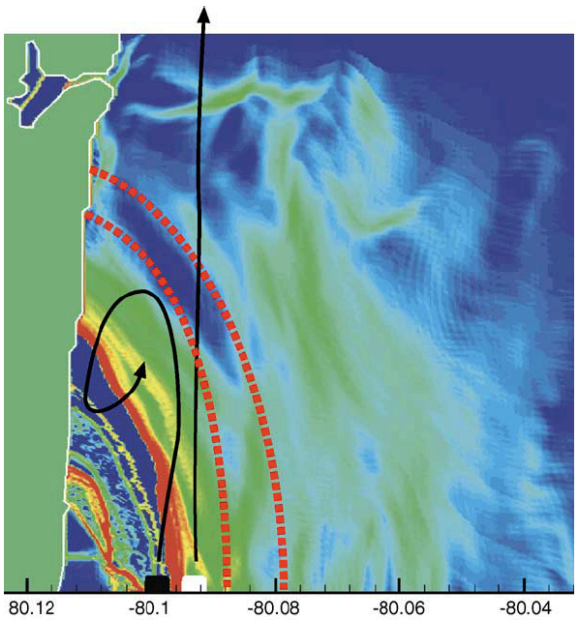

Figure 13: Alternative view of the FTLE field, with future positions of the LCS plotted as dashed lines. While the white particle escapes the coastline, the black particle will recirculate. Image from [4].

The LCS was found to remain attached to the coastline at all times. However, its point of attachment to the coast oscillated north and south approximately periodically. This result is particularly interesting when the behavior of the ocean flow on either side of the LCS is considered in Fig. 13. Surface-bound articles that begin 'outside' the LCS (relative to the coastline) are advected away from the coast as time progresses. However, those particles that begin inside the LCS recirculate and remain trapped in this region. The LCS has done exactly what we would hope - it has allowed us to identify a time-dependent curve that divides regions of qualitatively different behavior in an extremely large empirical data set.

The discovery of this LCS is particularly important for pollution control. Pollution deposited into the sea inside the LCS will recirculate near the coast of Florida, while pollution deposited outside the LCS will swiftly move out to sea. Lekien et al used the computed, time-varying LCS to design an optimal pollution strategy for a hypothetical pollution-dumping plant on the Florida coastline - that is, a real-time strategy that would minimize the resultant coastal pollution in Florida. The problem is more complicated than it might appear, as it is necessary to predict the position of the LCS up to eight hours in the future in order to release as much pollution as possible outside the LCS. However, it is shown that such a real-time system is feasible and can efficiently reduce the impact of a pollution source.

2.5.2 Other Applications

Even though the definition of a LCS was only formalized in 2005 [3], similar structures have been observed and studied in geophysical flows for almost 30 years. A brief summary of other uses is presented below.

- **Jellyfish Swimming and Feeding** - jellyfish swim by contracting and relaxing to generate vortex rings that produce thrust. They feed in a similar way, using vortices to direct prey and nutrients towards their tentacles. Shadden et al [5] aimed to study the fluid transport around a jellyfish to uncover any potentially useful techniques for making transport more effective. The FTLE field derived from fluid velocity measurements was computed (with $T = 13.3$s), and is shown in Fig. 14(a). The associated LCS, highlighted in Fig. 14(b), shows that a region of recirculating water is carried behind the jellyfish in contact with the subumbrellar region (the concave underside of the jellyfish body).

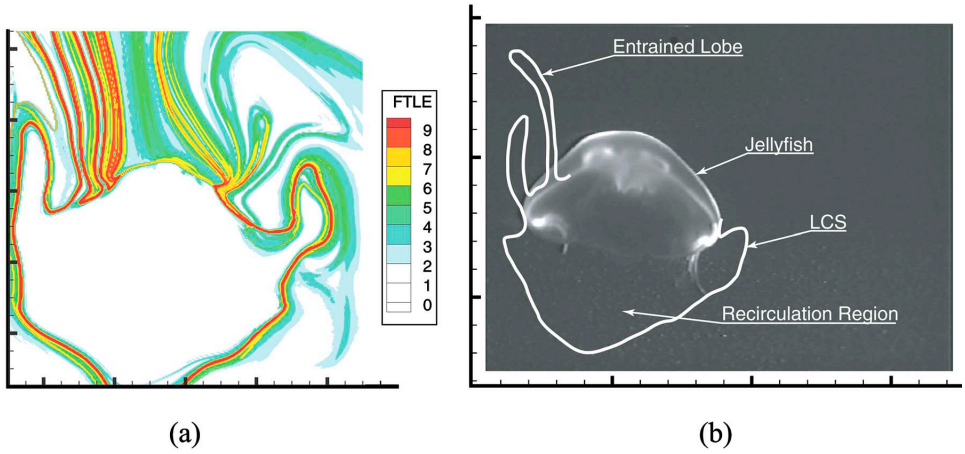

Figure 14: FTLE field and LCS for Aurelia Aurita jellyfish. Image from [5].

- **Separation over an Airfoil** - LCSs can be used to identify the *separation profile* of airflow over an airfoil. A separation profile corresponds to the streamline whose tangent is normal to the airfoil surface. In [3], the separation profile (LCS) is used to analyze the effectiveness of flow control, particularly in controlling the attachment point of the separation profile and hence the lift and drag of the wing.

3 Conclusions

We have seen that Lagrangian Coherent Structures are consistent extensions of stable and unstable manifolds to general dynamical systems. In particular, since their definition utilizes a finite time interval, they are well suited to studying a wide range of real fluid flows, especially those where complicated transport or mixing processes may not be obvious when studying the raw velocity data. The theory is also an area of active research - the authors of [3] have submitted an extension of the 2-dimensional results presented in [3] and referenced throughout this paper to N-dimensional systems, due for publication in 2008.

4 Appendix

The FTLE field in Fig. 6 (and others throughout this paper) was generated by direct numerical integration of the equations (8)-(9) using Matlab's adaptive Runge-Kutta routine `ode45`. Analytically one would expect the FTLE field to be smooth and symmetric in both θ and $\dot{\theta}$, but in Fig. 6 this is not the case - a result of truncation errors introduced by the numerical integration. This plot of the FTLE field includes some mild smoothing to counteract the numerical inaccuracies.

Computationally, these calculations are expensive even for moderately sized discretizations - for example, when using an n by n grid of initial conditions n^2 integrations must be performed. For $n = 35$ and $T = 20$ as in the example code below, the FTLE field took several minutes to compute. Also note that the Jacobian is computed using a centered-difference approximation to the derivative (lines 22-25) - another potential source of error. For systems where the governing equations are known (as in the pendulum examples above), one could find J by explicitly solving variational equations, which would improve the accuracy of the final numerical FTLE field.

```
1.   % initialize integration time T, f(x,t), discretization size n ----------------
2.   T = 20;
3.   f_x_t = inline('[v(2);-sin(v(1))]','t','v');
4.   grid_min = -3.4;
5.   grid_max = 3.4;
6.   grid_width = grid_max-grid_min;
7.   n = 35;
8.   grid_spacing = grid_min:(grid_width/(n-1)):grid_max;
9.  % integrate all initial points for t in [0,T] --------------------------------
10.  for i = 1:n
11.      for j = 1:n
12.          [t,x] = ode45(f_x_t,[0,T],[grid_spacing(i),grid_spacing(j)]);
13.          % store advected positions as they would appear in (x,y) coords ------
14.          advected_x(n-j+1,i) = x(length(x(:,1)),1);
15.          advected_y(n-j+1,i) = x(length(x(:,2)),2);
16.      end
17.  end
18.  % at each point in interior of grid, store FTLE -----------------------------
19.  for i = 2:n-1
20.      for j = 2:n-1
21.          % compute Jacobian phi ---------------------------------------------
22.          phi(1,1) = (advected_x(i,j+1)-advected_x(i,j-1))/(2*grid_width/(n-1));
23.          phi(1,2) = (advected_x(i-1,j)-advected_x(i+1,j))/(2*grid_width/(n-1));
24.          phi(2,1) = (advected_y(i,j+1)-advected_y(i,j-1))/(2*grid_width/(n-1));
25.          phi(2,2) = (advected_y(i-1,j)-advected_y(i+1,j))/(2*grid_width/(n-1));
26.          % find max eigenvalue of phi'*phi ----------------------------------
27.          lambda_max = max(abs(eig(phi'*phi)));
28.          % store FTLE -------------------------------------------------------
29.          sigma(i,j) = log(lambda_max)/abs(T);
30.      end
31.  end
32.  % plot FTLE field ----------------------------------------------------------
33.  contourf(grid_spacing,grid_spacing,sigma);
34.  colorbar('location','EastOutside');
```

5 Acknowledgments

I would like to thank Prof. Flaschka for overseeing this project, as well as Prof. Lin for suggesting an excellent topic and providing useful feedback.

6 References

[1] H. Poincaré, '*Sur le Problème des Trois Corps et les Équations de la Dynamique*', Acta Mathematica, 13, pp. 1-270, 1890.

[2] S. Strogatz, '*Nonlinear Dynamics and Chaos*', Westview Press, 2000.

[3] S. Shadden, F. Lekien and J. Marsden, '*Definition and Properties of Lagrangian Coherent Structures from Finite-Time Lyapunov Exponents in Two-Dimensional Aperiodic Flows*', Physica D, 212(3-4), pp. 271-304, 2005.

[4] F. Lekien, C. Coulliette, A. Mariano, E. Ryan, L. Shay, G. Haller and J. Marsden, '*Pollution Release Tied to Invariant Manifolds: A Case Study for the Coast of Florida*', Physica D, 210(1-2), pp. 1-20, 2005.

[5] S. Shadden, J. Dabiri and J. Marsden, '*Lagrangian Analysis of Fluid Transport in Empirical Vortex Ring Flows*', Phys. Fluids, 18(4), pp. 047105-047105-11, 2006.

Other Resources:

- `http://monet.unibas.ch/elmer/pendulum/spend.htm` - a damped, forced pendulum java applet.
- `http://www.cds.caltech.edu/~shawn/LCS-tutorial/` - a brief, less technical overview of Lagrangrian Coherent Structures.