

Detection of Lagrangian coherent structures in three-dimensional turbulence

By M. A. GREEN¹, C. W. ROWLEY¹ AND G. HALLER²

¹Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

²Department of Mechanical Engineering, Massachusetts Institute of Technology,
Cambridge, MA 02139, USA

(Received 14 June 2006 and in revised form 20 October 2006)

We use direct Lyapunov exponents (DLE) to identify Lagrangian coherent structures in two different three-dimensional flows, including a single isolated hairpin vortex, and a fully developed turbulent flow. These results are compared with commonly used Eulerian criteria for coherent vortices. We find that, despite additional computational cost, the DLE method has several advantages over Eulerian methods, including greater detail and the ability to define structure boundaries without relying on a preselected threshold. As a further advantage, the DLE method requires no velocity derivatives, which are often too noisy to be useful in the study of a turbulent flow. We study the evolution of a single hairpin vortex into a packet of similar structures, and show that the birth of a secondary vortex corresponds to a loss of hyperbolicity of the Lagrangian coherent structures.

1. Introduction

In this paper, we use direct Lyapunov exponents (DLE) to identify Lagrangian coherent structures in three-dimensional turbulent flows. Previous work on flow structure identification has been primarily *Eulerian*, i.e. it has been concerned with the spatial structure of quantities derived from the instantaneous velocity field and its gradient. The resulting Eulerian coherent structure criteria have been broadly used in flow structure identification, although none has emerged as a definitive tool of choice.

By contrast, *Lagrangian* methods identify flow structures based on the properties of fluid particle trajectories. An immediate advantage of these methods is their *objectivity*: they remain invariant with respect to rotation of the reference frame. A further advantage is their insensitivity to short-term anomalies in the velocity field. Computing Lagrangian quantities, however, can be computationally expensive. For this reason, the present study is partly aimed at weighing the efficacy of Lagrangian structure identification against its computational cost.

1.1. Eulerian methods

Eulerian coherent structure criteria are typically formulated in terms of the invariants of the velocity gradient tensor $\nabla\mathbf{u}$. The criteria to be discussed here include the Q -criterion and the swirling strength criterion.

The Q -criterion, developed by Hunt, Wray & Moin (1988), locates regions where rotation dominates strain in the flow. Letting \mathbf{S} and $\mathbf{\Omega}$ denote the symmetric and antisymmetric parts of $\nabla\mathbf{u}$, one defines Q as the second invariant of $\nabla\mathbf{u}$, given for incompressible flow by

$$Q = \frac{1}{2}(\|\mathbf{\Omega}\|^2 - \|\mathbf{S}\|^2), \quad (1.1)$$

where $\|\cdot\|$ is the Euclidean (or Frobenius) matrix norm. A coherent vortex is defined as a region where $Q > 0$.

The swirling strength criterion, employed by Zhou *et al.* (1999), seeks flow structures in regions where $\nabla\mathbf{u}$ has a complex pair of eigenvalues, which indicates locally spiralling streamlines. The swirling strength λ_{ci}^2 is then defined as the squared magnitude of the imaginary part of the complex eigenvalues. Coherent vortices are defined as areas where the swirling strength is greater than some positive threshold.

Other Eulerian criteria have also been used for structure identification, and some of these have been compared to Lagrangian criteria in Haller (2005). These include the Δ -criterion (Chong, Perry & Cantwell 1990), which defines the vortex as those regions where $\nabla\mathbf{u}$ has complex eigenvalues, and the λ_2 -criterion (Jeong & Hussein 1995), which identifies pressure minima within two-dimensional subspaces. Additionally, Chakraborty, Balachandar & Adrian (2005) proposed using the ratio of the real and imaginary parts of the complex eigenvalues of $\nabla\mathbf{u}$ to refine the definition of a vortex core.

These criteria identify similar structures in most flows, but they share several disadvantages. In particular, though they are invariant with respect to Galilean transformations, they are not invariant to time-dependent rotations, and thus are not *objective* (frame-independent). Furthermore, in practice all the above Eulerian criteria require a user-defined threshold to indicate the regions where a structure exists. In addition, the boundaries of the structure depend on the selected threshold, lending subjectivity to the definition of the size or boundary of the structure.

1.2. A Lagrangian method

The Lagrangian criterion used in this study is the direct Lyapunov exponent (DLE) method (Haller 2001). At each point in space, this scalar is a measure of the rate of separation of neighbouring particle trajectories initialized near that point. More precisely, if $\mathbf{x}(t, \mathbf{x}_0, t_0)$ denotes the position of a particle at time t , that began at position \mathbf{x}_0 at time t_0 , one defines a coefficient of expansion σ_T as the square of the largest singular value of the deformation gradient $\partial\mathbf{x}(t_0 + T, \mathbf{x}_0, t_0)/\partial\mathbf{x}_0$:

$$\sigma_T(\mathbf{x}_0, t_0) = \lambda_{\max} \left(\left[\frac{\partial\mathbf{x}(t_0 + T, \mathbf{x}_0, t_0)}{\partial\mathbf{x}_0} \right]^T \left[\frac{\partial\mathbf{x}(t_0 + T, \mathbf{x}_0, t_0)}{\partial\mathbf{x}_0} \right] \right). \quad (1.2)$$

The DLE field is then defined as

$$DLE_T(\mathbf{x}_0, t_0) = \frac{1}{2T} \log \sigma_T(\mathbf{x}_0, t_0). \quad (1.3)$$

Since the maximum eigenvalue is used in the definition of σ_T , direction information is not retained in $DLE_T(\mathbf{x}_0, t_0)$. A point \mathbf{x}_0 will have a high DLE value if there is a great amount of expansion in one direction, even if there is compression in all other directions.

Regions of maximum material stretching generate local maximizing curves (*ridges*) for the DLE field. The converse is not true: local maxima of the DLE field may indicate either locally maximal stretching or locally maximal shear. Trajectories that stretch relative to each other when integrated in negative time converge in forward time. This is analogous to passive scalars in a fluid flow collecting in coherent structures obtained from flow visualization. Therefore, if the DLE is calculated by integrating trajectories in backward time ($T < 0$), ridges in the DLE field may indicate attracting material lines, or attracting Lagrangian coherent structures (*attracting LCS*) (Haller & Yuan 2000). Integrating trajectories in forward time ($T > 0$) may produce DLE ridges that mark the location of *repelling LCS*. In order to confirm that the




FIGURE 1. DLE plots of Hill's spherical vortex, with white showing maximum values: (a) short integration time, and (b) long integration time.

ridges are indeed hyperbolic LCS (instead of regions of maximal shear), one may calculate the strain rate normal to the ridge, as discussed further in §2.2. Quantitative criteria for defining ridges in two-dimensional DLE fields have been given by Shadden, Lekien & Marsden (2005).

These positive-time and negative-time LCS delineate the boundary between qualitatively different regions in the flow (Shadden, Dabiri & Marsden 2006). For instance, figure 1 shows the backward-time DLE for a cross-section of Hill's spherical vortex, for two different integration times. The boundary of the spherical vortex is clear. In figure 1(a), the complete boundary is not seen, but as the integration time is increased, the whole boundary emerges and the LCS becomes sharper and clearer. The integration time can be increased or decreased depending on the amount of detail desired from the calculation, but the location of the ridge indicating the boundary of the vortex does not change.

Past applications of the DLE as a structure identification tool in two dimensions include LCS in two-dimensional quasi-geostrophic turbulence (Haller & Yuan 2000), LCS near the stratospheric polar vortex (Koh & Legras 2002), LCS in freely decaying two-dimensional turbulence (Lapeyre 2002), and LCS in a magnetically forced two-dimensional conducting fluid experiment (Voth, Haller & Gollub 2002). More recently, Lekien & Leonard (2004) used the DLE method to find coherent structures in Monterey Bay based on radar data, and Shadden *et al.* (2006) employed DLE to identify the structure of a piston-generated vortex ring and also captured the vortex ring wake structure of a jellyfish from two-dimensional DPIV data.

The DLE from discrete data has been shown to be robust and relatively insensitive to imperfect velocity data as long as the errors remain small in a special time-weighted norm (Haller 2002). Three-dimensional DLE has been computed by Haller (2005) on two established analytic flow solutions. Here, we apply DLE to physically relevant turbulent flows in the interest of detailing structures until now investigated using only Eulerian methods.

In this study, we focus on finding attracting material structures (which correspond to structures seen using flow visualization in experiments), and accordingly all calculations of DLE fields shown here are for backward time.

2. Applications

2.1. Isolated hairpin vortex

We begin by studying a single hairpin vortex, a structure commonly found in turbulent wall-bounded flows (Theodorsen 1955). For this case, the Eulerian and Lagrangian criteria may be compared to an unsteady flow in which the structure location and




FIGURE 2. An evolved isolated hairpin vortex generated by linear stochastic estimation, plotted using 10 % maximum λ_{ci}^2 . (a) Initial condition, (b) $t^+ = 63$, and (c) $t^+ = 171$.

qualitative shape are known *a priori* from previous numerical results and experiments, such as those of Head & Bandyopadhyay (1981) and Smith & Walker (1991).

The method we use to extract a single hairpin vortex was introduced by Zhou *et al.* (1999). First, a direct numerical simulation (DNS) of a fully developed turbulent channel flow was performed in a domain periodic in the streamwise and spanwise directions, using the method of Kim, Moin & Moser (1987). The calculation used a spectral collocation method with Fourier modes in the streamwise and spanwise directions and with Chebyshev modes in the wall-normal direction, and a second-order Adams–Bashforth time march. Our data were validated by comparing the mean profiles, Reynolds stresses, log law, and r.m.s. velocity fluctuations to the original calculation of Kim *et al.* (1987). The Reynolds number based on wall friction velocity and channel half-width δ was $Re_\tau = 180$, with a grid resolution of $128 \times 129 \times 128$ points and a domain of length $2\pi\delta$ in both streamwise and spanwise directions.

Next, statistics from this simulation are used to extract a single hairpin vortex. One signature of a hairpin vortex is a fluctuation velocity vector in the second quadrant ($u' < 0, v' > 0, w' = 0$) at the location of the vortex. In Zhou *et al.* (1999), linear stochastic estimation is used to identify the statistically most probable flow field that has a specified velocity, here $(u', v', w') = (-8.16, 3.45, 0)$, at a prescribed point in the flow (here, a wall-normal location of $y^+ = 49$). These values are the same as in Zhou *et al.* (1999), and the resulting most probable flow field is used as an initial condition for the DNS solver to study the evolution of the structure.

Figure 2(a) shows the iso-surface of the swirl criterion (10 % maximum value) for the initial condition generated by the procedure described above. This structure was evolved in time, and the structure based on 10 % maximum swirl is shown in figures 2(b) and 2(c), at two time instants, showing the development of a secondary hairpin. The threshold for these plots (i.e. the value of the level set of swirling strength λ_{ci}) was chosen to correspond to Zhou *et al.* As this hairpin vortex develops into a packet, there are only small differences in structure for different threshold values, but the size of the structure varies. Additionally, swirl is the only Eulerian criterion plotted because for appropriate thresholds there was little distinction among Eulerian criteria plots.

Figure 3(a–d) shows three two-dimensional plots of the negative-time DLE evaluated at $t^+ = 63$, as well as the location of these planes in the three-dimensional volume. Two-dimensional planes are used to study the Lagrangian structure as opposed to three-dimensional iso-surfaces because while we expect the LCS to depict the boundaries of the structure, these are not constant value surfaces. Therefore, we reconstruct the Lagrangian structure skeleton using data from a small number of two-dimensional planes. The resolution of the DLE data is higher than that of the Eulerian, as the grid of trajectories to be integrated is not restricted to be the




FIGURE 3. Two-dimensional DLE plots of the isolated hairpin evolved to $t^+ = 63$. (a) 10% maximum λ_{ci}^2 superimposed with the location of the three planes, (b) constant-streamwise cut, (c) constant-spanwise cut, and (d) constant wall-normal cut ($y^+ = 98$).

same size as the DNS grid, and linear interpolation is used to determine the velocity between grid points. For the plots of the isolated hairpin, the DLE plot resolution is greater by a factor of 6 in all three dimensions. Only a portion of the domain is plotted, including the entire streamwise domain span, half the channel width, and the middle third of the spanwise extent. All the calculations presented used on the order of 500–1000 instantaneous data sets for each plot.

The plots in figure 3 show the results from a DLE calculation using an integration time of $t^+ = 45$. A distinct LCS is seen as the boundary of the hairpin vortex. As with Hill's spherical vortex, using a larger integration time introduces greater detail, but the size and outer shape of the structure do not change. Here, the (backward) integration time may be at most $t^+ = 63$, as of course we cannot evolve backward in time earlier than the initial condition. At this moment in the hairpin evolution, not many differences exist between the DLE results using the longest possible integration time and those using an intermediate integration time.

Figure 4 illustrates a comparison between the LCS and the Eulerian criteria. In figure 4(a), a 'skeleton' of points where the DLE is greater than 60% of its maximum is shown over level sets of swirling strength λ_{ci} . The extent of the Eulerian structure depends on the threshold used, but the boundary indicated by DLE is independent of integration time or any other parameters. Although some of the Eulerian criteria discussed in the introduction approximate these structures better than others, none capture the amount of detail that DLE does. Comparable level sets of 60% maximum DLE are plotted in figure 4(b). The three-dimensional DLE field was calculated on one fifth the channel domain in the streamwise and spanwise directions with four times the resolution of the DNS. In the wall-normal direction, DLE was calculated on half the domain with six times the resolution. The three-dimensional structure marks the separation between fluid entrained into the hairpin structure, and that which convects with the outer flow. The two-dimensional 'skeleton' is again plotted to illustrate the internal LCS detail under the outer shell. As previously noted, the iso-surface is not an accurate depiction of the LCS, but at 60% maximum DLE value, the level set approximates the LCS surface.

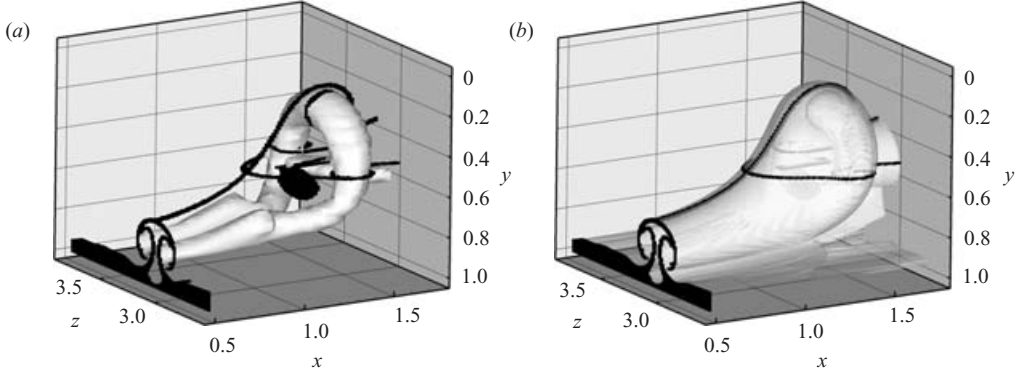


FIGURE 4. (a) Iso-surfaces of 10% maximum λ_{ci}^2 (white) superimposed on three planes of DLE $\geq 60\%$ maximum (black). (b) Translucent iso-surfaces of 60% maximum DLE (white) superimposed on three planes of DLE $\geq 60\%$ maximum (black).




FIGURE 5. Two-dimensional DLE plots of the isolated hairpin evolved to $t^+ = 171$. (a) 10% maximum λ_{ci}^2 superimposed with the location of the three planes, (b) constant-streamwise cut, (c) constant-spanwise cut, and (d) constant wall-normal cut ($y^+ = 95$).

As the single hairpin evolves further, it induces a secondary hairpin to form, and eventually forms a family or ‘packet’ of structures. The corresponding structures are shown in figure 5, at time $t^+ = 171$ (also for an integration time of $t^+ = 45$). Here, the 10% maximum swirl iso-surface is plotted for reference, and figure 5(b–d) shows three two-dimensional plots of negative-time DLE, demonstrating the ability of this Lagrangian method to capture the whole packet of hairpin vortices in great detail.

2.2. Formation of a secondary hairpin

In addition to structure identification and characterization, an analysis of hyperbolicity in the DLE field in this flow yields detailed information about how the hairpin vortex evolves into a packet, as described by Zhou *et al.* (1999). In particular, it is observed that the birth of a secondary hairpin structure corresponds to a loss of hyperbolicity (or a bifurcation) along the LCS.




FIGURE 6. Magnitude of $\langle \mathbf{n}, \mathbf{S}\mathbf{n} \rangle$ along spanwise-constant cross-sections of LCS surfaces at times (a) $t^+ = 45$, (b) $t^+ = 63$, and (c) $t^+ = 99$. Negative strain rate (compression normal to the surface) is black; positive strain rate (expansion) is white with black outline.

To investigate this phenomenon, we compute the rate of strain normal to the surface of the LCS, given by $\langle \mathbf{n}, \mathbf{S}\mathbf{n} \rangle$, where \mathbf{n} is the unit normal to the LCS. Theorem 3 of Haller (2002) states that the DLE ridge is an attracting material line if $\langle \mathbf{n}, \mathbf{S}\mathbf{n} \rangle < 0$. If this condition is not met, then the structure is not hyperbolic, and instead corresponds to a local maximum of shear.

To compute the normal rate of strain, we follow the procedure employed in Mathur *et al.* (2006). Specifically, we first find the locus of points on the LCS surface, using a two-dimensional gradient climb in regions near the local maxima of DLE. The direction of the normal vector is then approximated by calculating the Hessian of the (three-dimensional) DLE field and using the eigenvector associated with its eigenvalue of largest magnitude.

A plot of the rate of strain $\langle \mathbf{n}, \mathbf{S}\mathbf{n} \rangle$ is shown in figure 6 at three time instants in the development of the secondary hairpin vortex. Figure 6(a) shows the structures of the hairpin in the mid-span plane at time $t^+ = 45$, calculated from a negative-time DLE field which used an integration time of $t^+ = 45$. Here, the strain rates normal to the LCS are negative both upstream and downstream of the vortex head, indicating that this structure is indeed a hyperbolic attracting line.

In figure 6(b), taken at $t^+ = 63$ from a DLE field obtained using an integration time of $t^+ = 63$, a hump in the LCS above the hairpin legs is seen and a magnified picture is shown in the inset. On the downstream slope of this hump small white regions of non-negative strain rate exist, highlighted by the dashed oval, corresponding to a loss of hyperbolicity. The bifurcation of the structure coincides with the beginning of the formation of the secondary hairpin structure.

Lastly, figure 6(c) is calculated at time $t^+ = 99$ from a negative-time DLE field using integration time $t^+ = 99$. The secondary hairpin is clearly evolving and its




FIGURE 7. Two-dimensional DLE plots of the fully turbulent channel. (a) 1% maximum Q superimposed with the location of the three planes, (b) constant wall-normal cut, (c) constant-spanwise cut ($y^+ = 33$), and (d) constant-streamwise cut.

structure begins to resemble that of the primary hairpin, as the LCS begins to fold and roll up on itself. This is particularly clear in the magnified inset. Additionally, the sign of strain rate along the LCS alternates in a similar pattern in both hairpin heads, indicative of the shearing rotational flow within.

Zhou *et al.* discuss at length the physical mechanism that results in the development of this second and subsequent hairpin structures. The Lagrangian criteria when used in this way offer a quantitative method for recognizing their formation and interpreting their generation as a loss of hyperbolicity of the Lagrangian coherent structures. For the same initial condition as used in this study, Zhou *et al.* observe characteristic indications of hairpin development at $t^+ = 72$, whereas using LCS the bifurcation occurs and can be detected at least as early as $t^+ = 63$.

2.3. Fully turbulent channel

Finally, this analysis was applied to the fully turbulent channel data. In figure 7(a), the Q -criterion structures are shown as level sets of 1% maximum value. Also shown in this plot are the locations of three two-dimensional planes on which DLE was calculated. The results are shown in figure 7(b–d). As expected, the LCS plotted are clean and sharp. Several of these curves have the same shape as those of the isolated hairpin in the respective plane, and support the notion that the fully turbulent channel is populated with similar structures. These locations are highlighted by a white box, and can be compared with figures 3–5.

In figure 8(a, b), the DLE and Q -criterion are plotted at the same constant-streamwise location. This comparison highlights the fact that the LCS clearly depicts structures in locations where the Q -criterion would not if plotted with a large threshold. Two such structures are marked with a white box. Also, the finer resolution of the DLE plot yields more detail than possible with the Eulerian criteria, which




FIGURE 8. Constant-streamwise cuts of (a) DLE and (b) Q -criterion.

require derivatives of the velocity field and are thus restricted to the resolution of the original data.

3. Conclusions

Lagrangian coherent structures (LCS) have been identified for various flows in a plane channel, including an isolated hairpin vortex and a fully developed turbulent flow, by calculating the direct Lyapunov exponent (DLE). The Lagrangian method captures features of the flow that are familiar from flow visualization experiments, and are also described by various Eulerian criteria currently in use, but the DLE field yields greater detail than existing Eulerian criteria. This is partially because, unlike Eulerian criteria, the DLE may be evaluated on a finer grid than the original velocity data. Additionally, using the Lagrangian criteria, one may quantify the boundary of a vortex as a local maximum of the DLE field. Whereas the size and shape of the Eulerian structures depend on a user-defined threshold, the locations of the LCS are independent of such parameters. Increased integration time yields greater detail, but outer boundaries of the structures do not vary. Lastly, the DLE is truly independent of coordinate frame and would yield the same results for a non-Galilean-invariant coordinate transformation.

The development of an isolated hairpin vortex is studied, and it is shown that the birth of a secondary hairpin corresponds to a loss of hyperbolicity along the LCS. Thus, the Lagrangian criteria can provide a quantitative way of determining when these structures are born. Previously, such events have been identified using qualitative, visual methods (Zhou *et al.* 1999), but the Lagrangian method allows one to detect these events at their earliest stages.

The benefits of the Lagrangian method have a cost: Lagrangian calculations are more computationally intensive than any of the Eulerian criteria, as they involve integration of particle trajectories from each point at which the DLE value is desired, and hence large amounts of time-resolved data are necessary for the trajectory integration. However, as postprocessing represents a relatively small part of the computational time devoted to most flow calculations, this drawback is often not severe. Each DLE data set presented in this paper took on the order of 1–2 computational hours, using 500–1000 data sets, requiring approximately 150 GB of storage. Higher-Reynolds-number flows would naturally necessitate a greater cost for the same calculations.

These methods may be especially useful for time-resolved experimental data, for instance from particle image velocimetry, which are often too noisy to compute derivatives needed for the Eulerian criteria. DLE may also be useful as a tool to calibrate the more easily implemented Eulerian schemes, for instance to determine appropriate thresholds for the Eulerian criteria.

The authors wish to thank Dr Alexander Smits for his input on the results presented in this paper, Mingjun Wei and Milos Ilak for their work with the turbulent channel flow simulation, and Francois Lekien for his helpful thoughts on the Lagrangian structure analysis. This work was generously supported by an NSF Graduate Research Fellowship, NSF awards CMS-0347239 and DMS-04-04845, and AFOSR grant FA 9550-06-0092.

REFERENCES

- CHAKRABORTY, P., BALACHANDAR, S. & ADRIAN, R. J. 2005 On the relationships between local vortex identification schemes. *J. Fluid Mech.* **535**, 189–214.
- CHONG, M. S., PERRY, A. E. & CANTWELL, B. J. 1990 A general classification of three-dimensional flow fields. *Phys. Fluids A* **2** (5), 765–777.
- HALLER, G. 2001 Distinguished material surfaces and coherent structures in 3D fluid flows. *Physica D* **149**, 248–277.
- HALLER, G. 2002 Lagrangian coherent structures from approximate velocity data. *Phys. Fluids* **14** (6), 1851–1861.
- HALLER, G. 2005 An objective definition of a vortex. *J. Fluid Mech.* **525**, 1–26.
- HALLER, G. & YUAN, G. 2000 Lagrangian coherent structures and mixing in two dimensional turbulence. *Physica D* **147**, 352–370.
- HEAD, M. & BANDYOPADHYAY, P. 1981 New aspects of turbulent boundary-layer structure. *J. Fluid Mech.* **107**, 297–338.
- HUNT, J. C. R., WRAY, A. A. & MOIN, P. 1988 Eddies, stream, and convergence zones in turbulent flows. *Center for Turbulence Research Rep.* CTR-S88.
- JEONG, J. & HUSSEIN, F. 1995 On the identification of a vortex. *J. Fluid Mech.* **285**, 69–94.
- KIM, J., MOIN, P. & MOSER, R. 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.* **177**, 133–166.
- KOH, T. Y. & LEGRAS, B. 2002 Hyperbolic lines and the stratospheric polar vortex. *Chaos* **12** (2), 382–394.
- LAPEYRE, G. 2002 Characterization of finite-time Lyapunov exponents and vectors in two-dimensional turbulence. *Chaos* **12** (3), 688–698.
- LEKIEN, F. & LEONARD, N. 2004 Dynamically consistent Lagrangian coherent structures. In *AIP: 8th Experimental Chaos Conference, CP 742*, pp. 132–139.
- MATHUR, M., HALLER, G., PEACOCK, T., RUPPERT-FELSOT, J. E. & SWINNEY, H. L. 2006 Uncovering the lagrangian skeleton of turbulence. *Phys. Rev. Lett.* (submitted).
- SHADDEN, S., DABIRI, J. & MARSDEN, J. 2006 Lagrangian analysis of fluid transport in empirical vortex ring flows. *Phys. Fluids* **18**, 047105-1–047105-11.
- SHADDEN, S., LEKIEN, F. & MARSDEN, J. 2005 Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D* **212**, 271–304.
- SMITH, C. & WALKER, J. 1991 On the dynamics of near-wall turbulence. In *Turbulent Flow Structure Near Walls* (ed. J. Walker). The Royal Society (First published in *Phil. Trans. R. Soc. Lond. A* **336**, 1991).
- THEODORSEN, T. 1955 The structure of turbulence. In *50 Jahre Grenzschichtforschung* (ed. H. Görtler & W. Tollmien). Vieweg and Sohn.
- VOTH, G. A., HALLER, G. & GOLLUB, J. P. 2002 Experimental measurements of stretching fields in fluid mixing. *Phys. Rev. Lett.* **88**, 254501.
- ZHOU, J., ADRIAN, R. J., BALACHANDAR, S. & KENDALL, T. M. 1999 Mechanisms for generating coherent packets of hairpin vortices in channel flow. *J. Fluid Mech.* **387**, 353–396.