On the Finite-Time Scope for Computing Lagrangian Coherent Structures from Lyapunov Exponents

No Author Given

Abstract Lagrangian coherent structures (LCS) can be extracted from time-dependent vector fields by means of the finite-time Lyapunov exponent (FTLE). While the LCS approach has proven successful in many areas and applications for the analysis of time-dependent topology, it is to some extent still an open problem how the finite time scope is appropriately chosen. One has to be aware that the introduction of this finite time scope in the Lyapunov exponent, where the time scope was originally infinite, is largely responsible for the recent success of the FTLE in analysis of real-world data. Hence, there is no general upper bound for the time scope: it depends on the application and the goal of the analysis. There is, however, a clear need for a lower bound of the time scope because the FTLE converges to the eigenvalue of the rate of strain tensor as the time scope approaches zero. Although this does not represent a problem per se, it is the loss of important properties that causes ridges in such FTLE fields to lose the LCS property. LCS are time-dependent separatrices: they separate regions of different behavior over time. Thereby they behave like material constructs, advecting with the vector field and exhibiting negligible cross flow. We present a method for investigating and determining a lower bound for the FTLE time scope at isolated points of its ridges. Our approach applies the advection property to the points where attracting and repelling LCS intersect. These points are of particular interest because they are the loci where Lagrangian dynamics varies most and which are important in typical questions of Lagrangian topology. We demonstrate our approach with examples from dynamical systems theory and computational fluid dynamics.

1 Introduction

Traditional vector field topology [7] deals with several types of distinguished streamlines. Those that degenerate to isolated points play a special role in the form of critical points: if they exhibit saddle-type flow behavior in their linearized neigh-

borhood, they give rise to separatrices. Separatrices are another type of distinguished streamlines: those that converge to saddle-type critical points in forward or reverse time. Other cases include periodic orbits, i.e., isolated closed streamlines and invariant tori. Common to all these special cases is their derivation: these constructs are obtained as limit cases as integration time of the streamline approaches infinity [1]. This, except for technical issues, does not represent a problem. Steady vector fields are constant over time and hence do not require the notion of scope of time.

The traditional Lyapunov exponent (LE) shares many aspects with vector field topology. Despite additional technical issues regarding numerics and boundedness of the temporal domain, its properties are well defined and do not depend on additional parameters. It was, however, the boundedness of the temporal domain together with the aim of choosing a region of interest also in terms of time that lead to the finite-time Lyapunov exponent (FTLE). A main risk, however, is the choice of too short advection time for FTLE computation and resulting misinterpretation.

Lagrangian coherent structures (LCS) by means of the FTLE [6] have become a prominent alternative for the investigation of time-dependent topology of vector fields. In this paper, we present a method for validating and choosing the advection time parameter for 2D FTLE computation with respect to prescribed error measures. We base our approach on a main principle of LCS: their advection property [12], i.e., their behavior as material lines.

1.1 Finite-Time Lyapunov Exponent

Vector fields exhibit a spectrum of Lyapunov exponents. It is the largest exponent in this spectrum that has become a prominent tool for predictability analysis in time-dependent vector fields. The LE can be determined by computing two neighboring trajectories in phase space and measuring their separation rate as time approaches infinity. Since the LE was originally introduced for predictability analysis, it has to reflect properties along trajectories. Therefore, precaution has to be taken to assure that the "neighboring" trajectories do not separate too far, e.g., by renormalization [2].

Since the systems under investigation are often defined on a finite temporal domain only, or because it is often the objective of the user to restrict the analysis to a temporal region of interest, the FTLE has been becoming more and more popular. Again, there are techniques to assure proximity of the (implicitly) involved trajectories, such as the localized FTLE [8].

Whereas the LE and FTLE have been applied for a long time to predictability analysis, there is a recent trend in the visualization community to use FTLE for revealing the topology of time-dependent vector fields. Haller [6] showed that ridges present in the FTLE represent a time-dependent counterpart to separatrices from vector field topology [7]: they separate regions of qualitatively different behavior.

In the context of time-dependent vector field topology, the FTLE is typically computed from the *flow map* $\phi_{t_0}^{t_0+T}(\mathbf{x})$, mapping seed points $\mathbf{x}$ of trajectories to their

end points after advection for finite time T. According to Haller [6], the finite-time Lyapunov exponent $\sigma(\mathbf{x}, t_0, T)$ computes from the flow map ϕ as follows:

$$\sigma(\mathbf{x}, t_0, T) = \frac{1}{|T|} \ln \sqrt{\lambda_{\max} \left[\left(\nabla \phi_{t_0}^{t_0+T}(\mathbf{x}) \right)^{\top} \nabla \phi_{t_0}^{t_0+T}(\mathbf{x}) \right]}, \tag{1}$$

$\lambda_{\max}(\cdot)$ being the major eigenvalue.

2 Method

In this paper, we do not address the choice of the advection time parameter T in the context of specific applications and related application-oriented questions. In fact, we rather consider the fundamental advection property of LCS to find appropriate advection times inside a prescribed interval of temporal interest.

Let us examine the two extreme choices for T: infinity and zero. As $T \to \infty$, the FTLE converges[1] to the classical LE. Hence, this extreme case imposes no particular concerns. The interpretation and use of the LE is well established. If, on the other hand, $T \to 0$, it can be easily shown that the FTLE converges to the largest eigenvalue of the rate of strain tensor of the vector field. In this case, due to the instantaneous property of the rate of strain tensor, it cannot be assumed that the advection property [6] of ridges inside this field is satisfied, consistent with Shadden et al.'s Theorem 4.4 [12].

Shadden et al. measured fluxes of the instantaneous velocity field across FTLE ridges as a means of verifying the advection property of given FTLE ridges. However, zero flux alone is a necessary but not sufficient condition for advection: it does not capture tangential motion, i.e., the advection component along the ridge does not need to satisfy the advection property even if the flux across the ridge is zero. Unfortunately, it is hard to identify point correspondences between ridges from FTLE fields with different starting times t_0 because common ridge definitions, such as height ridges by Eberly [3], are purely geometric, i.e., they are not represented by identifiable particles that advect. We therefore follow a different approach: we measure the advection property only for distinguished points on the FTLE ridges, i.e., we identify point correspondences in terms of advection.

For sufficiently well defined LCS (according to Shadden et al.), the advection property holds for both ridges in the FTLE field computed from forward trajectories and ridges in the FTLE field computed from reverse-time trajectories. These ridges typically intersect. The intersection points give rise to hyperbolic trajectories [5, 11] and are important in the Lagrangian skeleton of turbulence [9]. Hence, if both types of ridges satisfy the advection property, this property must also hold

[1] We follow a conceptual, or continuous, argumentation here. In the discrete case, preacautions have to be taken such as renormalization [2] or evaluation of the velocity gradient along the trajectories [8].

for their intersections. In 2D vector fields, these intersections are isolated points (we exclude congruent ridge regions because these represent degenerate cases that can be avoided by perturbing the vector field). Thus, one needs to identify point correspondence between ridge intersections of successive FTLE time steps (at t_0 and $t_0 + \Delta t$).

As FTLE computation typically depends on time-dependent discretized vector fields, it is very hard to derive generic rules for parametrizing FTLE visualizations. FTLE computation is heavily dependent on the structure and position of the sampling grid, and the advection time T. Furthermore, FTLE ridges only tend to represent LCS if they are sufficiently sharp [12]. Providing automatic strategies for finding a good choice of parameters is very hard because many decisions directly depend on the goal of the user. Therefore, FTLE visualization, as other feature extraction procedures, is typically a trial-and-error procedure representing the basic exploration by visualization. Once the user has found a sampling grid that sufficiently captures the FTLE structures he wants to see, has found an appropriate threshold filtering out insufficiently sharp ridges (this can be accomplished by filtering by an eigenvalue of the FTLE Hessian [10]), and has found a minimum and maximum FTLE advection time T, our new technique takes over these parameters. Although our method can be further manually parametrized, it performs well with the default values (described below). As the result, our technique provides a plot of advection discrepancy with respect to FTLE advection time T. It can identify local and global optima inside the prescribed range of T, and in particular, provide a lower bound on T with respect to a prescribed accuracy in terms of average advection error of ridge intersections.

One contribution of this paper is a comparably robust technique for tracking FTLE ridge intersections. It is based on the assumption that the user already has chosen a meaningful time scope for the analysis, i.e., where the lower bound of the time interval already produces comparably sharp ridges. According to Shadden et al. [12], this will lead to FTLE ridges that already approximately satisfy the advection property. Therefore, the advection property of the intersections can be exploited for making the correspondence finding more robust. We then present a technique to quantify the advection property of the intersections and show how it can be used to optimize the FTLE advection time T for obtaining FTLE fields where the ridge intersections satisfy the advection property up to a prescribed average error. Interestingly, we observed for typical discretizations of FTLE sampling grids that the advection property is not a monotonic function of the finite FTLE advection time: we observed local optima. Hence, our findings indicate that once an appropriate interval of possible finite advection times is identified, it is desirable to choose an optimum inside this interval, possibly by our technique.

2.1 FTLE Ridge Intersection

FTLE sampling is a computationally expensive task. One is therefore typically limited to comparably low resolutions of the FTLE sampling grid. Although there are techniques that adapt the sampling grid to the vicinity of LCS [4] and ridges in general [10], they are, as many adaptive sampling techniques, susceptible to missed features, or would require a-priori knowledge about the data. We therefore base our analysis on a regular sampling of the FTLE and use regions of interest if high resolution is required.

The ridge lines in the resulting scalar FTLE fields are extracted according to Eberly's [3] criterion. The subsequent intersection procedure for obtaining the intersections is a trivial problem in the 2D context investigated in this paper. To assert sufficient quality of the geometric intersections, we impose a minimum intersection angle threshold.

2.2 Intersection Tracking

As illustrated in Figure 1, we extract forward and reverse FTLE fields for both t_0 and $t_0+\Delta t$. This leads to two sets of intersections, $\mathbf{i}(t_0)$ at time t_0 and $\mathbf{i}(t_0+\Delta t)$ at $t_0+\Delta t$. A straightforward approach would use a very small Δt. This would produce almost identical ridges and hence finding correspondences between their intersections would lead to a trivial tracking problem. Further, the limit case $\Delta t \to 0$ would be used for defining the intersection velocity

$$\mathbf{u}_i(t_0+\Delta t/2) = (\mathbf{i}(t_0+\Delta t)-\mathbf{i}(t_0))/\Delta t. \tag{2}$$

We have carried out such an analysis using the quad-gyre example, discussed in Section 3.1. Unfortunately, it turns out that ridge extraction tends to be only accurate in the order of the cell size of the FTLE sampling grid, and thus, using small Δt leads to poor accuracy of $\mathbf{u}_i$, because $\mathbf{i}(t_0)$ and $\mathbf{i}(t_0+\Delta t)$ are too close with respect to the resolution of the FTLE sampling grid. Thus, using larger Δt improves accuracy. The time span Δt can be estimated from the average speed $\bar{u}$ of the vector field and the cell size h of the FTLE sampling grid: $\Delta t = ch/\bar{u}$ with a constant $c > 1$. However, in this case, Δt is not small enough to allow the estimation of the intersection velocity by the linearization (2), i.e., the intersection point cannot be assumed to move at constant speed along a straight line at sufficient precision during Δt.

Using comparably large Δt leads to another problem: intersection correspondences are harder to identify because the ridge intersections move over larger distances; the FTLE ridges move and deform, and may even disappear or new ones might originate. However, since the input to our method is a desired parametrization (see Section 2) of FTLE and already satisfies the advection property to some extent, we can utilize the advection property to solve the correspondence problem: the intersection points at time t_0 are advected along path lines to time $t_0+\Delta t$ and

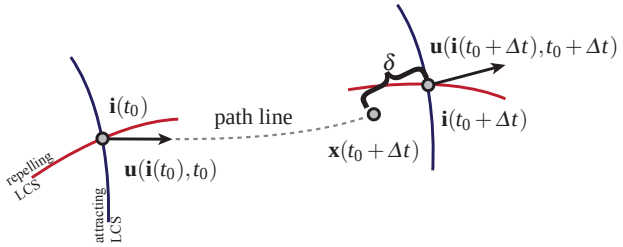

Fig. 1 FTLE ridge intersections $\mathbf{i}$ obtained with a fixed FTLE advection time T with starting times t_0 (left) and $t_0+\Delta t$ (right). The spatial discrepancy δ is defined as the distance between $\mathbf{i}(t_0+\Delta t)$ and the endpoint of the path line started at $\mathbf{i}(t_0)$ after advection time Δt.

these advected points are checked for correspondence with the ridge intersections at $t_0+\Delta t$ (Figure 1). To further avoid erroneous correspondences, a threshold limiting discrepancy δ is imposed and the correspondence is identified as the closest remaining candidate with respect to δ. The measure δ reflects the Lagrangian advection consistency of the FTLE ridge intersections and is denoted as *advection discrepancy* in our technique.

2.3 Measuring Advection Quality

Various evaluation approaches are possible for the advection discrepancy δ. For greedy considerations we determine $\delta_{\min}(T)$, the minimum δ of all intersections at a given value of T. We also compute $\bar{\delta}(T)$, the average over all intersections. Of course many other techniques, e.g., from statistics, can be applied.

Whereas $\delta_{\min}$ tends to show the best case in terms of advection, $\bar{\delta}$ can be used to get an overall picture of the advection quality of the FTLE ridges. All plots are in units of FTLE sampling grid cell size. This way one can easily choose a limit of advection discrepancy in terms of FTLE grid cells and then visually or numerically identify which regions of T satisfy this requirement.

A possible problem are intersections that do not or do only slowly move over time. It is likely that these intersections exhibit small δ and hence exhibit inappropriately high advection quality since they exhibit low discrepancy only due to the fact that they stand still. This can be addressed by an inverse weighting of δ with the length of the corresponding trajectory from Figure 1. However, it is unlikely that a time-dependent vector field exhibits a zero over extended time. We did not encounter this problem in the context of CFD simulations, although it could appear in vector fields from other domains.

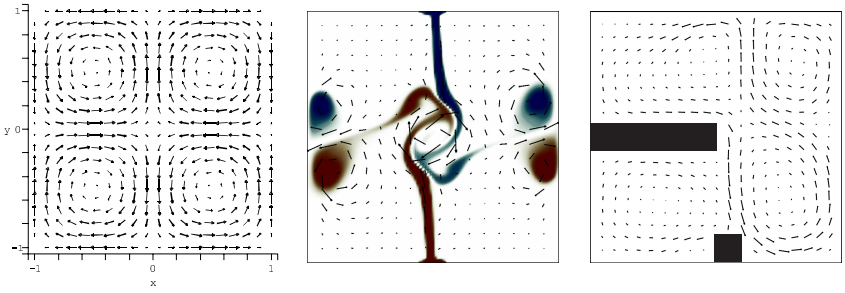

Fig. 2 Left: quad-gyre example. Center: buoyant plume dataset. Right: buoyant flow with obstacles, heated at its lower side and cooled at the top.

2.4 Finding Locally Optimal Advection Times T

Having now the building blocks at hand, it would be possible to run an optimization process over all feasible grid resolutions, advection times T, starting times t_0, and ridge sharpness thresholds for a given dataset. According to Theorem 4.4 in [12], this would likely result in infinite resolution and infinite advection time T for all t_0. On the other hand, there is a region of interest for all parameters of the analysis (Section 2). This is particularly important for the advection time scope T which depends on the questions of the user and the application (FTLE ridges typically grow with time and can fill the complete domain in, e.g., convective flows).

Therefore, we perform an analysis of the advection discrepancy by uniform sampling of T inside the prescribed advection time interval and provide a plot together with the T producing global minima of the different discrepancy measures $\underline{\delta}_{\min}$ and $\overline{\delta}$. Finding the global minimum inside the region of interest of T serves here as a straightforward example. Of course more sophisticated techniques for data analysis can be applied. An important question that can be interactively and visually answered from the plot of the advection discrepancy is the minimum required FTLE advection time T that makes sure that the advection property is in the average satisfied up to a given tolerance. In our current approach we still advocate visual inspection of the plots, in particular because they tend to exhibit outliers and are therefore susceptible to errors if simple automatic analysis techniques are applied.

3 Results and Evaluation

In the following, we demonstrate and evaluate our method on three time-dependent 2D examples: the analytic quad-gyre example (Section 3.1), and two CFD simulations of buoyant flow (Sections 3.2 and 3.3).

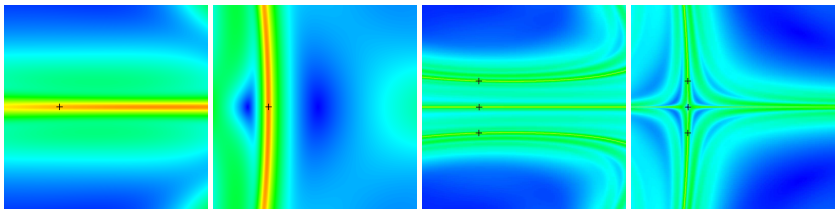

Fig. 3 FTLE visualization for the quad-gyre dataset at time $t_0 = 20$ in a region of interest at the center of the dataset. The FTLE advection times T vary: 4 (left), -4 (left center), 13 (right center), and -13 (right). The intersection points between forward-time and reverse-time FTLE ridges are marked by black crosses.

3.1 Quad-Gyre

The double gyre example was introduced by Shadden et al. [12] to examine FTLE and LCS and to compare them to vector field topology. This dataset consists of two vortical regions separated by a straight separatrix that connects two saddle-type critical points: one temporally oscillating horizontally at the upper edge and the other synchronously oscillating horizontally along the lower edge (Figure 2 (left)). This is a prominent example where vector field topology gives a substantially different result from that by FTLE. This dataset is temporarily periodic. To avoid boundary issues, we use a larger range of field, resulting in four gyres. Therefore, we call this example quad-gyre. Using

$$\begin{aligned} f(x,t) &= a(t)x^2 + b(t)x, \\ a(t) &= \varepsilon \sin(\omega t), \\ b(t) &= 1 - 2\varepsilon \sin(\omega t), \end{aligned} \tag{3}$$

the quad-gyre is defined in space-time as follows:

$$u(x,y,t) = \begin{pmatrix} -\pi A \sin(\pi f(x)) \cos(\pi y) \\ \pi A \cos(\pi f(x)) \sin(\pi y) \frac{df}{dx} \\ 1 \end{pmatrix}. \tag{4}$$

As recommended by Shadden et al., we use the configuration $\varepsilon = 1/4$, $\omega = \pi/5$, and $A = 1/10$. Figure 2 (left) shows a hedgehog plot of the vector field at $t = 0$ and Figure 3 depicts the FTLE inside a region of interest at the center.

Figure 5 shows a result plot from our method and Figure 4 a corresponding visualization. It is apparent that the plot of $\delta_{\min}$ is much lower than the plot of $\bar{\delta}$, mostly due to outliers, but both are similar with respect to local extrema and trends. It is also apparent how well the ridge sharpness criterion rejected outliers from the analysis. This fact is consistent with the statement by Shadden et al. that a FTLE ridge has to be sufficiently sharp to represent an LCS. Our test directly reflects the fact that as the analysis is restricted to sharp ridges, the advection principle is substantially better satisfied.

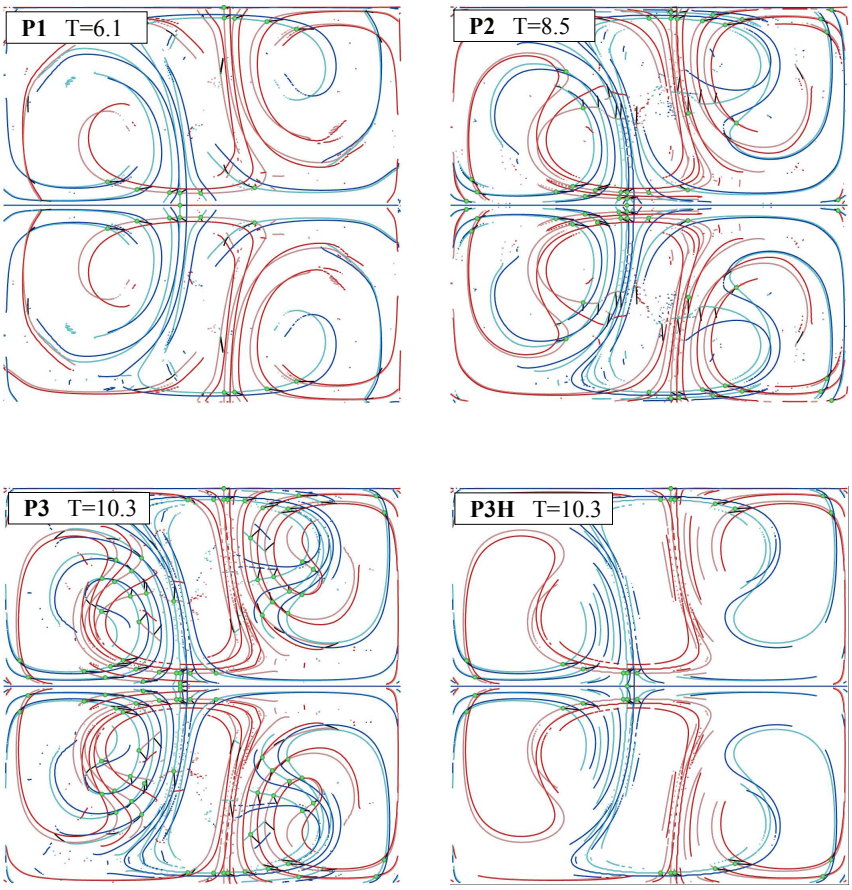

Fig. 4 Quad-gyre example. Forward FTLE ridges in red and backward in blue, low saturated starting at $t_0 = 20$ and high saturated at $t_0 = 20.3$ for various special T marked in the discrepancy plot in Fig 5. 'P3H' is the same as 'P3' but with a threshold suppressing ridges that are not sharp. For each intersection at $t_0 = 20$, a pathline is visualized (black). If the pathline leads to a corresponding intersection at $t_0 = 20.3$, the starting point is visualized by a green dot.

We can also well observe from the plot that the advection property is more and more violated as the FTLE advection time T reaches low values. This again supports Shadden et al.'s findings.

3.2 Buoyant Plumes

Our first CFD example is a time-dependent simulation of buoyant plumes inside a 2D box, depicted in Figure 2 (center). The domain is a square of 1 meter side length filled with air. The air is initially in rest and at 40°C. Gravity forces are

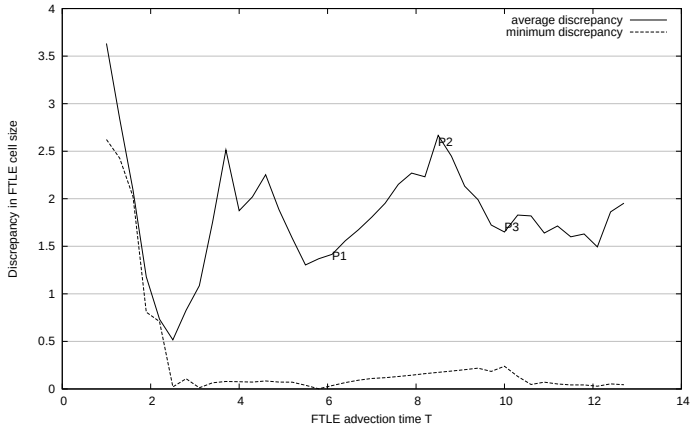

Fig. 5 Quad-gyre example. Discrepancy plots for a setup with FTLE starting time $t_0 = 20$ and ridge intersection advection to time $t = 20.3$ (Δt = 0.3). Global optimum in $\bar{\delta}$ plot at $T = 2.5$. For three special T, marked as $P1$ - $P3$, visualizations are given in Fig 4.

acting downwards. No-slip conditions are imposed on all boundaries. The left and the right walls both are supplied with adiabatic boundary conditions, making the walls neutral in the sense that there is no heat exchange with the outside. There is a region in the center of the lower wall that is heated to 75°C and a corresponding region on the upper wall that is cooled to 5°C.

Figure 6 shows the plot from our analysis using $T_0 = 20$ and $\Delta t = 1$. It can be seen that the average discrepancy has a decreasing trend. We added a plot of the number of used intersections for judgment of uncertainty. For validation and for further investigation, we generated images where the complete ridges are advected, not only their intersections (Figure 7). It is visible in the $\bar{\delta}$ plot that $T = 4$ exhibits high error, $T = 6$ reduced, and $T = 8$ already average error below the size of a FTLE sampling cell. The global optimum inside the $\bar{\delta}$ plot is at $T = 8.6$. The images in Figure 7 support this finding, it can be seen that for $T = 8$ the advected ridges and the ridges of the corresponding time fit well almost everywhere. It has to be noted that we advected here the repelling ridges, i.e., those from forward FTLE. We also did this for the attracting ridges and there the deviation was much smaller. Nevertheless, according to the initial motivation of our method, such a comparison with complete advected ridges instead of intersections can not detect tangential discrepancy between the motion of the ridges and the vector field behavior.

3.3 Buoyant Flow with Obstacle

Our second CFD example is again a buoyant unsteady flow (Figure 2, right). During the 80 seconds of simulation time, a convective flow evolves due to the effect of

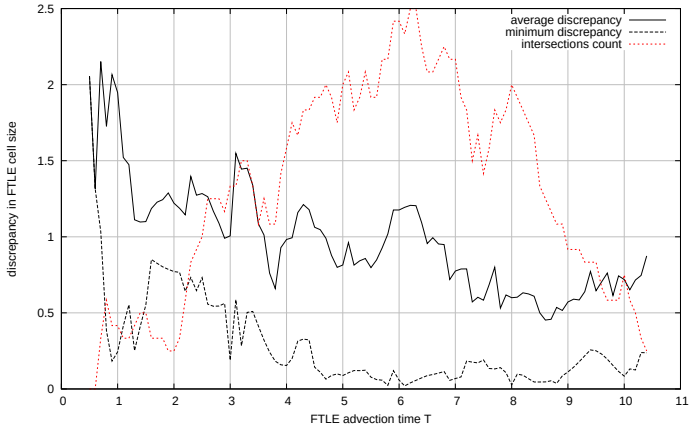

Fig. 6 Buoyant plumes dataset. Examining the plot of $\overline{\delta}$, we identify clearly too low advection times (until $T = 4$), medium quality ranges (around $T = 6$), and comparably high quality (lower than the size of a FTLE sampling cell) for $T = 8$ and larger. Please refer to Fig. 7.

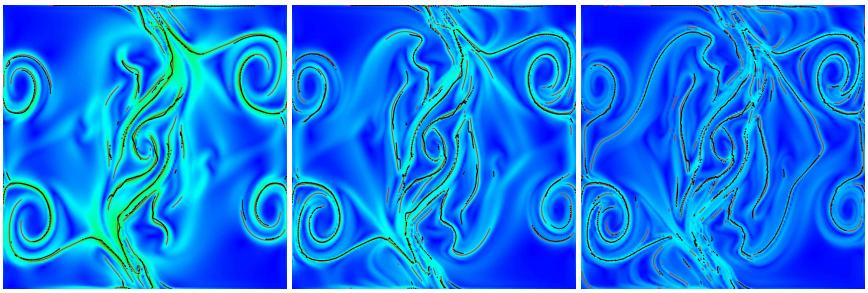

Fig. 7 Buoyant plumes dataset. Advected FTLE ridges (black dots) and corresponding ridges (red lines) of later FTLE visualizing ridge advection quality, for $T = 4$ (left), $T = 6$ (center), and $T = 8$ (right). The temporal difference Δt between the t_0 of the FTLE fields is 1 second (this is also the time that was used for the advection of the ridges). One can see that with $T = 8$ advection quality is getting sufficient for typical applications. Please refer to Fig. 6.

heating the lower wall to 75°C and cooling down the opposing upper wall to 5°C. Two rectangular obstacles, a small one on the bottom wall and a larger one on the side walls, prevent the onset of a simple circular flow. We use $t_0 = 10$ to study the early phase of turbulence development (Figure 2, right). Note that in the following description, time is given in seconds, length units are in meters.

Convective flows are known for their complex topology. We set up a FTLE visualization with $t_0 = 10$ and used a threshold for the eigenvalue of the Hessian to suppress many weak FTLE ridges caused by the turbulent flow. We applied our method for T in the interval $[0.4, 2.6]$, Figure 8 contains the resulting advection discrepancy plot and Figure 9 shows visualizations for advections times selected according to the plot. In this analysis, $T = 2.125$ produced the global minimum $\overline{\delta}$ inside the prescribed interval. If advanced data analysis techniques would be used,

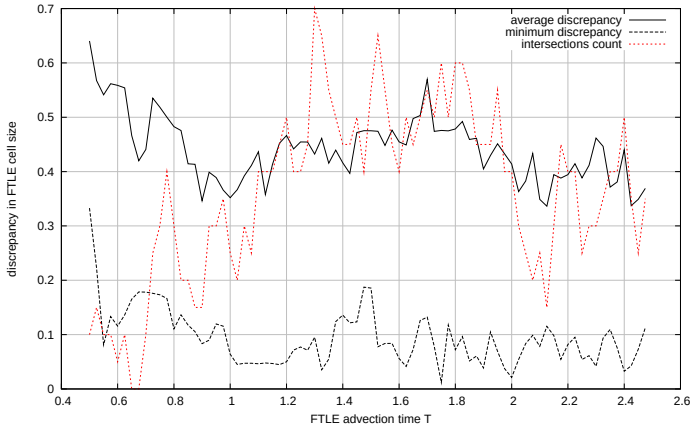

Fig. 8 Buoyancy dataset: average ($\bar{\delta}$) and minimum (δ_{min}) advection discrepancy plots. Again it is evident that increasing FTLE advection time improves the advection property. Please refer to Fig. 9 for the investigation of selected advection times T ($T = 0.2$, $T = 0.5$, and $T = 1.0$).

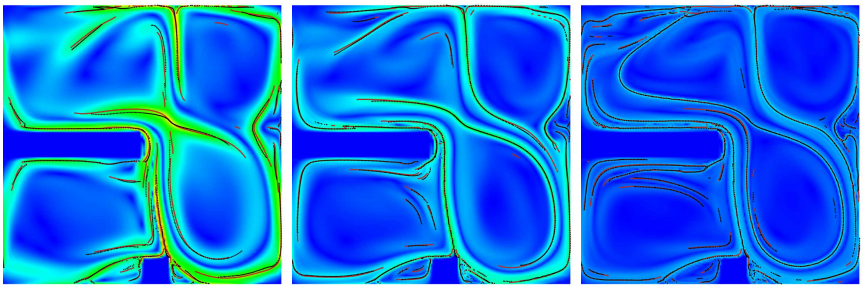

Fig. 9 Buoyancy dataset. Advected FTLE ridges (black dots) and corresponding ridges (red lines) of later FTLE visualizing ridge advection quality, for $T = 0.2$ (left), $T = 0.5$ (center), and $T = 1.0$ (right). The temporal difference Δt between the t_0 of the FTLE field is 0.05 seconds (this is also the time that was used for the advection of the ridges). One can see from the δ plot that with $T = 1$ advection quality is getting sufficient for typical applications. Please refer to Fig. 8.

that are able to address noise, it is likely that a lower value would be obtained, more consistent with our observation of a value of $T = 1$. As in our other experiments, the advection discrepancy first decays rapidly as T increases and once it has reached a certain quality, the decay slows down. However, due to the high LCS complexity in this flow, our chosen FTLE sampling grid resolution does not capture the intersections very robustly. This is no surprise since it is known that such flows exhibit very complex, i.e., massively folded LCS, and hence are very difficult to investigate (the ridges in Figure 9 (right) are folded, i.e., the line features consist of several ridges).

4 Conclusion

We presented a technique for measuring the advection property of FTLE ridge intersections in 2D vector fields. Our approach can be seen as dual to the flux-based approach by Shadden et al. By analyzing the temporal behavior of the ridge intersections, not only flow discrepancy orthogonal to the ridges, but also tangential to the ridges can be revealed. Our measurements support the theoretical behavior stated by Theorem 4.4 in [12], i.e., that the error in the advection property tends to decrease with increasing T. We also have observed that the advection property is better satisfied by sharp ridges. An apparent property of our approach is the noise in the order of the size of a FTLE cell. As mentioned, this error is seemingly introduced by the ridge extraction stage and hence related to the resolution of the FTLE sampling grid. A thorough analysis is, however, subject to future work. It would also be interesting to compare our approach with the flux-based approach by Shadden et al. and to compare their accuracy. Further, it seems that both approaches could complement each other, possibly leading to a more robust and more accurate technique.

References

1. D. Asimov. Notes on the topology of vector fields and flows. Technical Report RNR-93-003, NASA Ames Research Center, 1993.
2. G. Benettin, L. Galgani, A. Giorgilli, and J. M. Strelcyn. Lyapunov characteristic exponent for smooth dynamical systems and hamiltonian systems; a method for computing all of them. *Mechanica*, 15:9–20, 1980.
3. D. Eberly. *Ridges in Image and Data Analysis*. Computational Imaging and Vision. Kluwer Academic Publishers, 1996.
4. C. Garth, F. Gerhardt, X. Tricoche, and H. Hagen. Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Transaction on Visualization and Computer Graphics.*, 13(6):1464–1471, 2007.
5. G. Haller. Finding finite-time invariant manifolds in two-dimensional velocity fields. *Chaos*, 10(1):99–108, 2000.
6. G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D*, 149:248–277, 2001.
7. J. Helman and L. Hesselink. Representation and display of vector field topology in fluid flow data sets. *Computer*, 22(8):27–36, 1989.
8. J. Kasten, C. Petz, I. Hotz, B. Noack, and H.-C. Hege. Localized finite-time Lyapunov exponent for unsteady flow analysis. In M. Magnor, B. Rosenhahn, and H. Theisel, editors, *Vision, Modeling, and Visualization*, volume 1, pages 265–274. Universität Magdeburg, Inst. f. Simulation u. Graph., 2009.
9. M. Mathur, G. Haller, T. Peacock, J. E. Ruppert Felsot, and H. L. Swinney. Uncovering the Lagrangian skeleton of turbulence. *Physical Review Letters*, 98(14):144502+, Apr 2007.
10. F. Sadlo and R. Peikert. Efficient visualization of Lagrangian coherent structures by filtered AMR ridge extraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(5):1456–1463, 2007.
11. F. Sadlo and D. Weiskopf. Time-dependent 2-D vector field topology: An approach inspired by Lagrangian coherent structures. *Computer Graphics Forum*, 29(1):88–100, 2010.
12. S. C. Shadden, F. Lekien, and J. E. Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D Nonlinear Phenomena*, 212(3-4):271–304, 2005.