Comparing Finite-Time Lyapunov Exponents in Approximated Vector Field

Stefan Koch, Sebastian Volke, Gerik Scheuermann,
Hans Hagen and Mario Hlawitschka

Abstract In the context of fluid mechanics, larger and larger flow fields arise. The analysis of such fields on current work stations is heavily restricted by memory. Approximation limits this problem. In this paper, we discuss the impact of vector field approximation on visualization techniques on the example of Finite-Time Lyapunov Exponent (*FTLE*) computations. Thereby, we consider the results of three different vector field compression approaches and analyze the reliability of integration results as well as their impact on two different FTLE variants.

1 Introduction

Visualization is one of the most important tools for the investigation of complex flow fields. Many different visualization approaches have been developed. These include techniques ranging from the extraction and depiction of specific flow feature such as vortices [10, 11], to the computation of the topological structure of vector fields [6, 18]. There also exist many methods to cluster vector fields [3, 16] or compute simplified representations [1, 16]. Their main concern is to allow users to get a good overview of the overall flow behavior. A detailed summary of the most common visualization approaches can be found in these state-of-the-art reports [8, 12, 13].

Furthermore, to handle large and complex simulation results, compression approaches were introduced in the past [2, 9, 17]. These algorithms aim at a dataset size reduction to allow an easier transmission of datasets via networks and a fast

Stefan Koch, Sebastian Volke, Gerik Scheuermann, Mario Hlawitschka
Leipzig University, Augustusplatz 10, 04109 Leipzig, Germany,
e-mail: stefan.koch, volke, scheuermann, hlawit@informatik.uni-leipzig.de

Hans Hagen
University of Kaiserslautern, Erwin-Schrödinger-Str., 67663 Kaiserslautern, Germany
e-mail: hagen@informatik.uni-kl.de

evaluation even on low-end computers. While small-scale flow features could be lost during compression, large-scale features, as well as the global flow characteristic, are preserved in most cases.

Although, the previous works provide an evaluation of the quality of their results, often there is no direct comparison to other techniques and alternative approaches. Usually, methods are evaluated using Line Integral Convolution (*LIC*) [15] images as well as the topology or quantitative discussion of the overall approximation error and compression rates. In particular, the influence of the flow field changes on the computation of common flow visualizations, especially in case of integration based methods, is hardly discussed. Often, no further visualizations have been presented on the compressed datasets.

In this paper, we investigate the applicability of Finite-Time Lyapunov Exponent (*FTLE*) computations on compressed vector fields. Therefore, we present a quantitative discussion of the influence of the approximation error on streamline integration. The underlying method can be used to characterize the error distribution of the integration results between different compression techniques. Finally, we contrast the quality of LIC and FTLE visualizations of three different compression approaches. From our results, we derive first indications how good FTLE images can be obtained from compressed fields.

2 Related Work

The main focus of our work is the comparison of the results of two FTLE variants on vector fields that were compressed using different techniques. In the following, we briefly review related work on vector field compression, as well as foundations of the FTLE computation.

2.1 Vector Field Compression

Most vector field clusterings in literature aim at the computation of simplified representations of a vector field, rather than to compress its dataset size. Lodha et al. [9] first extended the clustering approach of Telea and Wijk [16] in order to compute an error-bounded vector field compression that preserves the main characteristics of the stationary points. Later, Theisel et al. [17] define two topologies to be equivalent if they (1) have an identical set of stationary points, including the first derivative at their positions; (2) have the same set of boundary switch points; and (3) all separatrices start or end in the same stationary point, respectively enter or leave the domain in the same in- or outflow region. They use this definition to formulate a topology-aware compression algorithm. The implementation of this idea iteratively applies an edge collapse to simplify the underlying grid of the input field. Thereby, the corresponding changes of the vector field are tolerated, as long as the topology of the

compressed field remains equivalent to the one of the input field. To avoid a constant tracking of topology changes and to reduce the running time, they presented a local test that can predict, whether a local change of the vector field could lead to a change of the topological skeleton. Although the topology-aware compression algorithm of Theisel et al. [17] achieves very high compression rates, it does not provide any error threshold that limits the rate of change of the local flow behavior. This can lead to large distortions in the compressed vector field.

To limit the maximal error introduced by the grid simplification, Dey et al. [2] presented an approach that computes a Delaunay triangulation for every two-dimensional dataset and locally removes vertex after vertex, as long as a certain error threshold is not exceeded by the local vector field change. In contrast to Theisel et al., their algorithm does not guarantee the preservation of the vector fields topology.

In order to provide an error-bound that retains the main characteristic of the local flow and additionally preserves the vector field topology as one important flow feature, we presented a new compression algorithm [20] that is based on a region-wise linear approximation of the input field. Thereby, the maximal approximation error within each local linearization is bound by a user-defined error threshold E_{max}. Note to the reviewers: A detailed description of this algorithm can be found in the appendix.

2.2 Variants of Finite-time Lyapunov Exponents

A common visualization approach, which has been investigated in detail in the past years, is FTLE as presented by Haller [5]. An FTLE field shows the rate of separation, respectively convergence, of closely neighbored particles over the time in space. Since its introduction, the basic FTLE computation was extended, for instances, to compute the separation of flow from surfaces [4] up to the extraction of the efficient separation of surfaces in three-dimensional fields [14].

Besides the common *flow-map based FTLE* (*F-FTLE*) computation we will use another variant: The *Localized Finite-time Lyapunov Exponent* (*L-FTLE*) of Kasten et al. [7]. It uses a computation along path lines that only depends on the first derivative. An important advantage of this FTLE variant is that it is more robust with respect to noise and thus might be well suited for compressed fields.

3 Data Acquisition

We use different types of compression algorithms for our comparison. We focus on topology-preserving and error-bounded compression techniques and want to study the influence of these properties on visualizations of the compressed fields. Thereby, we want to compare the different compressed fields against each other as well as against the original, uncompressed vector field.

To facilitate these comparisons, we apply two approaches: a quantitative, statistical analysis of the quality of streamline integration, and a qualitative, visual analysis of visualization results. Both approaches are described in the following.

3.1 Quantitative Analysis of Streamline Integration

The goal of the quantitative analysis is to assess the impact of the compression on the quality of integrations in the compressed field. This facilitates the comparison to the original vector field and allows to obtain a quality measure for the compression technique.

For this evaluation, we want to compare the deviation of streamlines in the compressed field from streamlines in the original vector field. A natural concept that allows to derive the deviation values is the *flow map*. The flow map stores the end point after a certain integration time for every position in the field. To assure comparability between the original field and the compressed field, we use the grid points of the original field as seed positions for both fields. The resulting flow map is stored on the grid of the original vector field. So, the deviation of streamlines can then be derived as a scalar field that contains the Euclidean distance of the particle end points between the two flow maps. The resulting field shows how much the flow behavior differs locally between the compressed and the uncompressed vector field.

As we are interested in the minimal and maximal deviation, as well as the distribution of integration errors, we use box plots [19] and histograms to visualize the flow map differences. Thereby, we can juxtapose the deviation distribution for multiple integration times. In this paper, we focus on integration times that are suitable for FTLE computations.

3.2 Qualitative Visual Analysis

For a qualitative comparison of different compression techniques, we compute and visually compare LIC and FTLE images. Thereby, the LIC images allow us a comparison of the characteristic, topological flow patterns. To evaluate how well convergence and divergence of flow is preserved, we compare FTLE images of compressed fields to the FTLE images of the original vector field. From the flow maps, we compute F-FTLE fields in a first step. Thereby, we use streamline integration instead of path line integration, to apply FTLE on steady fields. As a second FTLE variant, we compute the L-FTLE fields. In contrast to the original implementation of Kasten et al. [7], we do not use a precomputed and interpolated Jacobian field. We use the constant derivative that is given for each triangle cell, respectively for each linearly approximated region.

4 Results

We now discuss the results of the flow map and the FTLE computations described in Sec. 3 on the basis of two real-world datasets: the Kármán vortex street and a two-dimensional simulation of a jet stream. In the following, we first describe the used dataset and then show and discuss the results of (1) a topology-preserving non-error-bounded, (2) an error-bounded non-topology-preserving, and (3) a topology-preserving and error-bounded compression algorithm.

4.1 Datasets

4.1.1 Kármán Vortex Street

The Kármán vortex street is a well-known flow pattern that shows the flow separation from an obstacle. In this dataset, the flow goes around a bar, which is rotated by 45 degrees towards the inflow. After passing the obstacle, the flow shows a characteristic periodic swirling. The original dataset is given on a triangulated grid with 156 842 cells. Figures 1(a) and 1(c) show a LIC image and the F-FTLE image of this vector field.

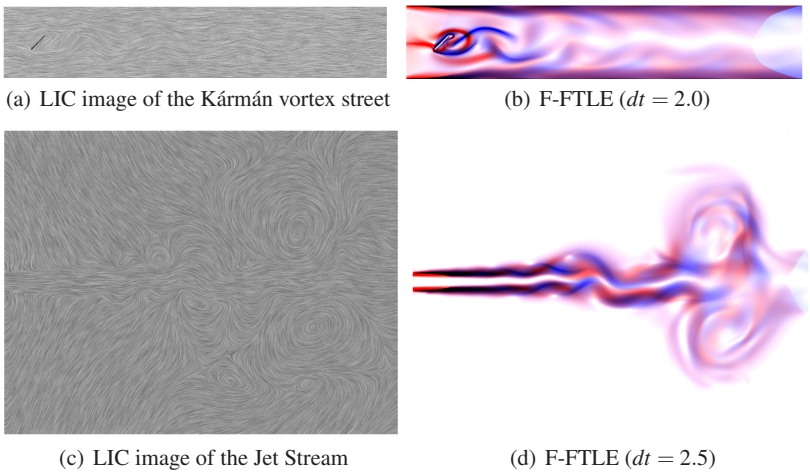

(a) LIC image of the Kármán vortex street (b) F-FTLE ($dt = 2.0$)

(c) LIC image of the Jet Stream (d) F-FTLE ($dt = 2.5$)

Fig. 1 LIC and F-FTLE images of the uncompressed, original Kármán vortex street (upper row) and the Jet Stream dataset (bottom row)

4.1.2 Jet Stream

The Jet Stream dataset shows a flow that enters the domain through a narrow opening on the left side at a velocity of Mach 0.1. This produces an expansion of the flow which results in a complex flow behavior with many vortices and small turbulent flow structures. The original vector field is given on a triangulated grid with 2922912 cells. Figures 1(c) and 1(d) show a LIC image and the F-FTLE image of this vector field, respectively.

4.2 Topology-preserving Compression without Error-bounds

As an example for a topology-preserving compression algorithm, we consider the algorithm of Theisel et al. [17] and achieve very high compression rates of 99.9% for the Kármán vortex street and 99.6% for the Jet Stream dataset.

First, we quantitatively analyze the integration errors. Fig. 2(a) shows box plots of the flow map differences for various integration times for the Kármán vortex street. The errors are distributed mostly uniformly, as can be seen from the underlying histograms. With increasing integration times the average as well as the maximal error increases nearly linearly. Given the dimensions of the vector field of 20 by 4, we have a deviation of about 10% of the domain diameter at an integration time of 2.5. Fig. 3(a) shows the results of the quantitative integration error analysis on the Jet Stream dataset. This fields has a dimension of 25 by 20. Therefore, the maximal error at an integration time of 2.5 is about 10% of the domain diameter. However, the overall distribution of the errors is not uniform, but concentrated near the minimum. Three quarters of the data points have an error of below 0.4.

For a qualitative analysis of the compression results, we consider the visual quality of LIC and F-FTLE images. Despite the high compression rate, both LIC images in Fig. 2(e) and Fig. 3(b) show that overall flow patterns are preserved very well. Especially in the Jet Stream dataset, the singularities, their region of influence as well as their interaction is clearly visible and the image is very similar to the LIC image of the uncompressed field (cf. Fig. 1). Only the inflow area in the Jet Stream dataset is distorted.

We see the reason for these good results in the preservation of the topological skeleton. When interpreting LIC images, the viewer focuses on well known flow patterns, e.g., laminar flow or characteristic linear flow in the vicinity of stationary points. Therefore, the perception of a LIC image is mainly sensitive to the existence and location of the stationary points. As these do not change during compression because of the algorithm design, the overall quality of the LIC images remains very good.

Figures 2(g) and 4(a) show the F-FTLE images of the Kármán vortex street and the Jet Stream respectively. In both cases, we see strong artifacts from the coarse underlying grid. Because Theisel et al. only aim at the construction of a topologically equivalent compressed vector field, they remove features that are not represented by

the topological skeleton. Thus, the typical flow separation and convergence of the Kármán vortex street has been removed as it can not be seen in Fig. 2(g). Only close to the bar and the topological skeleton, some features are still visible, but a distortion is also noticeable in those areas. The FTLE results of the Jet Stream dataset suffers from similar problems. By comparing Figures 1(a) and 3(b), one can see that the compression causes a shift of the inlet towards the lower left domain boundary. Therefore, the final FTLE image is not only too coarse to see small features but also shows high FTLE values in regions of nearly stagnating flow in the original field.

FTLE images show divergence and convergence in the flow and thus are highly sensitive to deviations in direction and magnitude of the vector field. Topology-preservation only guarantees a preservation of the overall flow behavior in the individual basins, but it does not necessarily preserve the course of the streamlines. High deviations in the flow direction are still possible without altering the topology. Therefore, the high compression rate has a significant impact on the quality of FTLE images.

4.3 Error-bounded Compression without Topology-Preservation

Since FTLE computation on topology-preserved vector fields is affected by large artifacts from the underlying coarsened grid, we study the results of the error-bounded vector fields approximation algorithm that is presented in [20]. In order to study only the influence of the error-bound E_{max} on the vector field approximation, we removed the topology-preservation from the algorithm. This approach achieves compression rates of 97.8 % for the vortex street and 99.2 % for the Jet Stream dataset that we used in Figures 2 and 3. Thereby, an error-bound of 0.1 is used for the Kármán vortex street, respectively 0.4 for the Jet stream dataset.

Again, we first quantitatively analyze the integration error-based on the flow map differences in the original and the approximated vector fields. The results in Figures 2(b) and 3(c) show strong differences to the results of the previous section. On both fields, the maximal error is only about 1 % of the field dimensions. On both fields, we see an over-linear growth of the maximal error with integration time. However, the average error grows much slower. When considering the distribution of the error values, we see that the range of possible errors is dominated by outliers and the majority of errors have a small magnitude. From the associated LIC images (cf. Fig. 2(f) and Fig. 3(d)), we can see that these outliers are concentrated around the main vortices and between the inflow and the neighbored nearly stagnating flow. We expect only small integration errors even with higher integration times, though we know that some artifacts will be present due to the outliers.

The LIC images in Figures 2(f) and 3(d) verify our assumptions. In case of the Kármán vortex street, the typically swirling flow behind the rotated bar is better preserved then in Fig. 2(e). This is due to the fact that the error-bounded algorithm preserves a higher number of cells (cf. Figures. 2(c) and 2(d)), which allows a more detailed approximation of the original field while decreasing the compression rate.

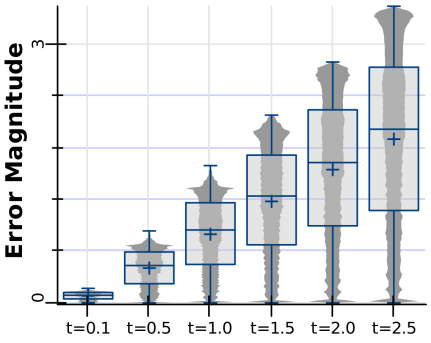

(a) Error distribution (topology-preserving)

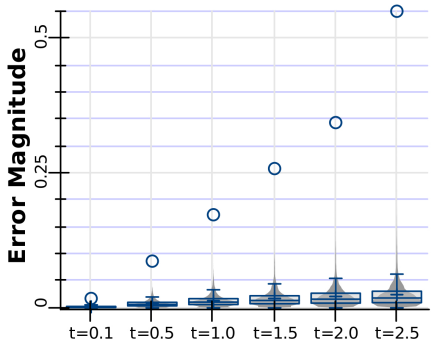

(b) Error distribution (error-bounded)

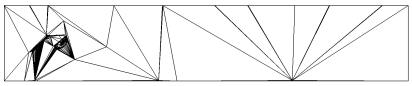

(c) Compressed grid (topology-preserving)

(d) Compressed grid (error-bounded, $E_{max} = 0.4$)

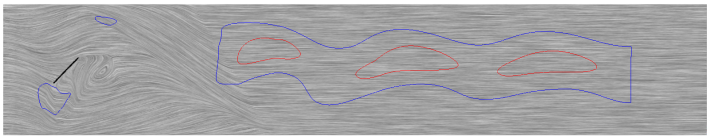

(e) LIC of compressed vector field (topology-preserving)

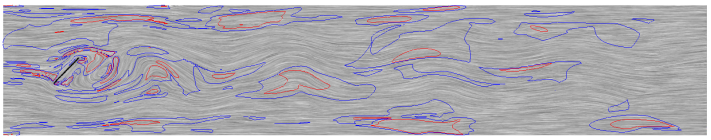

(f) LIC of compressed vector field (error-bounded, $E_{max} = 0.4$)

(g) F-FTLE of compressed vector field (topology-preserving, $dt = 2.0$)

(h) F-FTLE of compressed vector field (error-bounded, $E_{max} = 0.4$, $dt = 2.0$)

Fig. 2 In order to compare the integration results of the first two compression techniques, Figures 2(a) and 2(b) show the histogram of the error values in flow map differences for multiple integration durations. The box plots mark their quartiles, the plus their average and the circle their maximal value. Figures 2(e) to 2(h) show LIC images and F-FTLE images of the compressed vector fields. To analyze the distribution of the integration error in the compressed field (for an integration time of 2.0, we added iso-lines to emphasize the upper 25% (*blue*) and the upper 5%(*red*) of all error values in the domain.—Please note the different scale of the ordinate of the shown error plots. Although the areas of the higher approximation error seem to cover large parts of the domain, the overall error in Fig. 2(f) is much smaller compared to Fig. 2(e).

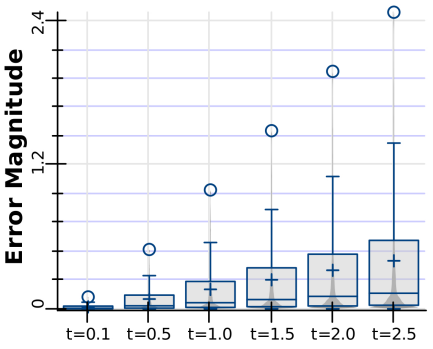

(a) Error distribution (topology-preserving)

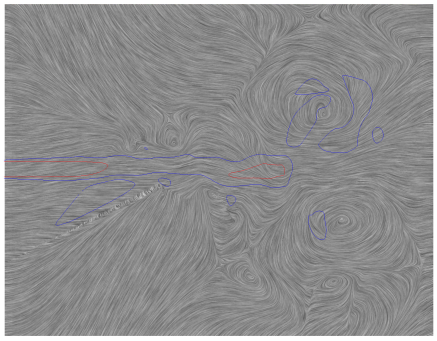

(b) LIC of compressed vector field (topology-preserving)

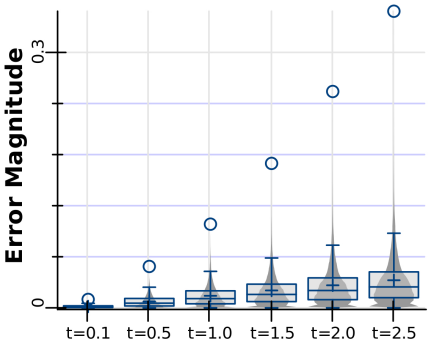

(c) Error distribution (error-bounded)

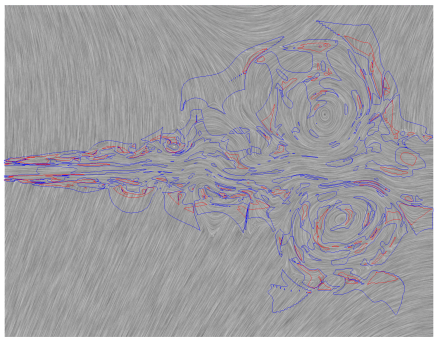

(d) LIC of compressed vector field (error-bounded, $E_{max} = 0.1$)

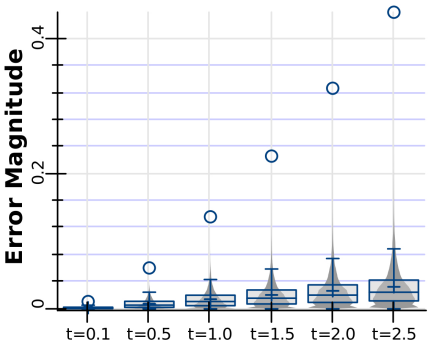

(e) Error distribution (topology-preserving and error-bounded)

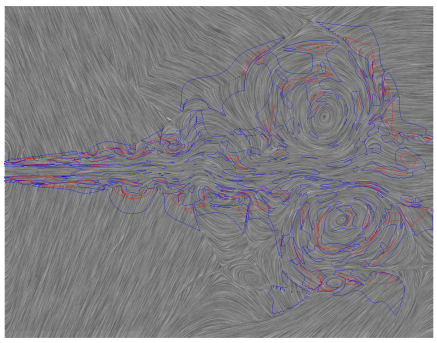

(f) LIC of compressed vector field (topology-preserving and error-bounded, $E_{max} = 0.1$)

Fig. 3 In order to compare the integration results of the used compression techniques, Figures 3(a), 3(c) and 3(e) show the histogram of the flow map differences for multiple integration times. The box plots mark their quartiles, the plus their average and the circle their maximal value. Figures 3(b), 3(d) and 3(f) show the LIC images the compressed vector fields. To analyze the distribution of the integration error in the compressed field (for an integration time of 2.5, we added iso-lines to emphasize the upper 25% (*blue*) and the upper 5% (*red*) of all error values in the domain.

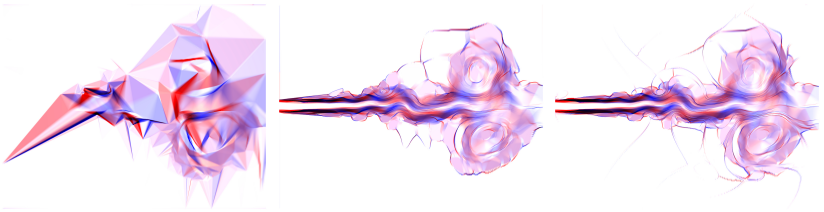

(a) topology-preserving (b) error-bounded (c) topology-preserving and error-bounded

Fig. 4 These three images show the F-FTLE results for the Jet Stream dataset. The results were computed for all compression techniques that are used in this work. Thereby, an integration time of 2.5 and an error-bound of $E_{max} = 0.1$ was used.

But in contrast to the simple flow characteristic of the vortex street, which is well preserved, the approximation of the Jet Stream shows strong differences. Figure 3(d) shows that the inflow as well as the main vortices are correctly approximated. But since the topology is very sensitive against vector field perturbations, one can clearly see changes of the vector fields topology. Especially in the lower part of the field, the weaker vortices begin to vanish at an error threshold of 0.1. So, depending on the chosen error threshold, the topology can change significantly and LIC images can no longer give reliable information of the main global flow behavior.

The FTLE results in Figures 2(h) and 4(b) are very similar to the corresponding FTLE images of the original fields in Fig. 1. All main features of the original FTLE fields are preserved. Especially in case of the Kármán vortex street, the characteristic flow separation and attachment behind the obstacle are very well perceptible. This is due to the fact that the local error threshold limits the allowed magnitude and angle changes that are introduced by the compression approach. These findings are also indicated by the small error magnitudes that we find in the flow map differences (Figures 2(b) and 3(c)). We did not expect this result because vanished saddle and center points should result in much higher integration errors theoretically. On the other hand, there are also clearly visible artifacts, especially in the case of the Jet Stream. These are probably caused by the outliers that are visible in the error plots. The position of the artifacts correspond to the areas of maximal error in the LIC images (Figures 2(f) and 3(d)). These areas also correspond to the boundaries of the linearly approximated regions that are used by the algorithm [20], which, by the design of the algorithm, tend to have the largest errors.

4.4 Error-bounded Topology-preserving Compression

As a synthesis from the previous sections, we apply the original approximation algorithm of [20] to the datasets. This algorithm combines the topology-preservation approach of Theisel et al. with an error-bound. We achieve compression rates between 92.3 % (for $E_{max} = 0.1$) to 94.8 % (for $E_{max} = 0.4$) for the Kármán vortex street and

compression rates between 99.5 % (for $E_{max} = 0.02$) to 96.0 % (for $E_{max} = 0.1$) for the Jet Stream dataset.

Similar to the error-bounded approach without the topology-preservation, the error distributions in Fig. 3(e) show fast increasing error outliers and a slow increasing average error over integration time. Also the location and the extent of the regions with the highest error are similar. Thus, the topology-preservation has no further influence on the error distribution in this case. However, the LIC images clearly show that all features of Fig. 1(c) could be preserved. The inflow could be correctly approximated as well as all topological features. Also the FTLE fields in Figures 4(b) and 4(c) look nearly the same. This is due to the fact that the same error measure was used, which locally limits the magnitude and angle differences between the approximated the original field.

We use this compression method to study FTLE on compressed fields in more detail. First, we investigate the influence of the integration times. Figure 5 shows a sequence of FTLE results of the Kármán vortex street for different integration times. Counterintuitively, longer integration times lead to smoother FTLE images. Whereas the highest integration time theoretically leads to the most error-prone image, these images also contain nearly no visual artifacts of the underlying coarse region-wise linearly approximated vector field. Compared to that, the smaller integration times are more precise and lead to a coarser looking result, since shorter integration times emphasize the transitions between to neighbored linearizations. Given the overall small integration error, we can still consider the results with high integration times reliable.

Second, we applied the L-FTLE algorithm to the compressed fields. L-FTLE is expected to compensate discontinuities at cell or linearly approximated region boundaries. The results for the Jet Stream dataset are shown in Fig. 6. While on the original field, L-FTLE leads to sharper images (cf. Figures 1(d) and 6(c)), on the compressed fields it appears to emphasize artifacts. These outweigh the flow features, so that the images cannot be interpreted correctly anymore. We assume that this is a consequence of the coarse, piece-wise linear vector field approximation, especially the piece-wise constant Jacobian.

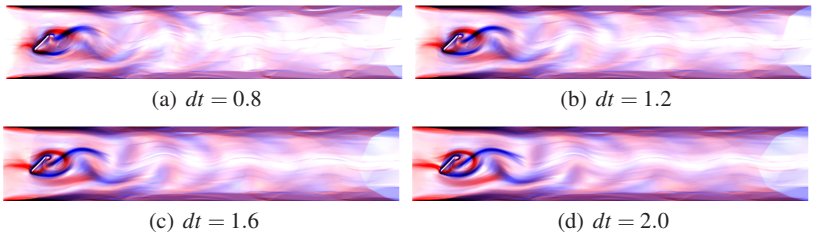

(a) $dt = 0.8$ (b) $dt = 1.2$

(c) $dt = 1.6$ (d) $dt = 2.0$

Fig. 5 F-FTLE images of our error-bounded, compressed vector field (the topology-preserving and error-bounded algorithm was used with $E_{max} = 0.1$) for different integration lengths. The main flow features can clearly be seen in all four images. However, the result appears to be much smoother with longer integration times.

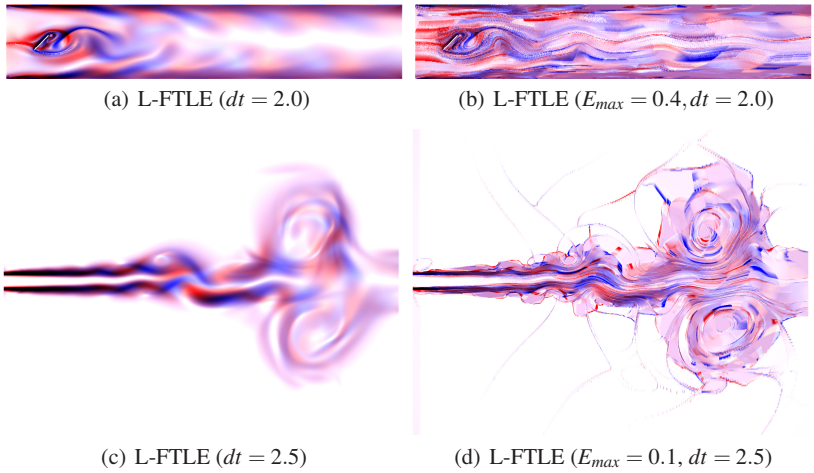

(a) L-FTLE ($dt = 2.0$) (b) L-FTLE ($E_{max} = 0.4, dt = 2.0$)

(c) L-FTLE ($dt = 2.5$) (d) L-FTLE ($E_{max} = 0.1, dt = 2.5$)

Fig. 6 These two images show a comparison of the L-FTLE results computed on the uncompressed Jet Stream dataset (left) and its error-bound compression (right). The result of the compressed field clearly shows artifacts of the region-wise approximation.

5 Discussion

We found that the two criteria, preservation of topology and a bounded approximation error, have a different impact on the resulting compressed field with respect to the achieved compression rate as well as on the preserved flow features. The best compression rates can be achieved when no error-bound is used. This is shown by the results of Theisel et al. [17]. The introduction of an error-bound has a large negative impact on the compression rate. The worst compression rates were achieved when fulfilling both criteria.

Depending on the flow features that should be preserved by the compression, the different compression algorithms have their advantages and disadvantages. In order to preserve the topology, either the topology has to be preserved explicitly, or an appropriately small approximation error has to be used (cf. Dey et al. [2]). When the compressed fields are mainly used for the generation of LIC images, the preservation of topology seems to be the more important compression criterion. For non-topological flow features, such as convergence and divergence, a suitable error-measure should be used, e.g., one that limits the deviation of flow direction.

The integration error that results from the compression not only depends on the vector field itself but on the design of the particular compression algorithm. Error-bounded algorithms from [20] show a characteristic small number of high error outliers, which are located at the transitions between the linearly approximated regions.

In every case, the integration error increases with longer integration times, because the error accumulates along the streamlines. On the other hand, longer integration times lead to smoother FTLE images with less visual artifacts in our exam-

ples. Since the overall error is higher in these images, they are also less reliable. Therefore, the integration time is not only is one of the most important parameters of the FTLE computation, here it also controls the trade-off between the visual appeal of the images and their correctness. As long as the integration error in the flow maps remains in the scale of the cell resolution, a meaningful preview of the original dataset can be obtained, as shown in the examples of Sec. 4.4.

6 Conclusion and Future Work

In this work we compared three different vector field compression algorithms with respect to their applicability in vector field visualization. One method with an unbounded approximation error results in very high compression rates while preserving the vector fields topology. The others limit the local vector field deviation—and also the compression rate—by using a user-defined error threshold. In particular, we compared the quality of the flow map computation in uncompressed and compressed vector fields as well as the results of FTLE computations.

Thereby we showed, that preservation of the topological skeleton is not sufficient to preserve flow features. Additionally, a mechanism to control the overall approximation error is needed. The quality of the integration results as well as the computed FTLE images of compressed vector fields largely depends on the used error-bounds, respectively the compression rates. Given a compressed field, the integration time not only controls the quality of the FTLE results, but also influences the error distribution in the resulting field. We showed this by the example of two real-world datasets.

For future work, more datasets and vector field compression techniques could be investigated to extend the comparison of the different algorithm approaches. Thereby, an interesting question is, whether there is a characteristic development of the error distribution with increasing integration times for a particular compression method and which properties of the method are causing it. Such research appears to be possible, since we have seen characteristic error distributions for the individual algorithms. Further, it could be possible to develop a guideline that helps to decide, which compression techniques can be used to facilitate certain visualizations or analysis tasks on compressed vector fields.

Acknowledgements We thank Markus Rütten, Guillaume Daviller, and Bernd Noack for providing the simulation datasets. Special thanks go to the FAnToM development group for providing the visualization software. We also thank Jens Kasten for the fruitful discussions. This work was partially supported by the European Social Fund (Appl. No. 100098251).

References

1. Auer, C., Kasten, J., Kratz, A., Zhang, E., Hotz, I.: Automatic, tensor-guided illustrative vector field visualization. In: Visualization Symposium (PacificVis), 2013 IEEE Pacific, pp. 265–272 (2013). DOI 10.1109/PacificVis.2013.6596154
2. Dey, T., Levine, J., Wenger, R.: A delaunay simplification algorithm for vector fields. In: Computer Graphics and Applications, 2007. PG '07. 15th Pacific Conference on, pp. 281–290 (2007). DOI 10.1109/PG.2007.34
3. Garcke, H., Preußer, T., Rumpf, M., Telea, A., Weikard, U., Van Wijk, J.: A continuous clustering method for vector fields. In: Proceedings Visualization 2000, pp. 351–358. IEEE Computer Society (2000)
4. Garth, C., Wiebel, A., Tricoche, X., Joy, K., Scheuermann, G.: Lagrangian visualization of flow-embedded surface structures. Computer Graphics Forum **27**(3), 1007–1014 (2008)
5. Haller, G.: Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence. Physics of Fluids (1994-present) **13**(11), 3365–3385 (2001). DOI http://dx.doi.org/10.1063/1.1403336
6. Helman, J.L., Hesselink, L.: Representation and Display of Vector Field Topology in Fluid Flow Data Sets. IEEE Computer **22**(8), 27–36 (1989)
7. Kasten, J., Petz, C., Hotz, I., Noack, B.R., Hege, H.C.: Localized finite-time lyapunov exponent for unsteady flow analysis. In: M. Magnor, B. Rosenhahn, H. Theisel (eds.) Vision Modeling and Visualization, vol. 1, pp. 265–274 (2009)
8. Laramee, R., Hauser, H., Zhao, L., Post, F.: Topology-based flow visualization, the state of the art. In: Topology-based Methods in Visualization, Mathematics and Visualization, pp. 1–19. Springer Berlin Heidelberg (2007)
9. Lodha, S.K., Renteria, J.C., Roskin, K.M.: Topology preserving compression of 2d vector fields. In: Proceedings of the Conference on Visualization '00, VIS '00, pp. 343–350. IEEE Computer Society Press, Los Alamitos, CA, USA (2000)
10. Lu, K., Chaudhuri, A., Lee, T.Y., Shen, H.W., Wong, P.C.: Exploring vector fields with distribution-based streamline analysis. In: Visualization Symposium (PacificVis), 2013 IEEE Pacific, pp. 257–264 (2013). DOI 10.1109/PacificVis.2013.6596153
11. Marchesin, S., Chen, C.K., Ho, C., Ma, K.L.: View-dependent streamlines for 3d vector fields. Visualization and Computer Graphics, IEEE Transactions on **16**(6), 1578–1586 (2010). DOI 10.1109/TVCG.2010.212
12. McLoughlin, T., Laramee, R.S., Peikert, R., Post, F.H., Chen, M.: Over two decades of integration-based, geometric flow visualization. Computer Graphics Forum **29**(6), 1807–1829 (2010)
13. Post, F.H., Vrolijk, B., Hauser, H., Laramee, R.S., Doleisch, H.: The State of the Art in Flow Visualization: Feature Extraction and Tracking. Computer Graphics Forum **22**(4), 775–792 (2003)
14. Sadlo, F., Peikert, R.: Efficient visualization of lagrangian coherent structures by filtered amr ridge extraction. IEEE Trans. Vis. Comput. Graph. **13**(6), 1456–1463 (2007)
15. Stalling, D., Hege, H.C.: Fast and resolution-independent line integral convolution. In: Proceedings of SIGGRAPH '95, pp. 249–256. Computer Graphics Annual Conference Series, ACM SIGGRAPH, Los Angeles, California (1995). DOI 10.1145/218380.218448
16. Telea, A., Van Wijk, J.: Simplified representation of vector fields. In: Visualization '99. Proceedings, pp. 35–507 (1999). DOI 10.1109/VISUAL.1999.809865
17. Theisel, H., Rössl, C., Seidel, H.P.: Compression of 2d vector fields under guaranteed topology preservation. Computer Graphics Forum **22**(3), 333–342 (2003)
18. Tricoche, X., Wischgoll, T., Scheuermann, G., Hagen, H.: Topology Tracking for the Visualization of Time-Dependent Two-Dimensional Flows. Computers & Graphics **26**(2), 249–257 (2002)
19. Tukey, J.: Exploratory Data Analysis. Addison-Wesley series in behavioral sciences. Addison-Wesley Publishing Company (1977)
20. X: Vector field approximation using linear neighborhoods. The paper is currently in submission. We provide a pre-print in the supplementary material.