Citation: Martinez, Q., Jagadeesh, C., Manolesos, M. & Omidyeganeh, M. (2026). Leading-edge vortex monitoring in dynamically stalled flows via persistent homology. Computers & Fluids, 306, 106931. doi: 10.1016/j.compfluid.2025.106931

This is the published version of the paper.

This version of the publication may differ from the final published version.

Computers and Fluids 306 (2026) 106931

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

Leading-edge vortex monitoring in dynamically stalled flows via persistent homology

Quentin Martinez [a,*], Chetan Jagadeesh [a], Marinos Manolesos [b], Mohammad Omidyeganeh [a,*]

[a] City St Georges, University of London, School of Science and Technology Northampton Square, London EC1V 0HB, UK
[b] National Technical University of Athens, School of Mechanical Engineering, Iroon Politechniou 9, 15780 Zografou Campus, Greece

ARTICLE INFO

Keywords:
Stall flutter
Dynamic stall
Vortex tracking
Persistent homology
Leading-Edge Vortex
Unsteady aerodynamics

ABSTRACT

A novel vortex core identification pipeline is developed based on template matching. Using persistent homology, a template similarity field is constructed from a sliding window template-target feature space distance. This scalar field is then used to accentuate localised regions of spanwise vorticity via nonlinear weighting. This method is successfully applied to track the leading-edge vortex trajectory in a stall flutter starting cycle for a pitching NACA 63(3)418 aerofoil. Trajectory results are compared with several user-based vortex core identifiers like local vorticity minimum, local Q-criterion maximum, local swirling strength maximum, and manual tracking. The results of this comparison are quite satisfactory as the developed method is capable of automatically monitoring the leading-edge vortex core through several critical stages of its lifecycle. The effects of template size and down sampling are also investigated with respect to the vortex core identification. It is found that a template radius of $r = 0.04c$ and down sampling factor $M = 10$ are sufficient for accurate vortex core monitoring in dynamically stalled flows. In general, this method acts primarily as a field-based filter that can be useful for isolating highly vortical regions like the leading-edge vortex core in stall flutter or dynamic stall scenarios.

1. Introduction

Tracking and monitoring of large-scale coherent structures in fluid flows is a ubiquitous task across various engineering and scientific disciplines. This is particularly true for vortex dominated flow scenarios like aerofoil ramp-up motion [1], stall flutter [2], and dynamic stall [3]. In these cases, a strong leading-edge vortex (LEV) is generated due to the traversal of a separation point upstream followed by the subsequent roll up of the leading-edge shear layer [4]. This particular vortex structure has been linked to characteristics such as unsteady lift augmentation, negative aerodynamic damping, and load hysteresis which can be regarded as favourable or unfavourable depending on the design context. For example, the LEV induced lift augmentation is exploited by birds during perching scenarios to maintain lift at high incidence angles [5]. For high aspect ratio aeroelastic structures though, this unsteady loading can cause fatigue and early wear. Within the wind energy community, many researchers aim to quantify the status of this LEV over a range of parametric conditions to develop a better understanding of its associated flow physics and effects on wind turbine blade integrity. To do this, it is common to measure the LEV core trajectory using metrics such as the Q-criterion, swirl criterion, and vorticity, to name a few [1,2,6,7] . Though, many Eulerian based definitions often suffer from spurious vortex identification in high Reynolds number flows that can make vortex monitoring highly challenging and case specific. Often, time-resolved monitoring of a particular vortex of interest requires an extended pipeline including methods such as thresholding, filtering, or clustering.

Several recent advances in the field of vortex core identification/monitoring have implemented tools like DBSCAN clustering [8], supervised machine learning [6], and computer vision-based techniques [9] to successfully identify and monitor vortices. A relatively niche solution space to this objective is by using template or pattern matching. In this case, a template is either identified or generated by the user that encompasses the approximate composition of a target. This type of pattern recognition task has roots in image-based segmentation/classification and is often quantified through a distance or similarity field

---

* Corresponding Authors.
E-mail addresses: martinezq95@gmail.com (Q. Martinez), Omid.Yeganeh@citystgeorges.ac.uk (M. Omidyeganeh).

https://doi.org/10.1016/j.compfluid.2025.106931
Received 16 September 2025; Received in revised form 31 October 2025; Accepted 27 November 2025

metric. Though, a similar objective can be applied to fluid flows where a template is provided to identify certain fluid dynamic events. Ebling and Scheuermann [10], Rodrigues et al. [11], and Heiberg et al. [12] all utilised vector valued templates corresponding to various scenarios like shearing, swirling, sinks, etc. to successfully identify and label fluid structures within a wide range of experimental and numerical case studies like jets, blood flow, and unsteady vortex shedding. This method requires the generation of a template library that can be applied individually or in superposition to compare with real world flows. In all cases, the convolution operation was utilised to generate a similarity measure between the template and target. One potential limitation of vector valued templates is that grid uniformity is required for the convolution operation, ideally with a comparable template-target size [13]. Second, template rotation is typically applied at each sliding window which can be computationally expensive depending on the angle resolution. Another template-based method was introduced by Elsinga et al. [14] that utilised a conditional eddy template of the average flow associated with a spanwise swirling event for a turbulent boundary layer via linear stochastic estimation (LSE). A distance map was then constructed using the cross correlation between this template and the spanwise swirling strength distribution. Overmeyer [15] utilised an adaptive template cross-correlation method to automatically monitor rotor tip vortices in their stereo image pairs. In all cases, the success of template-target pairing methods is inherently based on how well the template can objectively capture the target structure whilst minimising the computational cost.

In this work, we propose and demonstrate the use of a topological data analysis (TDA) informed template matching pipeline for tracking the LEV core trajectory in stall flutter simulations. The use of TDA within the engineering community has seen recent growth, granting a refreshing perspective on fluid flow analysis and modelling techniques. Suzuki et al. [16] utilised persistent homology to estimate the permeability of fracture networks in rocks to estimate flow characteristics without the need for fluid dynamic simulations. Likewise, Moon et al., [17] were able to infer transport relevant properties of their porous media samples such as the permeability, tortuosity, and anisotropy via statistical learning of vectorized persistence diagrams. Tymochko et al. [18] leveraged sublevel persistent homology to accurately monitor the diurnal cycle of a tropical cyclone using satellite images. Smith et al. [19] defined a persistent homology informed loss function for modelling dynamic stall events in a pitching flat plate. Specifically, it was shown that the state trajectory for a pitching flat plate is approximately homeomorphic to a unit circle. This is was shown to be useful as a learning constraint for data-driven modelling and also as a simplified geometric interpretation of the dynamic stall phenomenology. Chen and Lin [20], utilized persistent path homology to isolate and track vector field singularities in tropical storms. Different from traditional applications of point-cloud persistent homology, this method first requires the construction of a two-dimensional directed graph (digraph) for building a path-based complex rather than a simplicial one. This type of data structure encodes directional relationships rather than high dimensional geometric proximity. Furthermore, it was noted that the effects of graph resolution on isolating singular events in high-fidelity vector fields is not yet fully understood.

In this work, we utilise persistent homology to quantify the feature space similarity between a template-target pair via construction of a simplicial complex. The logic behind this strategy is that the LEV can have a wide range of strengths, sizes, and shapes owing to various dynamic stall related events. This makes direct 'image-like' pairing operations like convolution or RMSE potentially sensitive to template noise, orientation, and its absolute distance. Conversely, persistent homology provides a mechanism for mapping high-dimensional flow field data to a common low dimensional feature space that is both resilient to noise and invariant to template/target rotations. This abstraction also does not require dimensionality matching or grid uniformity for a direct template-target comparison which allows for applications in unstructured and non-uniform CFD domains. Finally, this method does not require a large template library to be generated, only a single user-input sample. To summarize, the primary difference between our proposed method and current template-matching techniques is that our template-target comparison is computed in a topological feature space that encodes the vortex core, rather than relying on vector or spatial correspondences. This abstraction makes it simple to implement in complex grids and minimises the amount of pre-processing overhead. To the best of the author's knowledge, the vortex identification pipeline demonstrated here is the first application of persistent homology for direct core identification via simplicial complex construction. It is for this reason that a comparison with established vortex core identification methods like the $Q-$ criterion [21], vorticity, swirling strength [22] and manual identification are utilised to validate its accuracy.

2. Setup

2.1. Numerical data set

Large eddy simulation data for a pitching NACA 63(3)418 aerofoil at a chord-based Reynolds number of Rec=100,000, reduced frequency $K = 0.4$, and mean incidence angle, $a = 15^\circ$, is used for the current investigation. A very fine, C-type grid with cell spacing 1761 × 300 × 81, in the chord-normal, chord-wise, and span-wise directions was generated using Pointwise to sufficiently resolve the aerofoil boundary layer. Operating in a rotating reference frame, the pitching motion of the aerofoil is modelled by a forced spring mass damper system with nondimensional mass moment of inertia, damping, and spring coefficients set to j*=0.0765 , u*=1000, and k*=0.0490.

The pitching moment forcing function is computed in time about the 0.5c location by integrating over the pressure and viscous forces along the aerofoil surface. See Martinez et al. [23] for a more detailed composition of the simulation setup conditions and grid convergence study. Starting from an initially static condition, this computational data set consists of field snapshots spanning a segmented window within the first aeroelastic pitching cycle indicated by the black dashed lines in Fig. 1. This window captures various dynamic stall events including the chordwise LEV migration, its detachment, and shedding behaviour at a temporal resolution of $\Delta t = 0.05\frac{t\ U_\infty}{c}$. Indeed, these events contribute to highly nonlinear aerodynamic loading during stall flutter starting cycles and transients. The dynamic status of an LEV, which includes traits like its size, shape, and relative position, are known to be critical indicators of this influence. In general, the current literature attempts to consolidate this departure from nominal static values based on various stages of the LEV-lifecycle [24]. Thus, dynamically stalled flow during a high angle of attack aeroelastic starting cycle is a relevant and important case study for of our proposed vortex identification method.

2.2. Persistent homology

In the field of topological data analysis, persistent homology is a tool

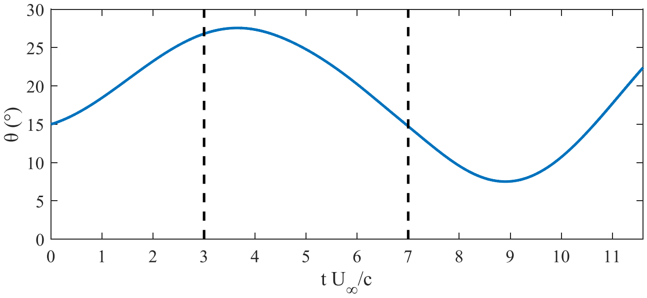

Fig. 1. Aeroelastic starting cycle in pitch degree of freedom. Black dashed lines indicate temporal limits where field data is sourced from.

for identifying the homology groups in a multivariate point cloud [19, 25]. This allows for a diagnosis of the data ‘shape’ that can be useful for modelling, classification, or segmentation purposes. This abstraction can be understood graphically through a persistence diagram which is constructed by monitoring the persistence, $p = \text{birth} - \text{death}$, of underlying topological features in the point cloud while varying the filtration level, ϵ. More rigorously, persistent homology is formulated using simplicial complexes. A simplicial complex is a collection of k-simplices such as vertices (0-simplices), edges (1-simplices), triangles (2-simplices), and their higher-dimensional analogues that are connected along shared faces. k-cycles, which are the subcomponents of this simplicial complex forming undirected closed loops, can be grouped into equivalence classes called homology groups. These homology groups formalize the number of holes in a given dimension of the complex by Eq. (1), where Zn is the number of k-cycles and im (∂_{k+1}) denotes the image of the boundary operator.

$$H_k = \frac{Z_k}{\text{im}(\partial_{k+1})} \tag{1}$$

In other words, H_k encodes the essential k-dimensional features of the simplicial complex by counting cycles that are not themselves boundaries of higher-dimensional simplices. Persistent homology applies this framework to point cloud data by building a nested sequence of simplicial complexes (filtration) indexed by a scale parameter and by tracking when elements of H_k appear (birth) and disappear (death). Here, an example of this construction is given canonically for the case of 2D point-cloud with noise.

First, a Vietoris-Rips complex is built using the Ripser Python library [26]. This procedure consists of initialising spheres of radius ϵ at each vertex in the point cloud. These spheres gradually grow, allowing for intersections with neighbouring spheres. The intersection of two spheres forms an edge between them. More generally, for $k + 1$ intersecting spheres, a k-dimensional simplex is formed between the corresponding vertices. A representation of the Vietoris-Rips complex is shown in Fig. 2 to demonstrate the vertex connectivity at an arbitrary filtration level, $\epsilon = 0.2$,.

At each filtration level, ϵ_i, the homology groups H_k are computed for the corresponding complex K_ϵ. The persistence of these homology groups is monitored as birth/death coordinate pairs and plotted as a diagram to indicate long-lived k-dimensional holes within the data set. This is shown in Fig. 3 corresponding to the point cloud data in Fig. 2. In this case, a highly persistent 1-dimensional hole or loop exists within the data set as indicated by the isolated H1 data point. This point implicates that a loop is born at a filtration level, $\epsilon = 0.22$ and dies at $\epsilon = 0.96$ which is the longest-lived feature in this 2D point cloud. This example serves as a simplified demonstration of persistent homology but follows the same intuition when dealing with high dimensional data associated with computational fluid dynamic (CFD) simulations.

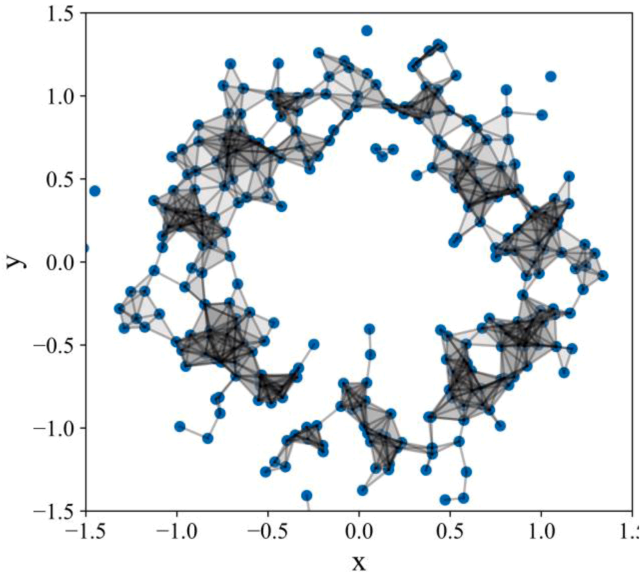

Fig. 2. Vietoris-Rips complex for 2D circular point cloud with noise at filtration level, $\epsilon = 0.2$.

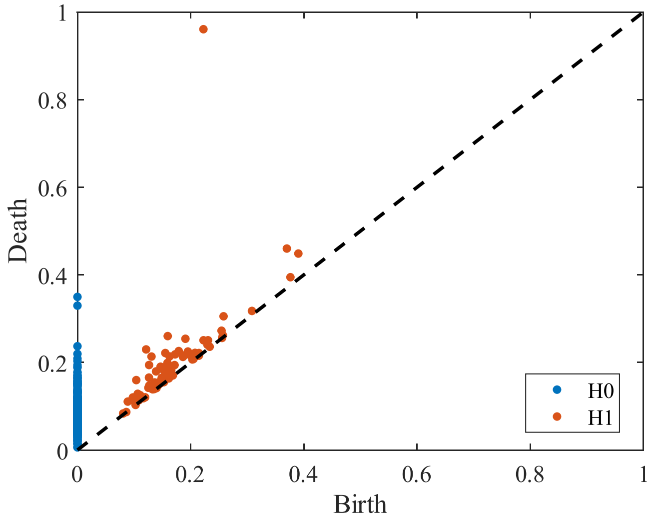

Fig. 3. Persistence diagram for point cloud data in Fig. 2 up to and including homology dimension $k = 1$.

2.3. Vortex template matching

In this section, we outline a novel post-processing pipeline for tracking the LEV core in stall flutter simulations. Starting from a developed LEV status, computed velocity components, pressure and spanwise vorticity field data u, v, p and ω_z are manually sourced within a circular region of interest (ROI) with radius r that approximately encloses the vortical structure core. In this case, the template is defined at the location, $x, y = (0.8024c,\ 0.3509c)$, with a radius, $r = 0.04c$. These sourced field variables are used to describe the LEV in R^4.

Next, we compute the persistence diagram for this high-dimensional point cloud up to and including the homology dimension, $k = 1$. Fig. 5 demonstrates the results of this computation given by the blue and orange dots ($H0 - t_0$, $H1 - t_0$), indicating a variety of topological features ranging between persistent holes and short-lived artifacts. An example of the manually identified LEV persistence diagram at a latter timestep, t_1, is also overlayed as the yellow and purple dots ($H0 - t_1$, $H1 - t_1$) in Fig. 5 and exhibits similar topological features. This abstraction serves as a single possible LEV template-target pairing between t_0 and t_1. The intuition is that this feature space pairing can be utilised to quantify the similarity between the LEV at two different instances. To do this, a second ROI of radius r scans over sampling points within the computational domain and computes the persistence diagram. In practice, it is not necessary to scan over every point as this is computationally expensive, and the attached LEV typically remains within one chord length of the aerofoil suction surface. In this work, our interrogation region (IR) is defined with radius, R, and has the same centroid as the template location.

Grid points outside of R are automatically assigned a maximum distance value to enforce incompatibility. The distance between these two persistence diagrams can be quantified by the Wasserstein distance [27] which is also known as the Earth mover’s distance [28]. Utilising the persim python library which is based on work by Adams et al. [29] the Wasserstein distance between two persistence diagrams is computed up to and including a homology dimension, $k = 1$. Mathematically, the p-Wasserstein distance between two persistence diagrams, D_1 and D_2, is given by Eq. (2) where p is the order of the Wasserstein distance. Here, the default value of $p = 2$ is utilised and φ is a bijection from D_1 to D_2. This operation attempts to find the best possible matching between the birth-death multisets by minimising the L_∞ norm of x and $\varphi(x)$ over all possible bijections.

$$W_p(D_1,\ D_2) = \left(\underset{\varphi: D_1 \to D_2}{inf} \sum_{x \epsilon D_1} ||x - \varphi(x)||^p \right)^{1/p} \tag{2}$$

This operation, applied to individually to the n^{th} grid point within an IR of N possible candidate locations, generates a template similarity field (TSF) indicating where the template-target distance is lowest at t_1. An example of this is shown in Fig. 6 where the ROI radius between two arbitrary but successive time steps is set to $r = 0.04c$. The effects of r will be discussed in the following sections. The template-target similarity at t_1 tends to be the highest in regions of high shear or vortex formation as indicated by the blue regions. Conversely, the red coloured regions indicate template-target dissimilarity which appears to be present in the regions immediately bounding strong fluid shearing or coherence. The highest values of W_p are situated about the leading-edge where vortex elements have not yet wholly formed.

Finally, this heatmap is then used to weight the out-of-plane vorticity field, ω_z, at each sampled grid point by Eq. (3). Here, τ is a small positive number to avoid division by zero.

$$\omega_{z,\text{scaled}} = \frac{1}{(W+\tau)^2} * \omega_z \tag{3}$$

This operation acts as both a form of error correction and filtration for regions of template dissimilarity. The LEV core location can be recovered as the local vorticity minimum in the TSF weighted vorticity field via thresholding. An example of this operation is presented in Fig. 7, where the TSF is utilised to weight the raw spanwise vorticity field at t_1 (Fig. 7a). Figs. 7b and 7c represent a visual comparison of the vortex core when the TSF weighting is withheld and applied, respectively. For both figures, a threshold is applied between $-100 \geq \omega_z \geq -20$ to highlight candidate core locations. It can be seen that the vortex core is more clearly depicted in Fig. 7c as most of the background vorticity is suppressed in regions where the template similarity field has a high Wasserstein distance. A schematic overview of this pipeline is given in Fig. 8 where additional smoothing, filtering, or thresholding methods can be added as auxiliary support for more precise LEV core identification and tracking.

3. Results

3.1. Effects of template radius

In this section, we explore the influence of the template size on vortex core identification during stall flutter simulations. More specifically, r is varied parametrically between 0.02c and 0.06c to quantify how accurately the TSF is able to discriminate between the LEV core and background noise. For this analysis, the immediate LEV vicinity is only considered such that $R = 0.4c$. Additionally, three successive but non-consecutive timestamps are used to assess the generality of this parametric study for identifying the LEV core. The template vortex used in this study is the same region that was sourced in Fig. 4 and is referenced as t_0. The manually identified vortex core is set as $x, y = (0.8024c, 0.3509c)$. The IR also has a fixed centroid at this same location. As a preliminary assessment, Fig. 9 demonstrates the effect of varying the template size between $r = 0.02c$ and $r = 0.06c$ at t_1. For a small template radius, the TSF weighted vorticity field permits a wide distribution of vorticity. This implicates that the manually defined template is insufficient for differentiating between a vortex core and the smaller eddy structures.

As such, the TSF exhibits a large degree of similarity (low Wasserstein distance) that acts to apply an approximately uniform weighting. In contrast, by using a template with $r = 0.06c$, the background vorticity tends to be suppressed. This is due to a high localised dissimilarity in the TSF that accentuates the vortex core. Thus, it can be expected that the scaled vorticity distribution at low radii has a high variance. Conversely, as the ideal template size is approached, the distribution should narrow about zero. To test this supposition, we compute the probability distribution of the TSF weighted vorticity field at three successive timestamps

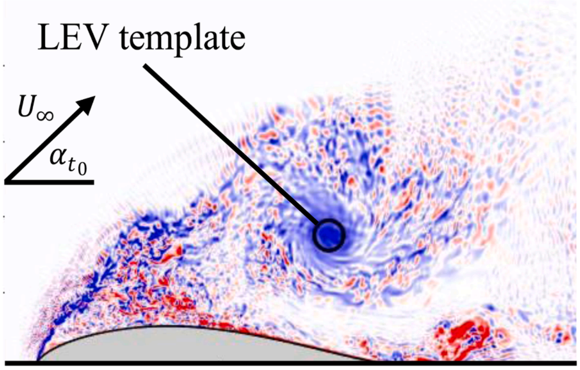

Fig. 4. ω_z field with manually defined LEV template source. ω_z range set between $-50 \leq \omega_z \leq 50$.

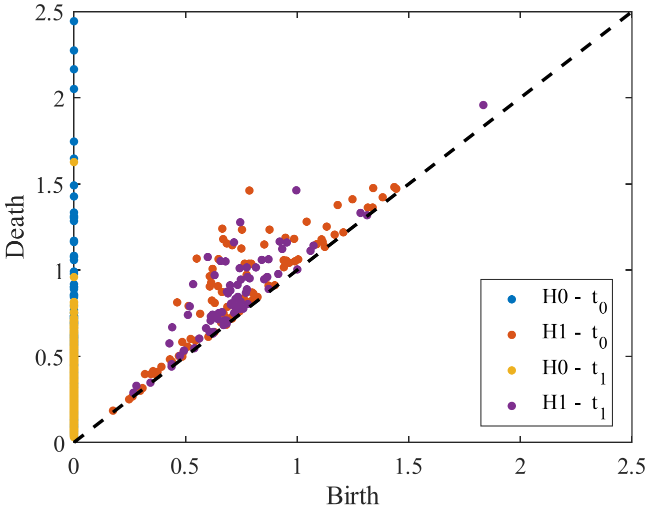

Fig. 5. Persistence diagram comparison between template sourced data and candidate core location at t_1.

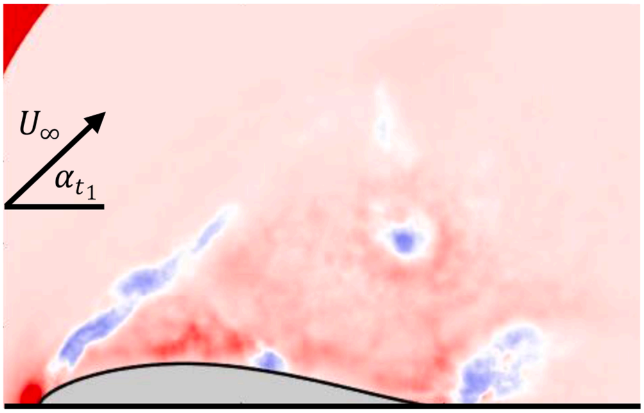

Fig. 6. Template similarity field with set between $0 \leq W_p \leq 15$.

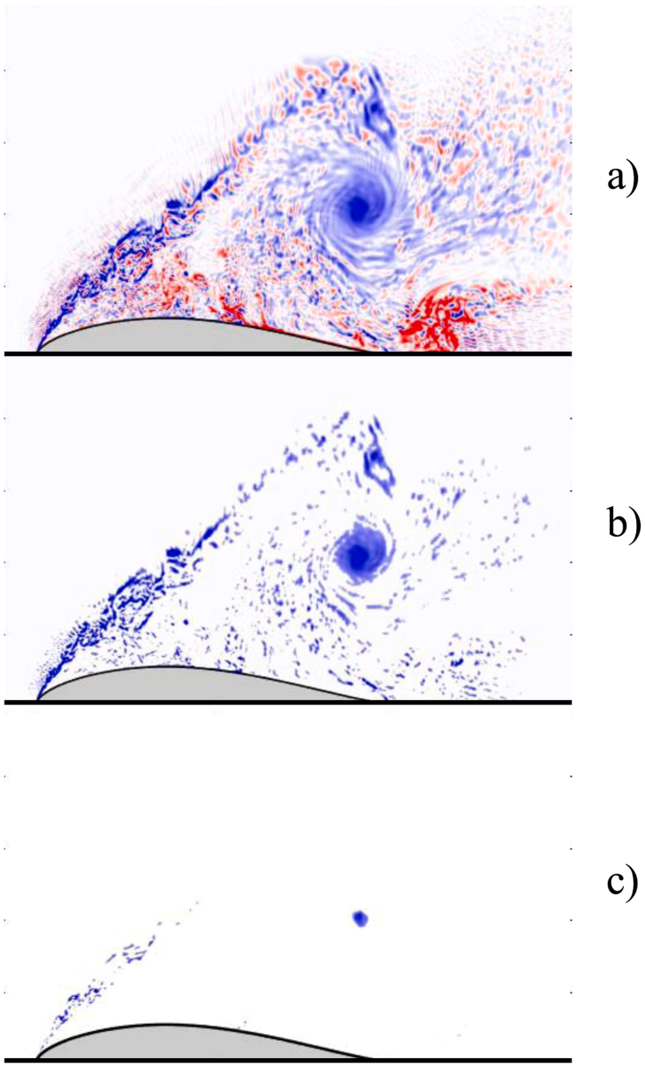

Fig. 7. a) Raw ω_z field with strong LEV present during high amplitude starting cycle. b) ω_z field with threshold applied between $-100 \geq \omega_z \geq -20$. c) TSF weighted ω_z field with threshold applied between $-100 \geq \omega_z \geq -20$.

using a logistic distribution fit. The spatial distribution mean, μ, and variance, σ^2, are monitored over the parametric range, $r = [0.02 : 0.01 : 0.06]$, and utilising the same previously defined template and IR. The results of this study are presented in Fig. 10 where t_i indicates the successive field snapshots. As expected, the weighted vorticity distribution tends to approach a mean value, $\mu_{\omega_z} = 0$ when r is varied between 0.02c and 0.06c across all cases. This is because more of the background vortical structures and noise are being suppressed due to a favourable template-target pairing. The variance also appears to shrink with r as per a larger disparity in vortex core candidate locations and noise. For both metrics, the change in distribution sensitivity appears to decrease dramatically at $r = 0.04c$. Beyond this template radius, the respective distributions between t_1, t_2, and t_3 are closely aligned.

An important consideration here, is that the construction of a Vietoris-Rips complex has a complexity of $O(2^n)$ where n is the number of data points sourced within r. Thus, the computational time and memory scaling for an ideal template is exponential in n and constant with the data dimension, R^d [30].

This potentially incurs significant computational challenges when automating the vortex core tracking in high fidelity simulations or particle image velocimetry (PIV) data. In some cases, the marginal increase in identification accuracy from utilising a larger template radius

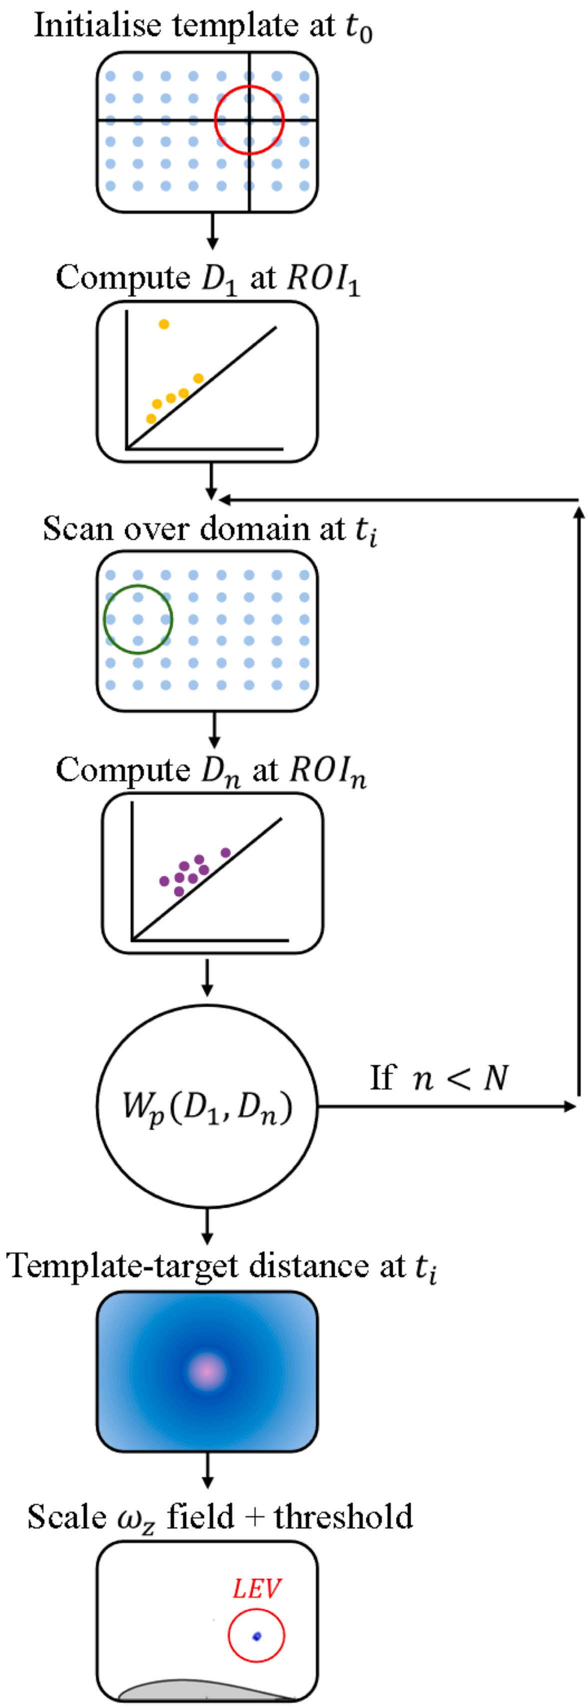

Fig. 8. Persistent homology-based template matching pipeline for LEV core identification.

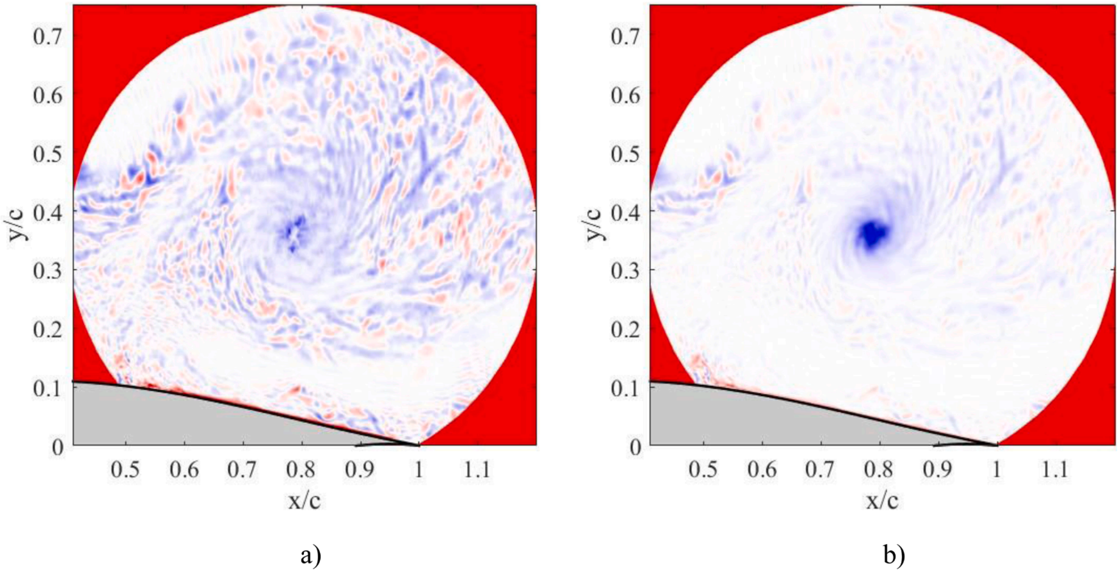

a) b)

Fig. 9. TSF weighted ω_z field comparison for $r = 0.02c$ (left) and $r = 0.06c$ (right). Both cases utilise an IR of $R = 0.4c$ and have the range set between $-50 \le \omega_z \le 50$.

may not warrant the computational overhead.

3.2. Effects of down sampling

To address the computational challenges associated with persistent homology, it is common to apply down sampling to the sourced point-cloud data. While an intuitive strategy to simplify the complex generation and filtration process, this ultimately changes the point-cloud topology and further complicates the feature space template/target 'definition' of a vortex core at each timestamp. A potentially important consideration is how the template/target data is down sampled (uniform, random, directional bias) to account for the vortex shape and spatial distribution. An alternative method is to simply apply down sampling to the prospective template/target ROI centroids as this does not alter its respective point cloud topologies, only the TSF spatial resolution which can be managed easily through interpolation. As such, we consider the effects of a uniform down sampling factor, M, within a fixed IR of $R = 0.3c$ on the vortex core identification.

For this test, down sampling is applied to the target snapshot by uniformly skipping over candidate centroid locations within the IR. All computations were conducted utilising an Intel i5 processor with 2.40 Ghz and multithreading over four physical cores. The Wasserstein distance at skipped locations is assigned by using a linear interpolation scheme on the unstructured grid. This ensures that the dimension of the TSF and vorticity field are the same. Finally, the vortex core is recovered here as the local minimum weighted vorticity within an IR that is centred by the template centroid location. Fig. 11 demonstrates the identified vortex core locations by varying the down sampling factor within the range, $M = [0:5:35]$. The template used here is defined at $x, y = (0.8024c,\ 0.3509c)$, with $r = 0.04c$.

It can be seen that the vortex core location is largely insensitive to the effects of down sampling as these predictions are contained within a range of $\pm 0.02c$ from the baseline case, $M = 0$. Though, these predictions tend to exhibit nonlinear behaviour as $M = 10$ and $M = 20$ yield the closest matches to our baseline definition. In general, down sampling appears to offer a significant advantage as the computational cost decreases dramatically when $M \ge 10$ (see Fig. 12). Furthermore, this advantage incurs a nominal decrease in prediction accuracy making it a suitable practice to implement in this vortex tracking pipeline.

3.3. Leading-edge vortex tracking

Utilising the developed vortex identification pipeline and a template definition with $x, y = (0.8024c,\ 0.3509c)$, $r = 0.04c$, and $M = 10$, the LEV trajectory is monitored during an aeroelastic starting cycle for a NACA 63(3)418 aerofoil.

As a preliminary demonstration, an IR with $R = 1c$ is used to provide a global comparison of the vortex core identification with and without the TSF weighting scheme. For this assessment, four successive but non-consecutive timestamps are considered. Both the TSF weighted and unweighted vorticity fields are computed and limited between $-100 \le \omega_z \le -20$. Spatial regions that satisfy this threshold criteria are then clustered together to highlight potential vortex candidate locations. For Fig. 13a, conventional thresholding of the raw spanwise vorticity field still yields significant noise at four distinct stages within the LEV lifecycle. This complicates the goal of vortex identification and tracking as there are many possible candidate core locations owing primarily to fluid shear interactions. When the TSF weighting scheme is applied though (Fig. 13b), we observe a significant filtration of these spurious vortex core locations. The leading-edge shear layer is still present at these timestamps but appears to be largely diminished compared to the conventional thresholding method. Finally, four vortex clusters are easily identifiable when the TSF weighting is applied. It is clear that this method provides an advantage in vortex core identification as it acts to filter out most of the turbulent background noise. This allows for a better qualitative and quantitative interpretation of the LEV dynamics during stall flutter.

To perform automated vortex tracking, an adaptive IR with $R = 0.2c$ is constructed based on the identified vortex core location. More specifically, the IR centroid at t_{i+1} is updated as the vortex core location, (x_c, y_c), at t_i. Fig. 14 presents the results of this LEV tracking study where the different trajectory traces correspond to the TSF weighting, local Q-criterion maximum, local vorticity minimum, local swirling strength maximum, and manual core definition methods respectively. The latter four methods all require user input whereas the blue trace (TSF) is completely automated aside from the template initialisation and first IR assignment. Note that all trajectories were smoothed using a sliding window filter of three timestamps. It can be seen that the TSF based tracking method is quite satisfactory in monitoring the core trajectory during its suction side migration and eventual severance when compared to the local $\omega_{z,min}$, local swirl strength maximum, manual core

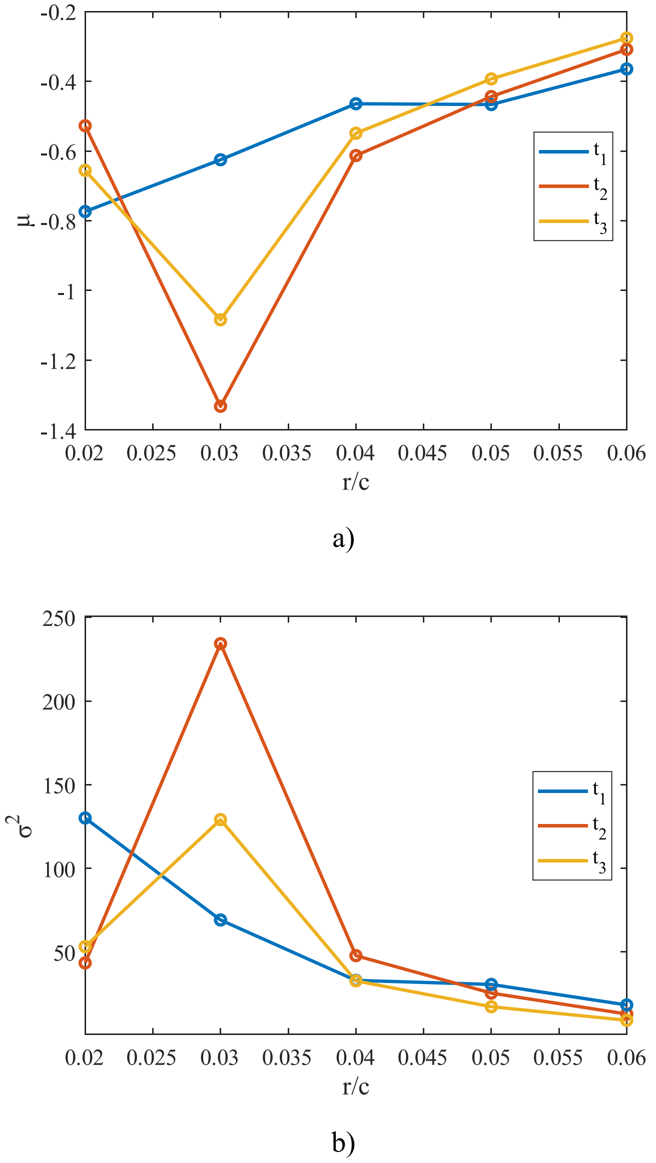

Fig. 10. Template size sensitivity analysis based on the TSF weighted ω_z distribution. a) distribution mean as a function of r. b) distribution variance as a function of r. Different colour plots correspond to the respective distribution at successive timestamps within the pitching cycle.

identification traces. The local Q-criterion maximum trace appears to exhibit stronger variations in the LEV core trajectory, especially in the vicinity leading up to LEV detachment. This suggests that TSF monitoring scheme is less sensitive to localised truncations of the flow field than the Q-criterion for the case study presented here.

Overall, the dynamic stall regime captured by these respective traces generally corresponds to a mature LEV status where the vortex structure is coherent and easily compatible with our chosen template. A distinct changepoint near $\frac{x}{c} = 0.9$ is captured in all trajectories. This trajectory shift away from the aerofoil has been previously linked to a trailing edge interaction with the LEV that is responsible for the LEV detachment and subsequent dynamic lift stall [2,31]. Thus, accurately identifying this tendency further demonstrates the capability of our method for monitoring the LEV as it progresses through various critical stages of its lifecycle. Overall, the TSF proves to be an effective tool for vortex core identification and monitoring in dynamically stalled flows.

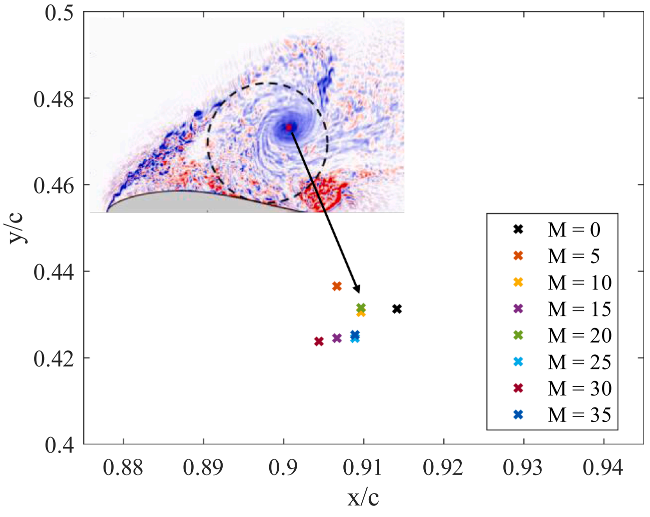

Fig. 11. Vortex core identification varied parametrically by a uniform down sampling factor, M. In all cases, the vortex core is quantified by the global minimum TSF weighted spanwise vorticity within the IR (black dashed boundary).

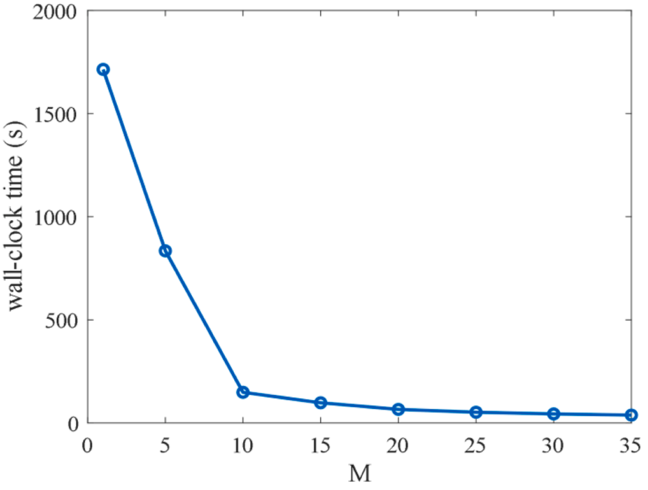

Fig. 12. Wall-clock time as a function of down sampling factor, M. All tests were conducted on a computer utilising an Intel i5 processor with 2.40 Ghz. Grid scanning is parallelised over 4 physical cores.

4. Conclusion

A novel vortex identification post-processing pipeline has been developed that leverages persistent homology for topology informed template matching. This pipeline has been successfully applied to the case study of a pitching NACA 63(3)418 aerofoil for the LEV detection and monitoring within a single pitching cycle. We have demonstrated that the TSF weighted vorticity field is quite effective at filtering out spurious vortex cores and background turbulent noise. This allows for a simpler interpretation of the dynamically stalled flow field when attempting to identify the LEV. We have demonstrated that the template-target pairing at various temporal snapshots is highly sensitive to the ROI radius when $r < 0.04c$. Though, larger templates incur a significant computational cost when computing persistent homology. This can be offset by utilising a uniform down sampling factor on the candidate core grid locations. It is shown that the computational time drops when $M \geq 10$ without significantly compromising the vortex core identification accuracy. Future work considers the application of this method to tracking multiple simultaneous vortices. This likely requires

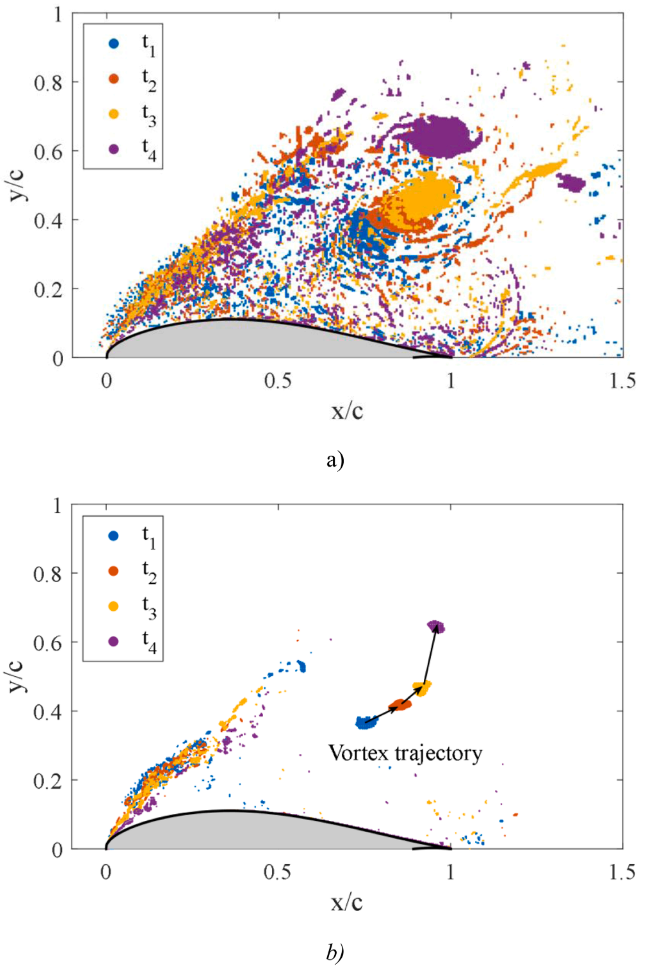

a)

b)

Fig. 13. a) Potential vortex core candidate locations clustered based on an applied threshold, $-100 \leq \omega_z \leq -20$, to the raw ω_z field. b) Potential vortex core candidate locations clustered based on an applied threshold, $-100 \leq \omega_z \leq -20$, to the TSF weighted ω_z field. Four successive temporal snapshots are utilised within the first starting cycle.

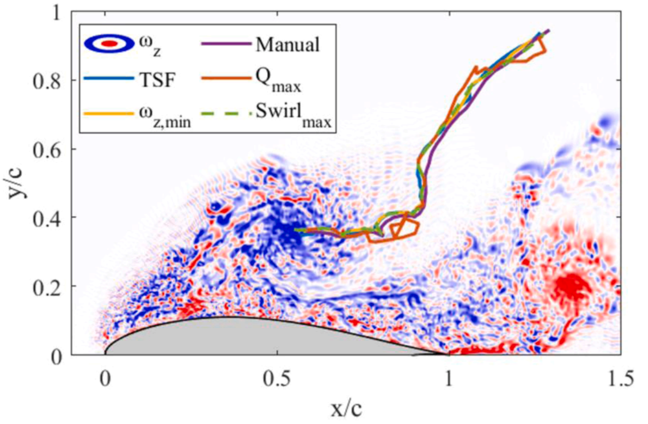

Fig. 14. TSF automated vortex core trajectory (blue) compared with several user-in-the-loop monitoring strategies. Trajectories are computed starting from the vortex state illustrated by the ω_z contour. Contour range is set between $-50 \leq \omega_z \leq 50$.

integration with current state-of-the-art tracking techniques like clustering or machine vision methodologies.

CRediT authorship contribution statement

Quentin Martinez: Writing – original draft, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Chetan Jagadeesh:** Writing – review & editing, Supervision. **Marinos Manolesos:** Writing – review & editing, Supervision. **Mohammad Omidyeganeh:** Writing – review & editing, Supervision, Resources.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.compfluid.2025.106931.

Data availability

Data will be made available on request.

References

[1] Rosti ME. Direct numerical simulation of an airfoil at high angle of attack and its control, Ph.D.Dissertation, School of Science and Technology, City, University of London, London, UK, 2016.

[2] Martinez Q, et al. Large Eddy simulations of a. NACA 2025;63(3):418. Airfoil Undergoing One-Shot Starting Cycles at Highly Unsteady Conditions, 59th 3AF International Conference on Applied Aerodynamics, Strasbourg, France, 24-26 March 2025. Association Aeronautique et Astronautiquede France (3AF).

[3] Mulleners K, Raffel M. Dynamic Stall Development. Exp Fluids 2013;54(2). https://doi.org/10.1007/s00348-013-1469-7.

[4] Eldredge JD, Jones AR. Leading-edge vortices: mechanics and modeling. Annu Rev Fluid Mech 2019;51(1):75–104. https://doi.org/10.1146/annurev-fluid-010518-040334.

[5] Jackson FA, et al. Experimental study of dynamic stall of avian inspired airfoils, AIAA SCITECH 2025 Forum. In: Orlando, United States, 6-10 January 2025. American Institute of Aeronautics and Astronautics; 2025. https://doi.org/10.2514/6.2025-1074.

[6] Lee H, et al. Leading edge vortex formation and wake trajectory: synthesizing measurements, analysis, and machine learning. Phys Rev Fluids 2022;7(7). https://doi.org/10.1103/physrevfluids.7.074704.

[7] Soto-Valle R, et al. Vortex identification methods applied to wind turbine tip vortices. Wind Energy Sci. 2022;7:585–602. 10.5194/wes-7-585-2022, 2022.

[8] Ibanez C, et al. Cluster-based tracking for the identification and characterisation of vortices, 59th 3AF International Conference on Applied Aerodynamics, Strasbourg, France, 24-26 March 2025. Association Aéronautique et Astronautique de France (3AF).

[9] Xu Z, et al. Vortex and core detection using computer vision and machine learning methods. Eur J Comput Mech 2023. https://doi.org/10.13052/ejcm2642-2085.3252.

[10] Ebling J, Scheuermann G. Segmentation of flow fields using pattern matching, Eurographics/IEEE-VGTC Symposium on Visualization, Lisbon, Portugal, 8-10 May 2006. Eurographics Association 2006:147–54. https://doi.org/10.2312/VisSym/EuroVis06/147-154.

[11] Rodrigues PSS, De A. Araujo A, Pinotti M. Describing patterns in flow-like images, Proceedings 10th International Conference on Image Analysis and Processing, Venice, Italy, 27-29 September 1999. IEEE. doi:10.1109/ICIAP.1999.797632.

[12] Heiberg E, et al. Three-dimensional flow characterization using vector pattern matching. IEEE Trans Vis Comput Graph 2003;9(3):313–9. https://doi.org/10.1109/tvcg.2003.1207439.

[13] Wong KY, Yip CL. Identifying centers of circulating and spiraling vector field patterns and its applications. Pattern Recognit 2009;42(7):1371–87. https://doi.org/10.1016/j.patcog.2008.11.037.

[14] Elsinga GE, et al. Tracking of vortices in a turbulent boundary layer. J Fluid Mech 2012;697:273–95. https://doi.org/10.1017/jfm.2012.60.

[15] Overmeyer A. A rotor tip vortex tracing algorithm for image post-processing. Annu AHS Forum Technol Disp Va Beach U S 2015. 5-7 May 2015. Langley Research Center.

[16] Suzuki A, et al. Flow estimation solely from image data through persistent homology analysis. Sci Rep 2021:11. https://doi.org/10.1038/s41598-021-97222-6.

[17] Moon C, et al. Statistical inference over persistent homology predicts fluid flow in porous Media'. Water Resour Res 2019;55(11):9592–603. https://doi.org/10.1029/2019WR025171.

[18] Tymochko S, et al. Using persistent homology to quantify a diurnal cycle in hurricanes. Pattern Recognit Lett 2020;133:137–43. https://doi.org/10.1016/j.patrec.2020.02.022.

[19] Smith L, et al. A cyclic perspective on transient gust encounters through the lens of persistent homology. J Fluid Mech 2024;980. https://doi.org/10.1017/jfm.2024.16. p.A18.

[20] Chen Y, Lin H. Analyzing singular patterns in discrete planar vector fields via persistent path homology. Comput Graph 2025:132. https://doi.org/10.1016/j.cag.2025.104354.

[21] Hunt JCR, et al. Eddies, streams, and convergence zones in turbulent flows. Proceedings of the 1988 Summer Program, Centre for Turbulence Research, NASA Ames/Stanford University, pp. 193–208. Available at: https://ntrs.nasa.gov/api/citations/19890015184/downloads/19890015184.

[22] Zhou J, et al. Mechanisms for generating coherent packets of hairpin vortices in channel flow. J Fluid Mech 1999;387:353–96. https://doi.org/10.1017/S002211209900467X.

[23] Martinez Q, et al. High amplitude stall flutter simulations of a NACA 63(3)418 aerofoil. AIAA AVIAT FORUM AND ASCEND 2025. https://doi.org/10.2514/6.2025-3839. 21-25 July 2025. American Institute of Aeronautics and Astronautics.

[24] Corke TC, Thomas FO. Dynamic stall in pitching airfoils: aerodynamic damping and compressibility effects. Annu Rev Fluid Mech 2015;47(1):479–505. https://doi.org/10.1146/annurev-fluid-010814-013632.

[25] Edelsbrunner H, et al. Topological persistence and simplification. Discrete Comput Geom 2002;28:511–33. https://doi.org/10.1007/s00454-002-2885-2.

[26] Bauer U. Ripser: efficient computation of Vietoris–rips persistence barcodes. J Appl Comput Topol 2021;5(3):391–423. https://doi.org/10.1007/s41468-021-00071-5.

[27] Wasserstein LN. Markov processes over denumerable products of spaces describing large systems of automata. Probl Inf Transm 1969;5:47–52.

[28] Agami S. Comparison of persistence diagrams. Commun Stat - Simul Comput 2021; 52(5):1948–61. https://doi.org/10.1080/03610918.2021.1894335.

[29] Adams H, et al. Persistence images: a stable vector representation of persistent homology. arXiv prepr 2016. arXiv:1507.06217. https://arxiv.org/abs/1507.06217.

[30] Somasundaram E, et al. Benchmarking R packages for calculation of persistent homology. R J 2021;13(1):184. https://doi.org/10.32614/rj-2021-033.

[31] Kissing J, et al. Insights into leading edge vortex formation and detachment on a pitching and plunging flat plate. Exp Fluids 2020;61(9). https://doi.org/10.1007/s00348-020-03034-1.