

Rotation-Minimizing Space Deformations for the Analysis of Lagrangian Vortex Corelines

Julio Rey Ramirez , Xingdi Zhang , Markus Hadwiger , Peter Rautek 




Figure 1: Pathlines computed in the lab frame depict particle motion typically leading to cluttered visualizations that do not provide insight regarding the rotational behavior of vortex regions (top-left, advected for $N = 60$ timesteps). Transforming pathlines with rotation-minimizing deformations allows us to analyze the rotating motion of particles in the vicinity of vortex cores. We define a vortex-centric reference frame for corelines as they evolve across time (bottom-left). This enables the transformation of pathlines into the vortex-centric reference frame, yielding curves that represent particle motion relative to the vortex core, removing the global drift and deformation of the core (middle $t_0 = 0.98$), enabling the visual verification and quantitative analysis of Lagrangian vortex cores. The insets show examples of insights possible with our method. Top-right: a vortex filament from the wing, encircled by weak arc-shaped cores. Bottom-right: Changing winding angles along vortex cores.

Abstract—Visualizing and verifying unsteady 3D vortex structures remains a fundamental challenge in flow visualization. While recent advancements enable the extraction of objective Lagrangian vortex corelines, visually verifying whether these candidate curves accurately represent the true center of swirling motion remains difficult. Standard visualization techniques, such as instantaneous streamlines, lack Galilean invariance and are highly dependent on the chosen frame of reference, often yielding misleading depictions. Conversely, while Lagrangian pathlines offer a physical ground truth, the local rotational component of particle motion is typically occluded by the convective motion of the flow. To address this, we propose a novel visualization technique that associates objective space-time pathlines with evolving Lagrangian corelines. Our method applies a space deformation induced by the deformation of the coreline that guarantees that no rotation component around the core's tangent is introduced. By applying this deformation to the tubular neighborhood of the coreline, we transform pathlines into the vortex-centric reference frame, effectively decoupling the local swirl of particles from the global drift and deformation. A primary application of this technique is the visual evaluation of vortex coreline candidates. By explicitly visualizing inversely transformed pathlines in context with the coreline curves, we directly reveal the swirling motion of the surrounding particles. Because this swirling motion is the physical ground truth for validating a vortex coreline, our method provides the first pathline-based visual assessment of candidate curves. We demonstrate our approach on multiple unsteady 3D flow datasets and methods for coreline extraction, revealing vortex-relative particle dynamics that are otherwise hidden by bulk transport.

Index Terms—Flow Visualization, Vortex Dynamics, Pathlines, Spacetime Deformation.

1 INTRODUCTION

The identification and analysis of vortex structures play a central role in understanding complex phenomena in fluid mechanics, since vortices segment flow into coherent structures that transport mass, momentum, and energy. Some relevant applications include turbulent mixing, aerodynamic drag, combustion, and climate modeling. Despite their relevance, the definition of a vortex has long been a subject of debate.

Objectivity and Material Properties. Historically, vortex detection relied on Eulerian criteria (e.g., Q -criterion, λ_2) applied to instantaneous velocity fields. However, these snapshot-based methods suffer

from a critical flaw: they lack *reference frame invariance*. The structures they reveal often depend on the velocity of the observer, leading to "phantom" vortices that disappear when the reference frame changes.

Recent advances have successfully addressed this by adopting a *Lagrangian* perspective. State-of-the-art methods can now extract *objective Lagrangian vortex corelines* directly from unsteady 3D flow data [11, 31]. These corelines represent material features that persist over time, tracking the centers of swirling particle bundles regardless of the observer's motion.

Visualization and Analysis of Vortex Structures. Detecting corelines is an essential step in the visualization and analysis pipeline. However, to understand the physics, we must analyze the *context* of the vortex. While the visualization community has made significant strides in detecting and representing vortices for visual and computational analysis, further understanding of the interaction between the core and

its surroundings is still missing. Paradoxically, while recent methods extract corelines that are Lagrangian as well as reference frame invariant, we lack a method to visualize the dynamic behavior of the flow *relative* to them.

Streamlines are commonly used to depict the local swirling motion. Günther and Theisel [11] choose to visualize streamlines that are traced in a flow field as it would be locally observed by a co-moving observer. Since their method computes different optimally co-moving observers at each point in space and time, the overall visualization uses an infinite number of co-moving observers. In general, this makes the interpretation of the rendered streamlines non-trivial. Zhang et al. [31] choose to visualize streamlines that are traced in a flow field as it would be locally observed by one co-moving observer. To show the swirling behavior around a coreline that deforms and moves with different speeds at different parts of the coreline, they show multiple visualizations side-by-side, each using their best-fitting co-moving observer. While this leads to interpretable visualizations, they face a tradeoff between the quality of the co-moving reference frames and the number of visualizations needed to show the same flow in multiple reference frames.

While these methods successfully visualize streamlines to verify the extraction quality of corelines, they remain locked in an instantaneous snapshot. They serve as a proxy for ‘what the algorithm sees,’ but they cannot describe the evolution of a material particle. In contrast, our method shifts the paradigm to Lagrangian analysis, where we visualize the actual material particles to reveal the dynamic evolution of the vortex itself.

In general, streamlines are not adequate to visualize and analyze the local swirling motion of material particles for multiple reasons:

Instantaneousness. Streamlines do not show the actual dynamical behavior of unsteady flow.

Frame Dependence. To see the swirling motion of streamlines, one must remove the convective velocity of the flow. In unsteady flows where the vortex accelerates or deforms, defining optimally co-moving reference frames is ambiguous.

Lack of Quantitative Measures. Streamlines represent the flow field at a frozen instant in time. In unsteady flows, this representation can be misleading and prevents quantitative measurements and analysis of effective swirling strength of a vortex.

On the other hand, pathlines provide the ground truth of particle motions. However, interpreting pathlines in the lab-frame is notoriously difficult. While for 2D unsteady flow, time can be mapped to the third dimension, no method has shown how to decouple the convective from the swirling motion for unsteady 3D flow. The rotational motion of the fluid is often obscured by the global translation of the vortex itself. When viewed in 4D spacetime, these trajectories form dense, tangled structures that are often impossible to interpret on a 2D screen.

Our approach employs a deformation of the space that is induced by the deformation of the vortex coreline over time. We propose a method to deform the spacetime neighborhood of a vortex minimizing the rotation of the vortex-centric reference frame around the tangent of the coreline. This way we construct a deformation operator that maps the curving and translating vortex core into a stationary domain, without introducing artificial twisting. By deforming the surrounding pathlines into this space, we decouple the global transport from the local rotation. This reveals the swirling motion of pathlines in the local space of each vortex core and allows us to combine the physical meaningfulness of Lagrangian pathlines with the reduced visual complexity of a vortex-centric reference frame. Further, by enforcing a smooth diffeomorphic mapping in the coreline’s tubular neighborhood and omitting artificial rotation around the vortex core, we ensure that rotational quantities (e.g., the winding angle) are correctly preserved.

In summary, our method allows us to analyze the dynamics of unsteady vortices, by providing the first pathline-based vortex-centric visualization of objective Lagrangian vortex cores.

Our contributions are:

- A method to construct a *twist-free frame evolution* along advected corelines using only the coreline geometry, via Rotation-Minimizing Frame transport, providing a consistent, singularity-free basis for expressing local particle displacements across time in a well-defined

tubular neighborhood of the coreline.

- A *vortex-centric coordinate transformation* that maps Lagrangian pathlines into a canonical reference frame anchored to the vortex coreline at an arbitrary time, reducing the visualization complexity from 4D to 3D and representing pathlines relative to their associated vortex coreline. This enables the visual analysis of vortex dynamics, such as swirling strength profiles around the vortex core, and longitudinal drifting along the vortex core.

- *Vortex-relative quantities*, such as the accumulated winding angle, become naturally expressible in the vortex-centric frame and provide quantitative insight into rotational behavior.

2 RELATED WORK

2.1 Coreline Extraction

The extraction and analysis of vortex structures is a fundamental task in flow visualization [10]. A vortex definition for 3D unsteady flows, where time-dependence complicates steady-flow intuitions [10, 21], remains an open research question. For a detailed overview of methods, we refer to state-of-the-art reports [3, 7, 10].

Early approaches focused on instantaneous Eulerian criteria, ranging from region-based measures like the Q -criterion [4, 15], the λ_2 -criterion [16], and the Δ -criterion [4], to eigenvector-based coreline definitions such as the method by Sujudi and Haines [27]. While widely used, these measures are not objective, meaning they depend on the observer’s frame of reference, which is problematic for the analysis of unsteady flow.

Reference Frame Optimization. To address this, Haller [13, 14] introduced the concept of objectivity for vortex definitions. Subsequent work proposed finding an optimal local reference frame that minimizes the unsteadiness of the observed flow [8]. This idea has been extended using differential geometry and Killing fields [12, 23] as well as locally adaptive moving least squares [22] solutions. Interactive tools have been shown to enable the exploration and detection of objectively determined vortex structures [24] as well as enhanced analysis and interpretation by explicitly changing the frame of reference [29, 30].

Objective Lagrangian Cores. A key challenge in unsteady flow analysis is ensuring that extracted corelines represent material structures that move with the fluid. Corelines extracted individually at instantaneous timesteps are typically not Lagrangian. However, these instantaneous curves can be advected over a finite time window to sweep out path-surfaces, effectively yielding a Lagrangian representation for that duration. Alternatively, recent advancements, such as the methods by Günther and Theisel [11] as well as Zhang et al. [31], bypass instantaneous extraction to directly compute objective vortex cores that are guaranteed to be Lagrangian.

Regardless of the approach used to obtain them, these objective Lagrangian solutions share a common fundamental problem: the resulting structures are notoriously difficult to verify, but crucial for further analysis. This difficulty stems from their complex space-time nature and the absence of a clear physical ground truth for isolated swirling motion in convective global transport. Theisel [11] as well as Zhang et al. [31] use streamlines computed in the co-moving reference frame. These streamline-based approaches are essentially ‘verification tools’: they allow experts to check if an extracted curve aligns with the surrounding velocity field. However, because streamlines are spatially integrated rather than temporally, they are fundamentally unable to characterize the actual particle motion. Our method is designed to address this. By transforming the surrounding spacetime pathlines into the stationary reference frame of the given vortex coreline, we provide a method that avoids the visual clutter of pathlines in the lab frame - yet allows for the analysis of true particle motions in the neighborhood of a Lagrangian vortex coreline in 3D unsteady flow.

2.2 Transformations and Deformations

To achieve this transformation into a coreline-local perspective, our method requires a mathematical mechanism to effectively invert the deformation and transport of the vortex core over time. We accomplish this through an inverse space deformation that is uniquely induced by




Figure 2: **Vortex-centric visualization pipeline.** The inputs are an unsteady flow field and a coreline advected for N time steps, from time t_0 to t_N . (a) input flow field represented as a line integral convolution (LIC) slice, alongside pathlines. (b) several timesteps of the advected coreline. The Lagrangian coreline defines a rotation-minimizing deformation of the tubular neighborhood. (c) shows how this deformation acts on cubic cells around a coreline, enabling the transformation of arbitrary spacetime geometry inside the tubular neighborhood to any reference time, without introducing additional rotation around the coreline. (d) shows the result of transforming the input pathlines from (a) to different reference times of the advected coreline. This vortex-centric representation enables us to compute quantitative measures, like the winding angle, without introducing artificial twisting.

the spatio-temporal evolution of the coreline candidate itself and valid in a tubular region around the coreline.

Space Deformations in Graphics. Smooth space deformation is a core problem in computer graphics, where various coordinate systems have been proposed to deform space relative to a cage or a skeleton. Notable examples include harmonic coordinates [17], mean value coordinates [28], and Green coordinates [20].

Curve Reformation Techniques for Visualization. In medical visualization, curved planar reformation (CPR) [18] and its variants are used to straighten curved structures like blood vessels. Similar curve-centric reformation techniques have been applied to volume data [5] and tubular flow [1] to reduce occlusion and facilitate comparison.

Transformations and Deformations in Flow Visualization. Reference frame transformations are commonly used to create insightful visualizations by moving with the flow features [23, 29, 30, 32]. Günther and Theisel [11] showed streamlines in a space deformed by pointwise optimal reference frames co-moving with corelines.

Straub et al. [26] recently introduced a feature-based deformation framework that straightens curved flow features to facilitate the observation of local dynamics.

While their method maintains important geometric properties, it is primarily designed for a per-timestep analysis of steady flow. Consequently, the resulting visualizations rely on instantaneous streamlines. Relying on streamlines in unsteady flows is problematic for structural analysis because they are sensitive to the chosen frame of reference, and might produce misleading visual artifacts (e.g., phantom vortices). Our method addresses this shortcoming by operating strictly on objective Lagrangian structures. By transforming physical spacetime pathlines rather than instantaneous streamlines, we yield a visualization of intrinsic material dynamics. This novel method allows the user to perform qualitative and quantitative analysis that instantaneous methods fundamentally cannot provide.

3 METHOD

The central idea of our method is to define a deformation of the space surrounding each coreline that allows us to map spacetime primitives (such as pathlines) to any target time, thereby isolating the local swirling dynamics from the global transport and deformation of the vortex core. The pipeline of our vortex-centric visualizations is illustrated in Fig. 2. The inputs to our method are an unsteady flow field $\mathbf{v}(\mathbf{x}, t)$ (visualized with a set of pathlines in Fig. 2 (a)) and a set of Lagrangian corelines that evolve over time (one pathline evolving over time is shown in Fig. 2 (b)). The key challenge is to define a deformation in a way that is *unique*, *twist-free*, and *invariant under rigid body transformations*, such that measurable quantities (e.g., winding angles) stay physically

meaningful. We achieve this by constructing reference frames along the coreline and evolving them over time, without introducing rotation around the coreline’s tangent. These reference frames induce a smooth and twist-free deformation of the space in a well-defined tubular neighborhood around the coreline. An example of such a coreline-induced deformation that is free of twist around the coreline’s tangent is shown in Fig. 2 (c). We can then use these reference frames and the associated space deformation for vortex-centric visualizations and quantitative analysis of vortex dynamics. In Fig. 2 (d) we show such a vortex-centric visualization where we map the pathlines (shown in Fig. 2 (a)) to different target times of the vortex coreline.

In the following, we first develop the continuum formulation (Sec. 3.1–3.3) and then present the discrete implementation (Sec. 3.4).

3.1 Lagrangian Corelines and Tubular Neighborhoods

Material curves in unsteady flow. Consider an unsteady velocity field $\mathbf{v}(\mathbf{x}, t)$ on a spatial domain $\Omega \subset \mathbb{R}^3$ over a time interval $[t_0, t_N]$. The flow map $\phi_{t_0}^t : \Omega \rightarrow \Omega$ sends every point $\mathbf{x}_0 \in \Omega$ at time t_0 to its advected position at time t :

$$\frac{d}{dt} \phi_{t_0}^t(\mathbf{x}_0) = \mathbf{v}(\phi_{t_0}^t(\mathbf{x}_0), t), \quad \phi_{t_0}^{t_0}(\mathbf{x}_0) = \mathbf{x}_0. \quad (1)$$

A *material curve* is a smooth curve $\mathcal{C}_{t_0} \subset \Omega$ at a seed time t_0 whose points are advected as material particles, yielding the time-dependent family

$$\mathcal{C}_t = \phi_{t_0}^t(\mathcal{C}_{t_0}), \quad t \in [t_0, t_N]. \quad (2)$$

We call such a material curve a *Lagrangian vortex coreline* (or, in this paper, simply *coreline*) if it is detected as the center of swirling motion of a vortex by a vortex detection method. Although our method is designed to work with material curves that are centers of swirling motion, it does not strictly depend on them and can be used to visualize fluid dynamics around material curves in general. However, since our primary application is the visualization of vortex dynamics, we will mostly use the term *Lagrangian coreline* even when *material curve* is often meant. We parameterize $\mathcal{C}_{t_0}$ by a parameter $s \in [0, L]$ and write the coreline position as $\gamma(s, t) = \phi_{t_0}^t(\gamma(s, t_0))$, where $\gamma(s, t_0) := \mathcal{C}_{t_0}(s)$. The spatial parameter s acts as a *label* when advecting fluid elements over time and provides a temporally stable *longitudinal coordinate*.

A typical choice in curve-centric methods [1, 5, 26] is to use the *arc-length* of the curve as the longitudinal coordinate. While arc-length parameterization distributes points uniformly along the curve at each instant, it discards the physical correspondence between material points across time: as the coreline stretches and compresses under advection, the arc-length labels of individual fluid elements change. By contrast,




Figure 3: Construction of the rotation-minimizing frame transport at a point $\gamma(s, t)$ on the coreline $\mathcal{C}_t$. The tangent $\mathbf{T}$ (blue) and the frame vectors $\mathbf{U}, \mathbf{V}$ (red, green) span a local coordinate system. As the coreline evolves, $\mathbf{T}$ changes direction; this change $\partial_t \mathbf{T}$ (orange) lies in the normal plane and is shown both at the tip of $\mathbf{T}$ (indicating the direction of tilt) and at the base point (where it decomposes into $\mathbf{U}$ and $\mathbf{V}$ components). The coefficient $c_1 = -\partial_t \mathbf{T} \cdot \mathbf{U}$ determines $\partial_t \mathbf{U}$ (purple), which points along $-\mathbf{T}$ to maintain orthogonality. The twist-inducing component c_3 is set to zero, ensuring the frame does not rotate around the tangent.

our material parameterization preserves this correspondence, so that the resulting deformation faithfully reflects how the material space evolves. In particular, regions of the coreline that are compressed or stretched by the flow produce a correspondingly compressed or stretched tubular neighborhood, capturing the true Lagrangian kinematics.

The tubular neighborhood of $\mathcal{C}_t$. Our goal is to define a well-posed coordinate system in a neighborhood of $\mathcal{C}_t$ that can be transported consistently across time. At each point $\gamma(s, t)$, the unit tangent

$$\mathbf{T}(s, t) = \frac{\partial_s \gamma(s, t)}{\|\partial_s \gamma(s, t)\|} \quad (3)$$

determines a two-dimensional *normal plane* $\mathbf{T}(s, t)^\perp$. The collection of all such normal planes along the curve forms the *normal bundle*

$$N\mathcal{C}_t = \{(\gamma(s, t), \mathbf{n}) : s \in [0, L], \mathbf{n} \in \mathbf{T}(s, t)^\perp\}. \quad (4)$$

The natural embedding $\Psi_t(\gamma(s, t), \mathbf{n}) = \gamma(s, t) + \mathbf{n}$ maps the normal bundle into $\mathbb{R}^3$. By the *tubular neighborhood theorem* [19], there exists $\varepsilon > 0$ such that Ψ_t restricts to a diffeomorphism from the ε -disk bundle

$$N_\varepsilon \mathcal{C}_t = \{(\gamma(s, t), \mathbf{n}) \in N\mathcal{C}_t : \|\mathbf{n}\| < \varepsilon\} \quad (5)$$

onto an open set $\mathcal{U}_\varepsilon(\mathcal{C}_t) \subset \mathbb{R}^3$ enclosing the curve. The maximum admissible radius ε is bounded by the *reach* of $\mathcal{C}_t$, which is the minimum of the reciprocal of the curve's maximum curvature and half its closest self-approach distance. We illustrate this in Appendix A.

Within $\mathcal{U}_\varepsilon(\mathcal{C}_t)$, every point $\mathbf{x}$ has a *unique* closest point $\gamma(s^*, t)$ on the coreline and a well-defined normal displacement

$$\mathbf{n} = \mathbf{x} - \gamma(s^*, t) \in \mathbf{T}(s^*, t)^\perp, \quad \|\mathbf{n}\| < \varepsilon. \quad (6)$$

To express $\mathbf{n}$ as a pair of scalar coordinates, we choose an arbitrary orthonormal frame $\{\mathbf{U}(s, t), \mathbf{V}(s, t)\}$ spanning the normal plane $\mathbf{T}(s, t)^\perp$ at each point along the curve. The displacement then decomposes as $\mathbf{n} = u\mathbf{U}(s^*, t) + v\mathbf{V}(s^*, t)$, and every point $\mathbf{x} \in \mathcal{U}_\varepsilon(\mathcal{C}_t)$ is uniquely described by the triple of scalar *curvilinear coordinates*

$$(s^*, u, v) \in [0, L] \times \mathbb{R}^2, \quad u^2 + v^2 < \varepsilon^2, \quad (7)$$

3.2 Rotation-Minimizing Temporal Frame Transport

The curvilinear coordinates (s^*, u, v) from (7) depend on the choice of orthonormal frame $\{\mathbf{U}, \mathbf{V}\}$ in the normal plane. At a single time t , an arbitrary assignment $s \mapsto \{\mathbf{U}(s, t), \mathbf{V}(s, t)\}$ yields a valid coordinate system. However, as the coreline evolves over time, the tangent

$\mathbf{T}(s, t)$ and hence the normal plane $\mathbf{T}(s, t)^\perp$ change at each material point. To compare coordinates across time (e.g., to transport points of pathlines), we need a uniquely defined *temporal* evolution of the frame. A naive transport rule could introduce artificial twist around the tangent direction, corrupting the measurement of physical winding angles. A seemingly natural candidate to advect the frame is the *pushforward* of the flow map: since ϕ'_t advects the coreline, its Jacobian $D\phi'_t$ transports tangent vectors and, after projection and re-orthonormalization, could define a frame evolution. However, this transport carries the *full* local deformation of the flow, including any rotational component around the coreline. A frame transported by the pushforward would co-rotate with the swirling fluid, causing the angular coordinate of a winding particle to remain constant, precisely canceling the swirl we wish to observe. To reveal the physical winding, we instead require a transport that is *agnostic* to the flow's local rotation: one that evolves the frame based solely on the geometric change of the coreline's tangent, with no rotational component about $\mathbf{T}$. This is exactly what rotation-minimizing transport achieves. Note that this also differs fundamentally from parallel transport on manifolds. While parallel transport relies on the intrinsic metric to propagate vectors along a path (effectively producing a 'parallel copy' in Euclidean 3D), our approach transports the frame through time explicitly using the coreline's tangent. This geometric transport is decoupled from the manifold's metric, allowing us to isolate local particle motion from the coreline's global convective evolution.

Temporal frame transport. At each material point $\gamma(s, t)$, the tangent $\mathbf{T}(s, t)$ evolves continuously in time as the coreline deforms. In Fig. 3 we illustrate a material curve $\mathcal{C}_t$, a point on the curve $\gamma(s, t)$, and the tangent $\mathbf{T}$ at this point in time. We will use the orthonormal frame $\{\mathbf{T}, \mathbf{U}, \mathbf{V}\}$ (also shown in Fig. 3) as a basis to show how the infinitesimal rotation of the frame is defined, which rotates the frame but without rotating around $\mathbf{T}$.

Since $\mathbf{T}$ has unit length, differentiating $\mathbf{T} \cdot \mathbf{T} = 1$ with respect to t gives $\partial_t \mathbf{T} \cdot \mathbf{T} = 0$: the rate of change of the tangent is orthogonal to the tangent itself, i.e., $\partial_t \mathbf{T}$ lies entirely in the normal plane $\mathbf{T}^\perp$. While $\partial_t \mathbf{T}$ can be visualized as a velocity vector at the tip of $\mathbf{T}$ indicating the direction in which the tangent is tilting (see Fig. 3, orange arrow at the tip of $\mathbf{T}$), it is, strictly speaking, a tangent-space vector based at $\gamma(s, t)$. We therefore also show $\partial_t \mathbf{T}$ projected onto the normal plane at the base point (orange arrow in the $\mathbf{U}$ - $\mathbf{V}$ plane), where it can be decomposed into components along $\mathbf{U}$ and $\mathbf{V}$.

While the evolution of $\mathbf{T}$ is prescribed by the Lagrangian coreline, the rotation of the frame $\{\mathbf{T}, \mathbf{U}, \mathbf{V}\}$ is in general under-determined. It must follow the evolution of $\mathbf{T}$ while remaining orthonormal. Since $\{\mathbf{T}, \mathbf{U}, \mathbf{V}\}$ is a basis, we can expand the time derivative of $\mathbf{U}$ in its three components:

$$\partial_t \mathbf{U} = \underbrace{(\partial_t \mathbf{U} \cdot \mathbf{T}) \mathbf{T}}_{c_1} + \underbrace{(\partial_t \mathbf{U} \cdot \mathbf{U}) \mathbf{U}}_{c_2} + \underbrace{(\partial_t \mathbf{U} \cdot \mathbf{V}) \mathbf{V}}_{c_3}. \quad (8)$$

We now determine each coefficient in such a way that the frame is uniquely defined and does not rotate around $\mathbf{T}$.

Coefficient c_2 (along $\mathbf{U}$). The unit-length constraint $\mathbf{U} \cdot \mathbf{U} = 1$, differentiated in time, gives $\partial_t \mathbf{U} \cdot \mathbf{U} = 0$, so $c_2 = 0$. This simply states that the rate of change of a unit vector has no component along itself.

Coefficient c_1 (along $\mathbf{T}$). The orthogonality constraint $\mathbf{T} \cdot \mathbf{U} = 0$, differentiated, gives $\partial_t \mathbf{T} \cdot \mathbf{U} + \mathbf{T} \cdot \partial_t \mathbf{U} = 0$, so $c_1 = \partial_t \mathbf{U} \cdot \mathbf{T} = -\partial_t \mathbf{T} \cdot \mathbf{U}$. This coefficient is fully determined by how much the tangent tilts toward $\mathbf{U}$: when $\mathbf{T}$ acquires a component along $\mathbf{U}$ (the orange dotted projection $\partial_t \mathbf{T} \cdot \mathbf{U}$ in Fig. 3), the vector $\mathbf{U}$ must compensate by tilting in the $-\mathbf{T}$ direction to stay orthogonal (purple arrow in Fig. 3).

Coefficient c_3 (along $\mathbf{V}$). The remaining coefficient $c_3 = \partial_t \mathbf{U} \cdot \mathbf{V}$ represents the rate of rotation of the frame *within* the normal plane, i.e., the twist. Crucially, this coefficient is *not* fixed by orthonormality: any value of c_3 yields a valid orthonormal frame. Setting it to zero defines the *zero-twist* or *rotation-minimizing* transport [2]:

$$c_3 = \partial_t \mathbf{U} \cdot \mathbf{V} = 0. \quad (9)$$

With $c_2 = 0$ and $c_3 = 0$, Eq. (8) reduces to (and by the same argument for $\mathbf{V}$):

$$\partial_t \mathbf{U} = -(\partial_t \mathbf{T} \cdot \mathbf{U}) \mathbf{T}, \quad \partial_t \mathbf{V} = -(\partial_t \mathbf{T} \cdot \mathbf{V}) \mathbf{T}. \quad (10)$$

The frame vectors $\mathbf{U}$ and $\mathbf{V}$ change only to the extent necessary to remain orthogonal to the evolving tangent, and specifically, they do not rotate around it. Given an arbitrary initial frame $\{\mathbf{U}(s, t_0), \mathbf{V}(s, t_0)\}$ at each material point, Eq. (10) uniquely determines the frame for all subsequent times.

In discretized form, the temporal transport over a time step Δt reduces to the *unique rotation of smallest angle* that aligns the old tangent with the new:

$$R_s^{t \rightarrow t + \Delta t} : \mathbf{T}(s, t) \mapsto \mathbf{T}(s, t + \Delta t). \quad (11)$$

This rotation acts about the axis $\mathbf{T}(s, t) \times \mathbf{T}(s, t + \Delta t)$ through the angle between the two tangents; when $\mathbf{T}$ does not change, R is the identity. Applying it to the frame vectors yields:

$$\mathbf{U}(s, t + \Delta t) = R_s^{t \rightarrow t + \Delta t} \mathbf{U}(s, t), \quad \mathbf{V}(s, t + \Delta t) = R_s^{t \rightarrow t + \Delta t} \mathbf{V}(s, t). \quad (12)$$

Because the minimal rotation has no component about $\mathbf{T}$, the transported frame does not twist around the tangent direction over time. This property of the minimal rotation frame transport (RMF transport) will later be the key property of the definition of the spatial deformation induced by the Lagrangian coreline: any angular change of a pathline's displacement in the normal plane, measured in the transported frame, reflects the genuine physical winding of the particle around the core and is not an artifact of how the frame is propagated. Note that this transport operates *independently* at each material point s , and does not depend on any spatial coherence of the frame along the coreline apart from a well-defined tangent.

Invariance under rigid body transformations. A central requirement for any meaningful vortex analysis is *objectivity*: the results must not depend on the observer's frame of reference. However, since our method works with material particles only, and their positions are invariant to reference frame transformations, this invariance carries over to our frame transport and hence to our space deformation. In this way, the vortex-centric particle motion and all measurements derived from it are genuine physical properties of the flow.

3.3 The Vortex-Centric Deformation

With the RMF transport over time, we define the tubular coordinates, and hence the deformation that maps the neighborhood of the coreline at any time t to the neighborhood at a reference time t_{ref} .

For a point $\mathbf{x} \in \mathcal{U}_\varepsilon(\mathcal{C}_t)$ with tubular coordinates (s^*, u, v) from (7), the *vortex-centric deformation*

$$\Phi_t^{t_{\text{ref}}} : \mathcal{U}_\varepsilon(\mathcal{C}_t) \longrightarrow \mathcal{U}_\varepsilon(\mathcal{C}_{t_{\text{ref}}}) \quad (13)$$

is defined by preserving the material label s^* and the normal-plane coordinates (u, v) :

$$\Phi_t^{t_{\text{ref}}}(\mathbf{x}) = \gamma(s^*, t_{\text{ref}}) + u \mathbf{U}(s^*, t_{\text{ref}}) + v \mathbf{V}(s^*, t_{\text{ref}}). \quad (14)$$

In summary, we decompose $\mathbf{x}$ into its material position along the core and its normal-plane coordinates (u, v) in the RMF at time t ; then reconstruct the point at t_{ref} using the same material label and the same coordinates, but now expressed in the transported reference-time frame $\{\mathbf{U}(s^*, t_{\text{ref}}), \mathbf{V}(s^*, t_{\text{ref}})\}$. Because the frame is rotation-minimizing, this deformation is twist-free: it introduces no spurious angular displacement of the mapped point relative to the core.

The map (Eq. 14) is a diffeomorphism within the tubular neighborhood, smoothly deforming one neighborhood onto another while preserving the radial and angular structure. This has two immediate consequences that are central to our visualization analysis method:

Pathline visualization relative to the core. By mapping pathline samples $\{\mathbf{x}(t_i)\}_{i=0}^N$ to a common reference time t_{ref} via $\Phi_{t_i}^{t_{\text{ref}}}$, we factor out the global translation and stretching of the coreline and isolate the local swirling dynamics.

Consistent winding angle measurement. The angular coordinate of a particle within its frame, $\theta = \text{atan2}(v, u)$, is invariant under changes to the reference time t_{ref} . While this measurement is a kinematic quantity and therefore inherently frame-dependent (meaning it accumulates precisely the angular velocity of the chosen reference frame), the change $\Delta\theta$ between consecutive pathline samples accurately isolates the local swirling motion of the particle around the coreline from the vortex's bulk translation and deformation. This guarantees that we can measure and compare winding angles in a physically consistent way for any given reference frame at any reference time.

3.4 Discretization

We now describe the discrete realization of the frame transport and the associated deformation of the tubular neighborhood, and describe how we use this to transform particles over time into a common frame of reference. The coreline $\mathcal{C}_t$ is represented by an ordered set of K vertices $\{\mathbf{c}_k(t)\}$, $k \in \{0, \dots, K-1\}$, each advected as a Lagrangian particle. These vertices are the discrete counterparts of the material parameterization $\gamma(s_k, t)$. We advect the coreline for N time steps, producing discrete time levels $t_0, t_1, \dots, t_{N-1}$.

3.4.1 Precomputed cumulative RMF rotations

As the coreline evolves over time, it is subject to translation and, more importantly, deformation. This translates to a temporal rotation of the local vortex-centric reference frames. Since the reference frame of the core depends only on its tangent, it suffices to track the temporal evolution of the core tangent at each vertex, which is the integral of the infinitesimal rotations shown in Eq. 10.

To transport particles into a common reference frame, a naïve implementation is to map the discrete timestep t to any timestep t_{ref} , as a sum of $t_{\text{ref}} - t$ sequential, incremental rotations. Since the incremental rotations depend only on the core's tangent evolution at each vertex index k , it is possible to pre-compute the *cumulative* rotation from the initial time t_0 to every timestep $t_j \in [t_0, t_{N-1}]$, for each $\mathbf{c}_k$. Denoting the incremental rotation from t_n to t_{n+1} at $\mathbf{c}_k$ as $R_k^{t_n \rightarrow t_{n+1}}$, the cumulative rotation $t_0 \rightarrow t_j$ is:

$$Q_k(t_j) = R_k^{(j-1) \rightarrow j} \circ \dots \circ R_k^{1 \rightarrow 2} \circ R_k^{0 \rightarrow 1} \quad (15)$$

with $Q_k(t_0) = I$ (the identity). In this way, for each $\mathbf{c}_k(t_s)$ at a starting time $t_s \in [t_0, t_{N-1}]$, the cumulative rotation to any final target time $t_f \in [t_0, t_{N-1}]$ reduces to:

$$Q_k^{t_s \rightarrow t_f} := Q_k(t_f) Q_k(t_s)^{-1} \quad (16)$$

which composes a single rotation in $O(1)$, reducing the total cost from $O(P^2)$ to $O(P)$ per pathline.

We represent each cumulative rotation as a unit quaternion $\mathbf{q}_k(t_n) \in \mathbb{H}$, where the incremental quaternion

$$\delta \mathbf{q} = \left(\cos \frac{\phi}{2}, \hat{\mathbf{a}} \sin \frac{\phi}{2} \right) \quad (17)$$

is constructed from the axis-angle pair

$$\hat{\mathbf{a}} = \frac{\mathbf{T}_i \times \mathbf{T}_{i+1}}{\|\mathbf{T}_i \times \mathbf{T}_{i+1}\|}, \quad \phi = \text{atan2}(\|\mathbf{T}_i \times \mathbf{T}_{i+1}\|, \mathbf{T}_i \cdot \mathbf{T}_{i+1}). \quad (18)$$

Quaternion composition $\mathbf{q}_k(t_{n+1}) = \delta \mathbf{q} \cdot \mathbf{q}_k(t_n)$ accumulates the chain, and the inverse in (16) is the quaternion conjugate $\mathbf{q}^{-1} = \bar{\mathbf{q}}$, which is exact for unit quaternions. This representation stores 4 floats per vertex per timestep, and quaternion composition preserves unitarity up to floating-point precision, making it numerically stable over long chains without periodic re-orthogonalization.

The precomputation is performed once per Lagrangian coreline, in $O(K \cdot N)$ time for K coreline vertices and N timesteps. Now, we will demonstrate the usage of the cumulative rotations for the vortex-centric transformation for a given point $(\mathbf{x}, t)$, e.g., a pathline sample.

3.4.2 Twist-free temporal transformation

Using the precomputed cumulative rotations from Sec. 3.4.1, we transform a pathline sample $(\mathbf{x}, t)$ to a reference time t_{ref} in three steps:

Index-based longitudinal coordinate: We find the closest point $\mathbf{c}$ on the coreline $\mathcal{C}_t$ to $\mathbf{x}$. Since we work with Lagrangian corelines, the discrete coreline vertices $\mathbf{c}_k$ correspond to material particles, and their indices naturally serve as a temporally stable discrete longitudinal coordinate by explicitly tracking material correspondence across time, ensuring that the transformed space remains anchored to the material evolution of the vortex.

Local coordinate extraction. We distinguish the two geometric cases: the closest point is either a vertex of the coreline or lies on a line segment of the coreline. If $\mathbf{c}$ coincides with a coreline vertex $\mathbf{c}_k$, we project $\mathbf{x}$ onto the two normal planes at $\mathbf{c}_k$ defined by the tangents of the two adjacent line segments. If $\mathbf{c}$ lies on a segment between $\mathbf{c}_k$ and $\mathbf{c}_{k+1}$, we project $\mathbf{x}$ onto the normal planes at both vertices. In both cases we compute the radial distances (r_k, r_{k+1}) , and the unit displacement vectors $(\mathbf{e}_k, \mathbf{e}_{k+1})$ in the two normal planes, as well as an interpolation weight $\beta \in (0, 1)$ as the normalized distance to the normal plane at $\mathbf{c}_k$. The vectors $\mathbf{e}_k$ and $\mathbf{e}_{k+1}$ will be transported over time using our frame transport method.

Frame transport: We transport the displacement vectors to t_{ref} using (16): $\mathbf{e}^* = Q_k^{t \rightarrow t_{\text{ref}}} \mathbf{e}$. Both $\mathbf{e}_k$ and $\mathbf{e}_{k+1}$ are transported independently at their respective vertex indices.

Position at Reference Time: The transformed position $\mathbf{x}^*$ in the neighborhood of the coreline $\mathcal{C}_{t_{\text{ref}}}$ at time t_{ref} is computed as a linear interpolation of the transported points in the normal planes:

$$\mathbf{x}^* = (1 - \beta) (\mathbf{c}_k(t_{\text{ref}}) + r_k \mathbf{e}_k^*) + \beta (\mathbf{c}_{k+1}(t_{\text{ref}}) + r_{k+1} \mathbf{e}_{k+1}^*) \quad (19)$$

3.5 Laplacian Smoothing of Corelines

We represent vortex corelines as piece-wise linear polylines and compute tangents at vertices as the central difference of the positions of the neighboring vertices. Some coreline extraction methods enforce smoothness constraints on the vertex positions, leading to well-behaved polylines. Nonetheless, we use coreline extraction algorithms as black boxes and found that, in practice, not all corelines are smooth. Since our coordinate system relies on the coreline tangent and closest-point projection, sharp spatial oscillations along the coreline produce erratic frame orientations and noisy local coordinates. Therefore, we apply an iterative Laplacian smoothing approach to the extracted corelines as a pre-processing step.

Given a coreline $\mathcal{C}_t = \{\mathbf{c}_0, \dots, \mathbf{c}_{K-1}\}$ with $K \geq 3$, a smoothing factor $\alpha \in (0, 1]$, and a number of iterations L , each iteration updates every interior vertex as:

$$\mathbf{c}_k^{(\ell+1)} = \mathbf{c}_k^{(\ell)} + \alpha \left(\frac{\mathbf{c}_{k-1}^{(\ell)} + \mathbf{c}_{k+1}^{(\ell)}}{2} - \mathbf{c}_k^{(\ell)} \right), \quad k = 1, \dots, K-2 \quad (20)$$

where the superscript ℓ denotes the vertex positions at the beginning of the current iteration. The endpoints $\mathbf{c}_0$ and $\mathbf{c}_{K-1}$ are held fixed to preserve the coreline extent. The parameter α controls the step size: $\alpha = 1$ performs full averaging, while smaller values damp the displacement toward the midpoint. The smoothing removes the extraction noise while preserving the large-scale geometry of the coreline.

3.6 Winding Angle and Color Mapping

Once pathlines are associated with the vortex core, we can quantify the rotational behavior of the flow around it. The cumulative winding angle $\theta(t)$ measures the angle a particle traces around a coreline over time. We compute this angle for each pathline directly by always mapping two consecutive vertices of the pathline into a common frame using the RMF transport as a twist-free reference.

At each timestep t_i , we project the pathline point $\mathbf{x}(t_i)$ onto the normal plane of the coreline at its closest point, and compute the normalized displacement direction $\mathbf{e}(t_i)$. We then RMF-transport the displacement direction $\mathbf{e}(t_{i-1})$ from the previous pathline vertex at the

tangent space at t_{i-1} to that at t_i , yielding a twist-free reference $\mathbf{e}^*_{t_{i-1}}$. The incremental winding is the signed angle between $\mathbf{e}(t_i)$ and $\mathbf{e}^*_{t_{i-1}}$:

$$\Delta\theta_i = \text{sign}(\mathbf{T}_i \cdot (\mathbf{e}^*_{t_{i-1}} \times \mathbf{e}(t_i))) \cdot \arccos(\mathbf{e}^*_{t_{i-1}} \cdot \mathbf{e}(t_i)) \quad (21)$$

where $\mathbf{T}_i$ is the coreline tangent at t_i and the sign is determined by the right-hand rule. The cumulative winding angle for a pathline with K vertices is $\theta_K = \sum_{i=1}^{K-1} \Delta\theta_i$. This formulation measures the particle's intrinsic rotation purely relative to the twist-free frame evolution of the vortex core.

To visualize winding behavior, each pathline is assigned one color based on its accumulated winding angle. We normalize the winding angle across all pathlines and map them to a color map. In some use cases, a common direction for measuring the orientation of rotation exists. For instance, in a vortex street, we can discern counterclockwise from clockwise rotating vortices. In these cases, the sign of the winding angle carries intrinsic meaning which can be visualized, for instance, with a divergent color map. In other cases, the sign does not carry any inherent meaning but is an artifact of the arbitrary ordering of the vertices of the coreline's polyline; hence, we visualize the absolute value of the accumulated winding angle.

4 RESULTS

We evaluate our method on two analytical 3D unsteady flow field with known ground truth (DEFORMING VORTEX), and four Lattice Boltzmann simulations that include a variety of flow phenomena, including a von Kármán vortex street behind an obstacle (HALF CYLINDER), and complex vortex filaments behind a delta wing (DELTA WING) and two aircraft (BOEING, RAPTOR).

Corelines are extracted with different methods that are discussed for each result. Each coreline is then advected using fourth-order Runge-Kutta (RK4) integration. Pathlines are seeded at concentric rings around each coreline vertex at t_0 and advected using RK4. For better reproducibility, our pathline seeding strategy is reported in Appendix B. In all experiments, we used Laplacian smoothing described in Sec. 3.5 with $\alpha = 0.4$ at t_0 before advection. This parameter setting effectively removes extraction noise from the corelines while preserving the coreline geometry. Smaller values of α with larger numbers of smoothing iterations K_{Lap} lead to similar results, while larger values of α may significantly alter the coreline geometry.

The dataset grid sizes, physical simulation domain sizes, experimental parameters, and computation times for each Figure are summarized in Table 2, Appendix C.

4.1 Performance

All results were obtained using a workstation with an AMD Ryzen Threadripper PRO 3975WX 32-core processor and an NVIDIA GeForce RTX 4090 GPU. The method was implemented in C++ using standard OpenMP parallelization [6]. To quantify the computational complexity of our method, we split it into a pre-computation step and an on-the-fly transformation step. The pre-computation step (S1) consists of coreline advection and cumulative RMF computation. The transformation step (S2 in Table 2, Appendix C) consists of pathline integration, winding angle computation, and pathline transformation. Pre-computing cumulative RMF rotations reduces the temporal frame transport to $O(1)$, making pathline transformation and winding angle computation significantly faster compared to an incremental approach. Further details about the runtime scaling of our method are presented in Appendix C, which also includes the runtimes for the paper results reported in seconds in Table 2.

4.2 Quantitative Validation via Rotating Reference Frames

To quantitatively evaluate the winding angle computation using our Rotation-Minimizing Frame (RMF), we construct an analytically defined ground truth dataset. We utilize the 2D four-centers function [9], which transforms a steady flow field of two counter-clockwise and two clockwise rotating centers into a rotating reference frame. Following Günther et al. [9] we construct an unsteady dynamic ground truth by




Figure 4: Ground truth validation using the analytical four-centers. (a) Pathlines integrated in the unsteady dataset (subjected to a global reference frame rotation Ω) for a time window of 2π (one full domain rotation). (b) Pathlines transformed into our vortex-centric RMF visualization. The local swirling motion is cleanly decoupled from the global drift, perfectly recovering the expected $\theta_{\text{steady}} + 2\pi$ (red pathlines) and $\theta_{\text{steady}} - 2\pi$ (cyan pathlines) kinematic shift.

introducing a global rotational reference frame transformation. The entire domain is rotated around its origin with a constant angular velocity Ω , such that the new time-dependent velocity field becomes unsteady.

In this experiment, it is crucial to distinguish between the objectivity of the vortex coreline and the inherent frame-dependence of material transport. While the extracted corelines represent frame-independent structural features, the accumulated winding angle of a material pathline is a kinematic measure that must be computed relative to a reference frame. In the steady frame, the pathlines exhibit a winding angle of θ_{steady} . When integrated within the unsteady, rotating frame over a time window Δt , the true kinematic winding angle must analytically reflect the global rotation, yielding:

$$\theta_{\text{unsteady}} = \theta_{\text{steady}} \pm \Omega \Delta t \quad (22)$$

For instance, for exactly one full domain rotation ($\Omega \Delta t = 2\pi$), the winding angles of the pathlines must shift by exactly 2π , adding to the co-rotating cores and subtracting from the counter-rotating cores.

As shown in Fig. 4, our RMF-integrated pathlines perfectly capture this exact kinematic shift. Our vortex-centric visualization recovers the pure local swirling motion, perfectly matching the analytical $\pm 2\pi$ difference, expected from the ground truth. The pathlines show the exact trajectory of the particles relative to their respective coreline, which are not perfectly circular as can be seen in Fig. 4. The exact numerical mapping confirms that the RMF isolates the true local flow kinematics and serves as a reliable quantitative tool in a given reference frame.

4.3 Comparison to Streamline-Based Approaches

When visualizing evolving vortex cores, one aspect of interest is the deformation of the vortex coreline itself. Another interesting aspect is the visualization of the local vortex dynamics. While visualizing the temporal evolution of a vortex can be achieved by showing a sparse set of coreline samples from different times, the visualization of the fluid dynamics relative to these moving cores is more difficult. To illustrate the differences between instantaneous streamlines and our vortex-centric pathlines, we design an analytical unsteady flow field featuring a vortex core that translates and deforms over time. In addition to a base translation, the convective speed of the vortex depends on the z -coordinate such that the central z -region moves faster, causing the vertical vortex tube to develop a curved bulge shape over time.

Therefore, the speed of an observer that co-moves with the vortex also depends on the z -position. The observer's speed is:

$$v_{\text{eff}}(z) = v_{\text{base}} + v_{\text{def}} e^{-\left(\frac{z^2}{2\sigma_z^2}\right)}, \quad (23)$$

where we set the base speed to $v_{\text{base}} = 0.5$, the extra speed to $v_{\text{def}} = 0.15$, and the scaling factor to $\sigma_z = 0.5$. This speed profile defines the location of the vortex center at each z -height and time $t \in [0, 10]$ as $\mathbf{c}(z, t) = (v_{\text{eff}}(z)t, 0, z)^T$. The complete 3D analytical velocity field $\mathbf{v}(x, y, z, t)$ of the deforming vortex incorporates this translation into a Gaussian vortex model with strength $\Gamma = 2.0$ and width $\sigma = 0.8$:

$$\mathbf{v}(x, y, z, t) = \begin{pmatrix} -\Gamma(y - c_y)e^{-\frac{r^2}{2\sigma^2}} + v_{\text{eff}}(z) \\ \Gamma(x - c_x)e^{-\frac{r^2}{2\sigma^2}} \\ 0 \end{pmatrix}, \quad (24)$$

where $(c_x, c_y) = (v_{\text{eff}}(z)t, 0)$ and the squared radial distance to the coreline is $r^2 = (x - c_x)^2 + (y - c_y)^2$. For simplicity, the flow is 2D rotational in the xy -plane and lacks a velocity component in the z direction ($w = 0$).

While streamline-based visualizations inherently require rendering streamlines at different time snapshots (potentially leading to visual clutter), our vortex-centric representation shows the full temporal dynamics compactly aligned at one single reference time.

The top row of Fig. 5 shows the method of Zhang et al. [31], who choose to visualize streamlines relative to one globally optimal rigid body motion observer. Fig. 5(a) shows streamlines computed relative to the observer that co-moves with the slowest part of the vortex, i.e., using the field $\mathbf{v}_{\text{obs_min}} = (v_{\text{base}}, 0, 0)^T$. Fig. 5(b) shows streamlines relative to the observer that co-moves with the fastest part of the vortex, $\mathbf{v}_{\text{obs_max}} = (v_{\text{base}} + v_{\text{def}}, 0, 0)^T$. Each of these rigid transformations has a clear physical interpretation. However, seeding streamlines with a constant distance from the vortex core in the respective observed field does not always result in perfectly matching streamline swirls across




Figure 5: Comparison of visualization approaches on the analytical deforming vortex dataset. We show corelines for the times $t = 0, 2, 4, 6, 8, 10$. The top row shows streamlines seeded near the coreline in a field observed by a rigid body motion observer. (a) Streamlines visualized relative to the observer co-moving with the slowest parts $\mathbf{v}_{\text{obs_min}}$ [31] of the vortex. (b) Streamlines visualized relative to the observer co-moving with the fastest parts $\mathbf{v}_{\text{obs_max}}$ of the vortex. (c) Streamlines seeded using infinitely many observers co-moving at speeds matching the coreline velocity locally [11]. (d) Our vortex-centric pathlines ($t_0 = 0$, $N = 10$ timesteps), which decouple the global deformation of the vortex from the swirling motion.




Figure 6: DELTA WING dataset, $t_0 = 0$, three corelines (green) extracted using adaptive LORF, visualized at two different times. For the two upper corelines, we compute our transformed pathlines (blue, $N = 100$) and streamlines (orange) integrated along the observed field from the adaptive LORF, using 400 steps (forward and backward, step size 0.001). For the chosen seeding and integration settings, both methods depict rotating motion around the upper cores. Zhang et al. [31] requires manual selection of reference frames.

the varying deformation of the core.

Fig. 5(c) displays the approach of Günther and Theisel [11], which seeds streamlines in an idealized observed field corresponding to the simultaneous measurement from infinitely many observers moving at all speeds ranging from $v_{\text{obs_min}}$ to $v_{\text{obs_max}}$. Specifically, their local observer velocity aligns perfectly with the coreline movement: $\mathbf{v}_{\text{obs_ideal}}(z) = (v_{\text{eff}}(z), 0, 0)^T$. This leads to streamline swirls that neatly correspond to the corelines in this carefully controlled analytical setup. However, it raises a significant interpretability challenge in generic flow cases where the space-time continuum of co-moving observers both translates and rotates in arbitrary ways. Without visualizing the underlying reference frames, it is unclear exactly how this approach should be physically interpreted.

In Fig. 5(d), we show our vortex-centric pathlines. Compared to instantaneous streamline-based methods, there are several key benefits: **Temporal integration over snapshots:** To visualize the vortex dynamics, we can choose a single reference time point and map the entirely transformed pathlines into a coreline-centered coordinate system at that time. Although we show the results for multiple time steps in Fig. 5(d) to demonstrate that they all perfectly match, rendering one timestep would be sufficient to analyze the local swirling behavior of the pathlines.

Meaningful physical quantities: Since our method operates on advected Lagrangian pathlines instead of instantaneous integration times, quantities such as the swirling strength are physically meaningful material properties rather than an artifact depending on a parameter such as the streamline integration time.

Capturing flow relative to the core: In complex flow fields, particles move along the coreline, swirl in- or outward over time, and have varying swirling strength depending on the distance to the coreline. This leads to them advancing either faster or slower in a longitudinal or radial direction relative to the vortex core. Such phenomena cannot be observed when analyzing instantaneous streamlines, but become visible with our approach.

In Fig. 6, we compare our method to streamlines transformed into a co-moving reference frame, as presented by Zhang et al. [31], using the DELTA WING simulation. We illustrate the necessity of multiple reference frames by visualizing streamlines relative to a single frame computed to fit the coreline’s movement at the tip of the vortex core at time t_0 . As shown in the left panel of Fig. 6, maintaining this frame over time results in significant drift, causing streamlines to deviate from the vortex core. This limitation is the reason Zhang et al. [31] utilize

different reference frames for various time steps. While this approach allows for the verification of swirling motion within localized frames, it requires users to manually identify and fit these frames, which can introduce artifacts and unwanted drift in other regions. In contrast, our unified Lagrangian approach avoids these complications, providing a comprehensive, vortex-centric visualization of the coreline evolution without the need for manual frame selection. We show the results of our method for different reference times on the right of Fig. 6.

4.4 Evaluation of Vortex Core Extraction Methods

Our method can be used to evaluate and analyze coreline extraction methods. In Fig. 7 we show extracted vortex cores from the DELTA WING dataset using two approaches: parallel vectors on the input field (lab frame) shown in the top row of Fig. 7, and parallel vectors on the observed field using adaptive locally optimal reference frame fields [22] shown on the bottom row of Fig. 7 (adaptive LORF). We show pathlines seeded around the extracted corelines visualized in the lab frame (left column of Fig. 7) and transformed pathlines with our method (middle column of Fig. 7). Our method removes the temporal complexity from the visualization, clearly revealing the rotation of particles around the extracted cores, and minimizing the pathline clutter compared to visualizations of the original pathlines. We observe that the locally optimal reference frame field enables the extraction of the complete vortex filament and two additional vortex cores (Zoom-ins on the right of Fig. 7). Furthermore, the winding angle in the vortex-centric reference frame provides strong clues where the parallel vectors method fails in the lab frame. While the winding angles largely match




Figure 7: Delta wing dataset, $t_0 = 0$, $N = 100$. We extracted corelines (dark green) using vtkVortexCore [25] from the input field (lab frame, first row), and from the observed velocity obtained via adaptive locally optimal reference frames (LORF) [22] (second row) with the instantaneous vorticity deviation (IVD) [14] as a guidance field. For both cases, we seed pathlines (left) around the extracted cores and transform them using our method (right). The third row displays zoom-ins of the transformed pathlines. The winding angle was computed in the vortex-centric reference frame and used for pathline color-coding with global (all cores) absolute winding angle normalization.




Figure 8: Multiple corelines extracted from HALF CYLINDER at $t_0 = 8.4$, using adaptive LORF. Winding angle color-coding for pathlines integrated over a time-window $\Delta t = 3$ ($N = 30$) enables the visual comparison of rotation direction and rotation strength of the vortex cores.

where both coreline extraction methods successfully detected a vortex core, we see very high winding angles at the tip of the vortex and very low winding angles on the two additional vortex corelines extracted with the adaptive LORF method.

4.5 Von Kármán Vortex Street

In a von Kármán vortex street, vortices are shed periodically from alternating sides of an obstacle, which makes neighboring vortices rotate in opposite directions. In Fig. 8 we use the signed winding angle for color-coding to visualize this phenomenon. In addition to the direction of rotation, we also visualize quantitative measures that show how the vortices evolve. By integrating pathlines over a short time-window ($\Delta t = 3$), we can observe how their swirling strength changes over time. As can be seen in Fig. 8, the swirling strength decreases farther away from the obstacle.

In Fig. 9 we show the full temporal evolution of one vortex coreline as it undergoes significant deformation. We seed particles equidistantly spaced along the coreline at $t = 0$. Interestingly, when the vortex coreline is far away from the obstacle, it deforms significantly, and the transformed pathlines move upward along the Z direction, indicating that the material particles moved along the coreline.

4.6 Analysis of Vortex Dynamics

We extracted vortex corelines from the BOEING flow using locally adaptive reference frame fields with IVD [22]. We visualize pathlines in the vortex-centric reference frame for 78 corelines, as depicted in Fig. 1. We observe multiple long corelines, for instance, long vortex filaments behind the wing surrounded by arc-shaped cores, and multiple cores shed from the tail. In the insets on the right, we highlight the winding angles that change along the filaments. Even for a large number of cores, strongly swirling regions can clearly be discerned from low-rotating regions, quantitative information that is more relevant than simply the presence of a coreline.

For the RAPTOR dataset, at $t = 0.095$, we extracted corelines from the observed field obtained via LORF [8]. The coreline surfaces (top) reveal highly symmetrical fluid behavior. Visualizing pathlines in the lab frame (bottom-left of Fig. 10) fails to reveal the rotational patterns of particle motion in the vortices' vicinity. Our approach (bottom-right of Fig. 10) enables a visualization of pathlines in the vortex-centric frame and reveals the regions of highest rotational strength.




Figure 9: Coreline extracted from HALF CYLINDER at time $t_0 = 8.4$, using adaptive LORF. We seed pathlines close to the coreline, integrate for $N = 50$ time steps, and transform them to different reference times. Away from the obstacle, the vortex coreline undergoes severe deformation, and the particles congregate in the upper portion of the core.




Figure 10: Corelines extracted from the dataset RAPTOR at $t_0 = 0.095$ using uniform LORF [8], and advected for $N = 35$ timesteps. The results are strongly symmetric, as seen in the coreline surfaces (top). For the same time window, the bottom-left and bottom-right images show original and transformed pathlines, respectively. The visualization of original pathlines makes the localization of strong particle swirling around the extracted cores difficult. After transformation with our method, the intricate vortex dynamics become clearly visible.

5 CONCLUSIONS, LIMITATIONS AND FUTURE WORK

Our vortex-centric reference frame transformation enables the visual analysis of rotating particle motion around vortex cores. Our rotation-minimizing frame transport method ensures that no artificial rotation about the coreline's tangent is introduced. We demonstrate that our approach can serve as an evaluation tool for coreline extraction methods, as it facilitates the visual confirmation of swirling particle motion around coreline candidates. Furthermore, we demonstrate that the vortex-centric reference frame can be used to compute pathline winding angles, enabling the quantitative and visual identification of specific rotational behavior.

The method provides a smooth and diffeomorphic coordinate system only in the *tubular neighborhood* of the coreline, i.e., it is not defined for the space outside the *reach* of the coreline and restricted in regions of high curvature. In such regions, the orthogonal planes from different coreline points intersect, causing the coordinate mapping to become multi-valued. While we have not experienced this as a strong limitation when analyzing pathlines close to the coreline, we plan to compute neighborhood geometries and use them to improve usability.

Furthermore, we want to exploit our framework as a versatile tool for different kinds of vortex dynamics visualizations through application-specific color coding and animation techniques. For instance, a time-based color mapping on pathlines will facilitate the direct visualization of particle motion relative to the coreline, allowing experts to analyze how flow evolves along the vortex trajectory over time. Additionally, color-coding local swirling strength will reveal sub-filament variations and variations over time, pinpointing regions of intensified vortex dynamics. Together, these visualization strategies will enable experts to identify specific temporal events and structural evolutions within complex, unsteady fluid environments.

Further, as a future research direction, we anticipate that an objective choice of reference frame could lead to a reference frame invariant measure of the relative winding angle. This novel measurement could then be used to formulate an inverse problem to find maxima of relative rotation for the robust and reference frame independent extraction of vortex corelines.

ACKNOWLEDGMENTS

This work was supported by King Abdullah University of Science and Technology (KAUST) baseline funding.

REFERENCES

- [1] P. Angelelli and H. Hauser. Straightening tubular flow for side-by-side visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2063–2070, 2011. doi: [10.1109/TVCG.2011.235](#) 3
- [2] R. L. Bishop. There is more than one way to frame a curve. *The American Mathematical Monthly*, 82(3):246–251, 1975. doi: [10.1080/00029890.1975.11993807](#) 4
- [3] R. Bujack, L. Yan, I. Hotz, C. Garth, and B. Wang. State of the art in time-dependent flow topology: Interpreting physical meaningfulness through mathematical properties. *Computer Graphics Forum*, 39(3):811–835, 2020. doi: [10.1111/cgf.14037](#) 2
- [4] M. S. Chong, A. E. Perry, and B. J. Cantwell. A general classification of three-dimensional flow fields. *Physics of Fluids*, 2:765–777, 1990. 2
- [5] O. Daae Lampe, C. Correa, K.-L. Ma, and H. Hauser. Curve-centric volume reformation for comparative visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1235–1242, 2009. doi: [10.1109/TVCG.2009.136](#) 3
- [6] L. Dagum and R. Menon. Openmp: an industry standard api for shared-memory programming. *IEEE Computational Science and Engineering*, 5(1):46–55, 1998. doi: [10.1109/99.660313](#) 6
- [7] T. Günther and I. Baeza Rojo. Introduction to vector field topology. In *Topological Methods in Data Analysis and Visualization VI*, pp. 289–326. Springer, Cham, Nyköping, Sweden, 2021. doi: [10.1007/978-3-030-83500-2_15](#) 2
- [8] T. Günther, M. Gross, and H. Theisel. Generic objective vortices for flow visualization. *ACM Transactions on Graphics*, 36(4):141:1–141:11, 2017. doi: [10.1145/3072959.3073684](#) 2, 9
- [9] T. Günther, M. Schulze, and H. Theisel. Rotation invariant vortices for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):817–826, 2016. doi: [10.1109/TVCG.2015.2467200](#) 6
- [10] T. Günther and H. Theisel. The state of the art in vortex extraction. *Computer Graphics Forum*, 37(6):149–173, 2018. doi: [10.1111/cgf.13319](#) 2
- [11] T. Günther and H. Theisel. Objective lagrangian vortex cores and their visual representations. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):76–85, 2025. doi: [10.1109/TVCG.2024.3456384](#) 1, 2, 3, 7, 8
- [12] M. Hadwiger, M. Mlejnek, T. Theußl, and P. Rautek. Time-dependent flow seen through approximate observer killing fields. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):1257–1266, 2019. doi: [10.1109/TVCG.2018.2864839](#) 2
- [13] G. Haller. An objective definition of a vortex. *Journal of Fluid Mechanics*, 525:1–26, 2005. doi: [10.1017/S0022112004002526](#) 2
- [14] G. Haller, A. Hadjighasem, M. Farazmand, and F. Huhn. Defining coherent vortices objectively from the vorticity. *Journal of Fluid Mechanics*, 795:136–173, 2016. doi: [10.1017/jfm.2016.151](#) 2, 8
- [15] J. C. Hunt, A. A. Wray, and P. Moin. Eddies, streams, and convergence zones in turbulent flows. In *Proceedings of the Summer Program*, pp. 193–208. Center for Turbulence Research, 1988. 2
- [16] J. Jeong and F. Hussain. On the identification of a vortex. *Journal of Fluid Mechanics*, 285:69–94, 1995. doi: [10.1017/S0022112095000462](#) 2
- [17] P. Joshi, M. Meyer, T. DeRose, B. Green, and T. Sanocki. Harmonic coordinates for character articulation. In *ACM SIGGRAPH 2007 Papers*, SIGGRAPH ’07, pp. 71–es. Association for Computing Machinery, New York, NY, USA, 2007. doi: [10.1145/1275808.1276466](#) 3
- [18] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and M. E. Gröller. Advanced Curved Planar Reformation: Flattening of Vascular Structures. In *IEEE Visualization 2003*, pp. 43–50, Oct. 2003. 3
- [19] J. M. Lee. *Introduction to Smooth Manifolds*, vol. 218 of *Graduate Texts in Mathematics*. Springer, 2nd ed., 2013. doi: [10.1007/978-1-4419-9982-5](#) 4
- [20] Y. Lipman, D. Levin, and D. Cohen-Or. Green coordinates. *ACM Trans. Graph.*, 27(3):1–10, Aug. 2008. doi: [10.1145/1360612.1360677](#) 3
- [21] H. J. Lugt. The dilemma of defining a vortex. In *Recent Developments in Theoretical and Experimental Fluid Mechanics*, pp. 309–321. Springer, 1979. doi: [10.1007/978-3-642-67220-0_32](#) 2
- [22] J. R. Ramirez, P. Rautek, T. Günther, and M. Hadwiger. Locally adapted reference frame fields using moving least squares. *IEEE Transactions on Visualization and Computer Graphics*, 32(1):495–505, 2026. doi: [10.1109/TVCG.2025.3634845](#) 2, 8, 9
- [23] P. Rautek, M. Mlejnek, J. Beyer, J. Troidl, H. Pfister, T. Theußl et al. Objective observer-relative flow visualization in curved spaces for unsteady 2D geophysical flows. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):283–293, 2021. doi: [10.1109/TVCG.2020.3030454](#) 2, 3
- [24] P. Rautek, X. Zhang, B. Woschizka, T. Theußl, and M. Hadwiger. Vortex lens: Interactive vortex core line extraction using observed line integral convolution. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):55–65, 2024. doi: [10.1109/TVCG.2023.3326915](#) 2
- [25] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit (4th ed.)*. Kitware, 2006. 8
- [26] A. Straub, F. Sadlo, and T. Ertl. Feature-based deformation for flow visualization. *Journal of Visualization*, 27:795–817, oct 2024. doi: [10.1007/s12650-024-00963-5](#) 3
- [27] D. Sujudi and R. Haimes. Identification of swirling flow in 3-d vector fields. In *Proceedings of the 12th Computational Fluid Dynamics Conference*, pp. 792–799, 1995. doi: [10.2514/6.1995-1715](#) 2
- [28] J.-M. Thiery, P. Memari, and T. Boubekeur. Mean value coordinates for quad cages in 3d. *ACM Trans. Graph.*, 37(6), art. no. 229, 14 pp., Dec. 2018. doi: [10.1145/3272127.3275063](#) 3
- [29] X. Zhang, A. Ageeli, T. Theußl, M. Hadwiger, and P. Rautek. Exploring 3d unsteady flow using 6d observer space interactions. *IEEE Transactions on Visualization and Computer Graphics*, 32(1):517–526, 2026. doi: [10.1109/TVCG.2025.3642506](#) 2, 3
- [30] X. Zhang, M. Hadwiger, T. Theußl, and P. Rautek. Interactive exploration of physically-observable objective vortices in unsteady 2D flow. *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS 2021)*, 28(1):281–290, 2022. doi: [10.1109/TVCG.2021.3115565](#) 2, 3
- [31] X. Zhang, P. Rautek, and M. Hadwiger. Generic variational spacetime optimization of vortex core manifolds. In *Proceedings of ACM SIGGRAPH (to appear)*, 2026. 1, 2, 7, 8
- [32] X. Zhang, P. Rautek, T. Theußl, and M. Hadwiger. Enhancing Material Boundary Visualizations in 2D Unsteady Flow through Local Reference Frame Transformations. *Computer Graphics Forum (Proc. Eurovis 2025)*, 2025. doi: [10.1111/cgf.70128](#) 3

A APPENDIX A: THEORETICAL FOUNDATIONS OF THE TUBULAR NEIGHBORHOOD

In this work, the formulation of the coreline is based on the definition of a tubular neighborhood situated within the ambient manifold $M = \mathbb{R}^3$. Let the coreline be modeled as a smooth, embedded, one-dimensional submanifold, represented by a unit-speed curve $\gamma : I \rightarrow \mathbb{R}^3$. At any arbitrary evaluation point $p = \gamma(s)$ along the curve, the tangent space to the submanifold constitutes a one-dimensional linear subspace $T_p C \subset T_p \mathbb{R}^3$, spanned entirely by the unit tangent vector $\mathbf{t}(s) = \gamma'(s)$.




Figure 11: Geometric construction of a tubular neighborhood U (blue cylinder) around a coreline curve γ (green). At any evaluation point p , the curve's tangent vector $\mathbf{t}(s)$ (yellow) uniquely defines the orthogonal normal plane $N_p C$ (orange).

Using the Riemannian metric of the ambient space $\mathbb{R}^3$, we define a strict orthogonal decomposition of the ambient tangent space as $T_p \mathbb{R}^3 \cong T_p C \oplus N_p C$. The normal space $N_p C$ is the two-dimensional orthogonal complement to the tangent vector, defined mathematically as $N_p C = \{\mathbf{v} \in \mathbb{R}^3 \mid \langle \mathbf{v}, \mathbf{t}(s) \rangle = 0\}$. Geometrically, this forces the tangent of the curve to serve as the unique normal vector to the cross-sectional normal plane at each point. The disjoint union of these distinct normal spaces across the entire spatial extent of the curve forms a smooth rank-2 vector bundle known as the normal bundle, denoted $NC = \bigsqcup_{p \in C} N_p C$. Figure 11 illustrates this construction. The tubular neighborhood U is shown as a cylinder around the coreline curve $\gamma(s)$ (green). A normal plane $N_p C$ (orange) intersects the curve at an evaluation point p , and the curve's tangent vector $\mathbf{t}(s)$ (yellow) is strictly orthogonal to this plane.

To map this theoretical bundle structure into the ambient space to form a realized volumetric tube, we invoke the normal exponential map $E : NC \rightarrow \mathbb{R}^3$, which in Euclidean space simplifies to $E(p, \mathbf{v}) = p + \mathbf{v}$. By the Tubular Neighborhood Theorem [1], the differential of E evaluated at the zero section constitutes a non-singular linear isomorphism. Consequently, by the Inverse Function Theorem, there exists a continuous, strictly positive radius function $\varepsilon : C \rightarrow (0, \infty)$ defining an open neighborhood $V_\varepsilon \subset NC$ upon which E acts as a global diffeomorphism. The image $U = E(V_\varepsilon)$ forms the well-defined, non-intersecting static tubular neighborhood surrounding the coreline curve. If the curve is compact, ε can be tightly restricted to a uniform constant radius bounded by the maximum local curvature and the global geometric reach of the curve.

Figure 12 illustrates the reach and the focal point constraints. It shows a 2D projection of the curve γ (green) with offset neighborhoods to either side of the curve. On one side, the neighborhood (blue band) of the curve is dictated by the reach: as the curve approaches itself, the blue band is strictly limited to half the width of this approach distance. On the other side of the curve, an orange neighborhood band is shown together with the normal spaces $N_p C$ (orange lines extending from the green curve). They intersect at a focal point [3] (for curvature greater zero), which dictates the width of the orange band on this side. Neighborhoods can either be constructed as the global minimum of the distances given by reach and curvature, or by the local minimum which varies along the curve.

Furthermore, our model accommodates the continuous deformation of corelines over a time parameter t . Let this spatial deformation be parameterized by a smooth isotopy of embeddings $\gamma_t : I \rightarrow \mathbb{R}^3$. While γ_t describes the motion of the submanifold itself, extending this motion to the surrounding spatial volume requires the extension of the isotopy




Figure 12: Geometric constraints determining the width of a tubular neighborhood. The maximum radius is bounded by the curve's global reach (blue band) to prevent self-intersection, and by local focal points where adjacent normal spaces $N_p C$ (dark orange lines) intersect due to high curvature.

to the neighborhood [2]. By integrating the time-dependent vector field generated by the deforming curve using partitions of unity, the theorem guarantees the mathematical existence of a globally smooth, compactly supported ambient isotopy $H_t : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ with the initial condition $H_0 = \text{Id}$. Applying this global diffeomorphism to the static tubular neighborhood yields a continuously evolving time-dependent neighborhood $U_t = H_t(U_0)$. This mathematical formulation guarantees that as the curve moves and deforms over time, the volumetric tubular neighborhood surrounding it updates continuously without undergoing topological changes.

B APPENDIX B: PATHLINE AND STREAMLINE SEEDING

For each coreline, we seed one pathline nearby every k -th coreline vertex, a distance r along a direction $\mathbf{d}$ in the plane perpendicular to the coreline tangent $\mathbf{t}$. In the plane, the seeding direction is $\mathbf{d} = \frac{\mathbf{v} \times \mathbf{t}}{|\mathbf{v} \times \mathbf{t}|}$, where $\mathbf{v}$ is the input field velocity at the coreline vertex. For the results in Fig. 6, we use $(0, 1, 0)$ instead of $\mathbf{v}$ to compute $\mathbf{d}$, so that pathlines and streamlines are seeded at the same position, which would not be the case using $\mathbf{v}$ due to the reference frame transformation required for streamline computation. The parameters r and k for all results on simulated data are reported in Table 1.

C APPENDIX C: RUNTIME SCALING

We computed runtime scaling curves by measuring the wall-clock runtime of the two stages described in Sec. 4.1 as we independently vary the three parameters that govern the cost of our method: the number of pathlines P , the length of the advection time window N , and the coreline vertex count K . All timings are raw (non-logarithmic) durations, and each configuration is reported as the best of three repetitions to suppress scheduling noise. Increasing P (Fig. 15) leaves the pre-computation S1 unaffected and grows the transformation S2 linearly, since each additional pathline is seeded, integrated, and transformed independently. Increasing the time window N (Fig. 14) grows both stages linearly: S1 because the coreline advection and cumulative-RMF pre-computation is $O(K \cdot N)$, and S2 because a longer window adds proportionally more samples per pathline. This exposes the true cost unit of the transformation stage: every per-sample operation (RK4 integration, local-coordinate extraction, RMF transport, and the winding-angle increment) is $O(1)$, so S2 scales with the *total number of pathline vertices* $P \cdot N$ rather than with P alone. Isolating the effect of the coreline

Table 1: Pathline and streamline seeding parameters.

Fig.	1	2	7	6	8	9	10
r	0.013	0.035	0.015	0.015	0.04	0.04	0.02
k	1	4	3	5	5	7	5

Table 2: Overview of the dataset grid size and spatiotemporal domain, parameter settings for the results, and runtime measurements. P is the total number of advected and transformed pathlines. K_{Lap} is the number of iterations for the Laplacian smoothing. # cores is the number of corelines. N is the number of advection steps that are used to pre-compute RMF transport. Δt is the integration time span for pathlines. S1 and S2 are bounds for the computation time in seconds.

Fig.	grid size $\times$ time steps	space $\times$ time domain	# cores	P	K_{Lap}	N	Δt	S1 [s]	S2 [s]
1	$225 \times 451 \times 090 \times 199$	$[-2.49, 2.51] \times [-5.00, 5.02] \times [-0.99, 1.01] \times [0, 1.96]$	78	2356	200	60	0.594	< 9	< 4
5	$61 \times 41 \times 11 \times 20$	$[-4.00, 8.00] \times [-4.00, 4.00] \times [-1.00, 1.00] \times [0, 10.0]$	1	39	0	10	1.052	< 1	< 1
7	$110 \times 629 \times 110 \times 234$	$[-0.99, 1.00] \times [-5.71, 5.71] \times [-0.99, 1.00] \times [0, 1.73]$	4	652	200	100	0.742	< 2	< 1
8	$640 \times 240 \times 080 \times 151$	$[-0.50, 7.50] \times [-1.50, 1.50] \times [-0.50, 0.50] \times [0, 15.0]$	6	366	400	30	3.000	< 1	< 1
10	$295 \times 472 \times 148 \times 159$	$[-2.49, 2.51] \times [-5.00, 3.00] \times [-0.99, 1.52] \times [0, 0.19]$	70	3154	200	35	0.042	< 6	< 3

resolution K requires care, because our seeding places pathlines at coreline vertices sampled with a fixed stride; naively increasing K would therefore also increase P and conflate the two effects. To measure the influence of K alone, we arc-length-resample the coreline to the target vertex count while scaling the seeding stride proportionally, so that the number of seeded pathlines P remains approximately constant across the sweep. Under this protocol, S2 still grows roughly linearly with K (Fig. 13): the closest-point search that anchors each pathline sample to its material coreline vertex is a brute-force $O(K)$ scan executed once per pathline vertex, and S1 likewise grows linearly with K . Our scaling measurements delineate the practical limits of the method on substantially larger datasets. The vortex-centric transformation itself is $O(1)$ per pathline sample and never becomes the bottleneck; instead, the runtime is dominated by advecting corelines and pathlines. Consequently, the principal limitation for large CFD data is not compute but *field residency*: both stages require random space-time access to $\mathbf{v}(\mathbf{x}, t)$ over the entire time window, and for grids that exceed available memory this forces out-of-core streaming and renders advection I/O- and bandwidth-bound. The remaining costs scale benignly and are parallelizable: runtime grows linearly with the number of cores (processed independently), with the number of pathlines, and with the time window, while the per-sample closest-point search introduces a linear dependence on the coreline vertex count K that could be reduced to $O(\log K)$ with a spatial acceleration structure.




Figure 13: Runtime scaling as a function of K , the number of coreline vertices. The pathline seeding stride k was scaled proportionally with K to keep the number of pathlines P approximately constant while varying K .

REFERENCES

- [1] J. M. Lee, *Introduction to Riemannian Manifolds*, 2nd ed., Graduate Texts in Mathematics, vol. 176. Cham, Switzerland: Springer International Publishing, 2018. 11
- [2] M. W. Hirsch, *Differential Topology*, Graduate Texts in Mathematics, vol. 33. New York, NY: Springer-Verlag, 1976. 11
- [3] V. Guillemin and A. Pollack, *Differential Topology*. Englewood Cliffs, NJ: Prentice-Hall, 1974. 11




Figure 14: Runtime scaling as a function of the time window length N , the number of advection steps. Both the pre-computation stage S1 (coreline advection and cumulative RMF) and the transformation stage S2 (pathline integration, winding angle, and transformation) grow linearly with N ; the number of pathlines P and the coreline vertex count K are held constant.




Figure 15: Runtime scaling as a function of the number of pathlines P . The transformation stage S2 (pathline integration, winding angle, and transformation) grows linearly with P , while the pre-computation stage S1 is independent of P and is therefore omitted. The time window length N and the coreline vertex count K are held constant.