

2D Vector Field Simplification Based on Robustness

Primož Skraba*
Jozef Stefan Institute
Slovenia

Bei Wang†
Scientific Computing and
Imaging Institute
University of Utah

Guoning Chen‡
University of Houston

Paul Rosen§
Scientific Computing and
Imaging Institute
University of Utah

ABSTRACT

Vector field simplification aims to reduce the complexity of the flow by removing features in order of their relevance and importance, to reveal prominent behavior and obtain a compact representation for interpretation. Most existing simplification techniques based on the topological skeleton successively remove pairs of critical points connected by separatrices, using distance or area-based relevance measures. These methods rely on the stable extraction of the topological skeleton, which can be difficult due to instability in numerical integration, especially when processing highly rotational flows. These geometric metrics do not consider the flow magnitude, an important physical property of the flow. In this paper, we propose a novel simplification scheme derived from the recently introduced topological notion of robustness, which provides a complementary view on flow structure compared to the traditional topological-skeleton-based approaches. Robustness enables the pruning of sets of critical points according to a quantitative measure of their stability, that is, the minimum amount of vector field perturbation required to remove them. This leads to a hierarchical simplification scheme that encodes flow magnitude in its perturbation metric. Our novel simplification algorithm is based on degree theory, has fewer boundary restrictions, and so can handle more general cases. Finally, we provide an implementation under the piecewise-linear setting and apply it to both synthetic and real-world datasets.

Keywords: Vector field data, topology-based techniques, flow visualization.

1 INTRODUCTION

Vector fields and their analysis are indispensable for many applications in science and engineering. With the increasing gap between the size and complexity of the vector field data from real-world applications and the limited bandwidth of our visual perception channel, it is more and more challenging for domain experts to interpret their data in detail or as a whole. This challenge is prominent in 2D turbulence flows, where features are everywhere and feature sizes differ by a few orders of magnitude. Vector field simplification aims at reducing the complexity of the flow by removing features in order of their relevance and importance, revealing prominent behavior, obtaining a compact representation for interpretation, and giving a consistent and multiscale view of the flow dynamics.

A considerable amount of research has been focused on vector field simplification based on the notion of a topological skeleton [14, 16]. A topological skeleton consists of critical points connected by special streamlines called *separatrices*, which provide a condensed representation of the flow by dividing the domain into regions of uniform flow behavior. However, existing simplification techniques rely on the stable extraction of the topological skeleton,

which can be difficult due to instability in numerical integration, especially when processing highly rotational flows, e.g. Figure 1. Furthermore, the distance and area-based relevance measures that are commonly used to determine the cancellation ordering of critical points typically rely on geometric proximities and do not consider the flow magnitude, an important physical property of the flow.

In this paper, we propose a new vector field simplification scheme derived from the recently introduced notion of *robustness*. Robustness, a notion related to *persistence* [10, 20], is used to represent the stability of critical points and assess their significance with respect to perturbations of the vector field. Intuitively, the robustness of a critical point is the minimum amount of perturbation, with respect to a metric encoding flow magnitude, that is required to cancel it within a local neighborhood. Our contributions are:

- We propose a new simplification strategy based on robustness, which is hierarchical and enables the pruning of sets of critical points according to a quantitative measure of their stability.
- Our method, compared to the distance-based metric, provides a complementary view to topological-skeleton-based simplification. It requires neither the topological skeleton nor heuristic parameters, but it does provide comparable simplification results for typical scenarios. Furthermore, the method can handle more general situations and the simplification can be computed efficiently based on sublevel sets for large, complex datasets when separatrices are difficult to integrate.
- Our strategy is built on a novel simplification algorithm based on degree theory. The algorithm can remove critical points in any connected subregion of the vector field whose degree is zero. It can handle more general boundary configurations without requirements on the Conley index. We provide an implementation under the piecewise-linear (PL) setting and apply it to a number of synthetic and real-world datasets.

We *do not* intend to show that the robustness-based method is necessarily better than topological-skeleton-based methods across all scenarios. In the typical situation involving pairs of critical




Figure 1: Topological skeleton: Sinks (and saddle-sink separatrices) are red, sources (and saddle-source separatrices) green, and saddles blue. (a) A highly rotational flow field where the pointed critical points are close to Hopf-bifurcations. Numerical inaccuracies may accumulate during integration and separatrices may intersect or switch. (b)-(c) Instability of separatrices under a small perturbation: The upper right sink is not connected with the saddle on the left in (b), but is after a small perturbation in (c).

*e-mail: primož.skraba@ijs.si

†e-mail: beiwang@sci.utah.edu

‡e-mail: chengu@cs.uh.edu

§e-mail: prosen@sci.utah.edu

points connected by separatrices, such methods offer comparable visual results. Rather, the robustness-based method can handle more general situations, is scalable, and gives a novel, mathematically rigorous hierarchical simplification scheme. Our method finds, in the space of all vector fields, the one that is closest to the original vector field with a particular set of critical points removed, according to a metric based on the L_∞ norm (the maximum point-wise modification to the vector field). Our results are optimal in this norm, that is, there exists no simplification with a smaller perturbation.

2 RELATED WORK

Vector field simplification can be classified into topology-based and non-topology-based techniques [32]. Non-topology-based techniques typically focus on Laplacian smoothing of the potential of a vector field [18, 25, 31]. Topology-based techniques modify the vector field topology explicitly by merging or cancelling nearby critical points based on the notion of a topological skeleton [5, 14, 16, 32]. De Leeuw and Van Liere [7, 8] made use of a geometry-based relevance measure (e.g., with respect to distance or area proximity) to determine the pair of critical points to be cancelled. Tricoche et al. [26] focused on a piecewise analytic description for the simplified field, which was later extended to time-dependent 2D flows [28]. Theisel et al. [24] presented a topology-preserved compression and simplification of vector fields. Zhang et al. [32] introduced a framework for fixed point pair cancellation based on Conley index theory. Chen et al. [5] extended this idea to include periodic orbits and presented a more complete pairwise cancellation framework. Recently, Chen et al. [4, 6] introduced a multiscale hierarchy of the vector field topology based on the Morse Connection Graph (MCG) computed from Morse decomposition [6]. This work was extended to address piecewise constant vector fields by Szymczak et al. [23, 22]. Such representations could be used to simplify vector fields by iteratively merging pairs of Morse sets that are adjacent in the MCG. The order of the pairs for cancellation depends only on the geometric characteristics of the Morse sets, i.e., the pairs that lead to smaller merged Morse sets will be cancelled or merged first. Weinkauff et al. [30] introduced a topological simplification technique for 3D vector fields based on the extraction of higher-order critical points. The simplification is assisted by a derived auxiliary 2D vector field on a closed surface surrounding each higher-order critical point.

Simplifications have also been proposed in a combinatorial setting [19, 21]. Edelsbrunner et al. [9, 10] performed pair cancellation on scalar fields defined on surfaces by changing the values of the scalar function near the fixed point pair. This is equivalent to simplifying the gradient vector field of the scalar function. Finally, scale space techniques [15, 20] have also been proposed to assess the importance of a critical point for topology-based simplification.

Robustness is closely related to the notion of persistence [10]. While persistence has been used successfully for scalar field visualization, robustness, first introduced in [11], is specifically designed for vector-valued data [3, 12]. Recent work [29] assigns robustness to critical points in both stationary and time-varying settings and obtains a structural description of the vector field. Such a structural description implies the existence of a hierarchical simplification strategy based on robustness, which is the focus of this paper.

In general, topology-based simplification techniques pair the topological features for cancellation via the computation of separatrices, which can be numerically unstable [6]. In contrast, the proposed robustness-based method does not require this computation and, thus, is insensitive to numerical error. The simplification hierarchy obtained from topology-based methods is typically invariant to scaling (multiplying the vector field with a scalar field), whereas our technique is sensitive to the change of vector field magnitude as it directly corresponds to our perturbation metric (Section




Figure 2: Figure adapted from [29]. Suppose the vector field is continuous, where sinks are red, sources are green, and saddles are blue. From left to right: vector fields f , relations among components of $\mathbb{F}_r$, and augmented merge trees. f contains four critical points, a sink x_1 , a source x_3 , and two saddles x_2 and x_4 . We use β , γ , ω , etc. to represent components of certain sublevel sets.

3). The robustness-based method achieves comparable results to the topology-based simplification and can handle more challenging cases in which the topology-based methods may fail (Section 5).

3 BACKGROUND

We provide relevant background in degree theory and robustness by reviewing previous work [3, 29] with minimal algebraic definitions and illustrating the related concepts through an example (Figure 2 adapted from [29]). We also provide introductory descriptions of isolating neighborhoods and Laplacian smoothing [5, 32].

Degrees. For a critical point x in 2D, its degree $\deg(x)$ equals its (Poincaré) index, that is, the number of field rotations while traveling along a closed curve centered at x counter-clockwise. Sources, sinks, centers, and saddles have indices $+1$, $+1$, $+1$ and -1 , respectively. Furthermore, for a (path-)connected component C that encloses several critical points, its degree $\deg(C)$ is the sum of the respective degrees of those critical points [3]. For our robustness-based simplification strategy, we rely on a corollary of the Poincaré-Hopf theorem (which is also employed by topological-skeleton-based simplification, e.g., [27]), which states that if a connected component C in 2D has degree zero, then it is possible to replace the vector field inside C with a vector field free of critical points.

Merge tree. To analyze a continuous 2D vector field $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, we define a corresponding scalar function (referred to as the flow magnitude function) $f_0: \mathbb{R}^2 \rightarrow \mathbb{R}$ which assigns for each point the magnitude (Euclidean norm) of the corresponding vector, $f_0(x) = \|f(x)\|_2$. We use $\mathbb{F}_r = f_0^{-1}(-\infty, r]$ to denote the *sublevel set* of f_0 for some $r \geq 0$. $\mathbb{F}_0$ is precisely the set of critical points of f .

Increasing r from 0, the space $\mathbb{F}_r$ evolves and we can construct a graph that tracks the (connected) components of $\mathbb{F}_r$ as they appear and merge. This is called a *merge tree* (or *join tree* as described in [1]). The root represents the entire domain of f_0 and the leaves represent the creation of a component at a local minimum. An internal node represents the merging of two or more components. We further record an integer at each node, which is the degree of the corresponding component in the sublevel set, and refer to the result as an augmented merge tree. An initial computation of the degrees of critical points is sufficient to determine the degree of any component of any sublevel set by computing the sum of the degrees of the critical points lying in it [3]. An example is shown in Figure 2¹. The merge tree on the right shows how the components of the sublevel sets $\mathbb{F}_r$ evolve. At $r = 0$ there are four components that correspond to the four critical points, each with nonzero degree. At $r = r_1$, components that contain x_1 and x_2 merge into a single component β_1 , which has zero degree. When $r = r_2$, components β_1 and β_2 merge into a single component γ_1 with degree $+1$, while β_3 grows into γ_2 . Finally at $r = r_3$, the single component ω_1 has zero degree.

Static robustness and its properties. The (static) robustness of a critical point is the height of its lowest degree zero ancestor in

¹We do not show any components that appear after $r = 0$ as they have zero degrees and do not correspond to critical points of the vector field.

the merge tree [2, 29]. The static robustness quantifies the stability of a critical point with respect to perturbations of the vector fields through the following lemmas explicitly stated in [29].

We first define the concept of *perturbation*. Let $f, h: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be two continuous 2D vector fields. Define the distance between the two mappings as $d(f, h) = \sup_{x \in \mathbb{R}^2} \|f(x) - h(x)\|_2$. A continuous mapping h is an r -*perturbation* of f , if $d(f, h) \leq r$.

Lemma 3.1 (Critical Point Cancellation [29]) Suppose a critical point x of f has robustness r . Let C be the connected component of $\mathbb{F}_{r+\delta}$ containing x , for an arbitrarily small $\delta > 0$. Then, there exists an $(r + \delta)$ -perturbation h of f , such that $h^{-1}(0) \cap C = \emptyset$ and $h = f$ except possibly within the interior of C .

Lemma 3.2 (Degree & Critical Point Preservation [29])

Suppose a critical point x of f has robustness r . Let C be the connected component of $\mathbb{F}_{r-\delta}$ containing x , for some $0 < \delta < r$. For any ε -perturbation h of f where $\varepsilon \leq r - \delta$, the sum of the degrees of the critical points in $h^{-1}(0) \cap C$ is $\deg(C)$. If C contains only one critical point x , we have $\deg(h^{-1}(0) \cap C) = \deg(x)$. That is, x is preserved as there is no ε -perturbation that could cancel it.

Revisiting the example in Figure 2, the robustness of the critical points x_1, x_2, x_3 , and x_4 is r_1, r_1, r_3 , and r_3 , respectively. Since the robustness of x_3 is r_3 , for any $\delta > 0$, we consider a component $C \subseteq \mathbb{F}_{r_3+\delta}$ that is slightly larger than ω_1 and contains x_3 (in fact, ω_1 contains all four critical points). Lemma 3.1 implies the existence of an $(r_3 + \delta)$ -perturbation that cancels x_3 by locally modifying the component C . Now consider another component $C' \subseteq \mathbb{F}_{r_3-\delta}$ where $r_2 < r_3 - \delta < r_3$, then C' has degree +1. Lemma 3.2 states that any $(r_3 - \delta)$ -perturbation preserves the degree of C' .

Isolating neighborhood and Laplacian smoothing. Previously, topology-based simplification has focused on cancelling pairs of critical points that are connected by separatrices. Zhang et al. [32] and Chen et al. [5] propose to compute an *isolating neighborhood* surrounding a pair of critical points, where a critical-point-free vector field can be found by solving a constrained optimization problem, referred to as a vector-valued *Laplacian smoothing* [32].

Based on Conley index theory, every boundary point of an isolating neighborhood can be classified as either an *entrance* or *exit* point. If an isolating region C in the domain contains multiple critical points and has a trivial Conley index, the flow inside C can be replaced with a new field free of critical points [32]. A typical situation for C to have a trivial Conley index is when its boundary ∂C consists of a single inflow and a single outflow component. As shown in the later examples, such an isolating neighborhood is not always easy to construct. The robustness-based method has no such constraint and only requires the degree of C to be zero.

4 ROBUSTNESS-BASED SIMPLIFICATION ALGORITHMS

In robustness-based simplification, we first locate sets of critical points that share the lowest zero-degree ancestors in the merge tree and sort them based on their robustness values. For each set with a common robustness r , we compute the corresponding component of the sublevel set $C \subseteq \mathbb{F}_r$. Since by construction $\deg(C) = 0$, our strategy can simplify C , whereas the distance-based strategy requires an isolating neighborhood with trivial Conley index.

4.1 Preliminary

First we introduce the relevant constructions in a smooth setting, and then translate the corresponding language into the PL setting.

Given a 2D vector field restricted to a degree-zero component C , $f: C \rightarrow \mathbb{R}^2$, we define the *image space* of C , $\text{im}(C)$. For each point $p \in C$, we have a vector $v_p = f(p) \in \mathbb{R}^2$. $\text{im}(C) \subset \mathbb{R}^2$ is constructed by mapping p to its vector coordinates v_p . The origin in $\text{im}(C)$ corresponds to the critical points (0 vectors) in C . Since $C \subseteq \mathbb{F}_r$, it follows that $\forall p \in C, \|v_p\|_2 \leq r$, therefore $\text{im}(C)$ is contained within a disc of radius r in $\mathbb{R}^2$. We denote the boundary of this disc by S .




Figure 3: (a)-(b): Illustrative examples for uncovered (a) and covered (b) boundaries of $\text{im}(C)$. (c): A component and its image space with a few mappings highlighted.

Now suppose the boundary of C , denoted as ∂C , is a simple closed curve². Note that the above maps ∂C to S , obtaining the image, $\text{im}(\partial C)$. We refer to the boundary of $\text{im}(C)$ as *uncovered*, if $\text{im}(\partial C) \subset S$, otherwise, as *covered*. Figures 3(a)-(b) illustrate these concepts. Note that both examples have zero degree. In 3(a), the region C encloses a saddle-sink pair connected by a separatrix. By traversing counter-clockwise along ∂C and observing how its image $\text{im}(\partial C)$ wraps around S , we see that the boundary of $\text{im}(C)$ is uncovered. In 3(b), the region C encloses a saddle-sink pair not connected by separatrix and the boundary of $\text{im}(C)$ is covered.

In the PL setting, the vector field f is restricted to a triangulation K of C , $f: K \rightarrow \mathbb{R}^2$, where the support of K , $|K| = C$. We construct the image of C by mapping each vertex $p \in K$ to its vector coordinates $v_p = f(p)$. Through linear interpolation, this construction also maps edges and triangles in K to edges and triangles in $\text{im}(C)$ (Figure 3(c)). The concept of *covered* and *uncovered* boundaries of $\text{im}(C)$ can be defined similarly up to a small additive constant.

4.2 Algorithm Overview

Our simplification strategy consists of four operations:

- **Smoothing**(C): Perform Laplacian smoothing on C ;
- **Cut**(C): Deform the vector field in its image space $\text{im}(C)$ to remove critical points in C ;
- **Unwrap**(C): Modify the vector field in its image space $\text{im}(C)$ so part of its boundary is uncovered;
- **Restore**(C): Set the boundary to its original value.

Three cases are classified by the Conley index of C , denoted as $\text{CH}_*(C)$. The operations to simplify each case are:

- If $\text{CH}_*(C)$ is trivial, return $C_1 = \text{Smoothing}(C)$.
- If $\text{CH}_*(C)$ is nontrivial and the boundary of $\text{im}(C)$ is uncovered, then $C_1 = \text{Cut}(C)$, and return $C_2 = \text{Smoothing}(C_1)$.
- If $\text{CH}_*(C)$ is nontrivial and the boundary of $\text{im}(C)$ is covered, then $C_1 = \text{Unwrap}(C)$, $C_2 = \text{Cut}(C_1)$, $C_3 = \text{Restore}(C_2)$ and return $C_4 = \text{Smoothing}(C_3)$.

By construction, $\deg(C) = 0$ in all three cases. Indeed, $\deg(C) \neq 0$ is a sufficient condition such that there exists no simplification.

4.3 Algorithm Details

We describe the **Cut** and **Unwrap** operations in detail and discuss the maximum amount of perturbation needed due to these operations. **Smoothing** is only used to achieve visually appealing results. **Cut operation.** Suppose the boundary of $\text{im}(C)$ is uncovered. The idea behind the **Cut** operation is to deform $\text{im}(C)$ such that there is a small neighborhood surrounding the origin that is not covered by $\text{im}(C)$. This corresponds to the situation where there is no critical point in C after the deformation. As shown in Figure 5(left), we

²This is not needed, but it simplifies the algorithm and exposition.




Figure 4: (a)-(b) Locating a cut point for the **Cut** operation: (a) Track the angle (a.k.a. phase) of a point in $\text{im}(\partial C)$ along S as we move along ∂C counter-clockwise. (b) The corresponding phase plot (a.k.a. angle-valued function) is shown in blue. The result of phase-unwrapping is shown in red. (c)-(e) Locating an unwrap point in **Unwrap** operation: (c) Track the angle of a point in $\text{im}(\partial C)$ along S as we move along ∂C counter-clockwise. (d) The corresponding angle-valued function (shown in blue), the result of phase-unwrapping (shown in red), and the optimal unwrap point c^* corresponding to phase ϕ^* . (e) The modified boundary of $\text{im}(C)$ (shown in purple), which becomes uncovered.




Figure 5: **Cut** operation. Left: The projection of edges that intersect ℓ during the **Cut** operation. Right: After **Cut**, the light blue region represents $\text{im}(C)$, which no longer contains (covers) the origin and so is critical point free.

choose a point c^* on the uncovered part of the circle S (this point is referred to as the *cut point*) and define the line ℓ as the line segment beginning at the origin O and terminating at c^* . Define another line ℓ' , which is orthogonal to ℓ and is ε away from the origin. The point $s \in \ell'$ is at a distance ε from the origin. Next, we find all the mesh edges $v_p v_q$ (corresponding to the edge $p q$ in K) in the interior of $\text{im}(C)$ that intersect with the line ℓ , and project their end points onto ℓ' , forming the projected edge $v'_p v'_q$. In the original domain, the vectors at $p, q \in K$ are deformed from v_p and v_q to the vectors v'_p and v'_q , respectively. Third, we locate all the mesh edges $v_x v_y$ where $x \in \partial C$ (and so v_x is on the boundary of $\text{im}(C)$ and v_y is in the interior). We move the point v_y to s so that the edge $v_x v_y$ no longer crosses ℓ and the boundary vector remains unchanged. Since the boundary of $\text{im}(C)$ is uncovered, there is no edge that intersects ℓ whose end points are both located on the boundary of $\text{im}(C)$ (i.e., whose corresponding points are both in ∂C). This operation creates an empty wedge around the origin (Figure 5 right), which ensures that there are no critical points in C after the modification. By construction, the amount of perturbation is less than $r + \varepsilon$. When doing Laplacian smoothing, we keep the projected vertices (end points of the cut edges) fixed to ensure that the origin is not recovered.

The procedure to find a cut point c^* is shown in Figure 4(a)-(b). In (a), by traversing counter-clockwise along ∂C and observing how its image $\text{im}(\partial C)$ (blue curve) wraps around S , we define the angle θ of a point along S to be its *phase*. In (b), we showcase (in blue) the corresponding phase plot (a.k.a. angle-valued function), that is, a function $h : \partial C \rightarrow \theta$ where $\theta \in [-\pi, \pi]$. Traversing ∂C again, we can use *phase-unwrapping* to give us a continuous function $\phi : \partial C \rightarrow \phi$ for $\phi \in \mathbb{R}$ (shown in red) using the following equation

$$\phi(i) = \lfloor \theta(i) - \phi(i-1) + 1/2 \rfloor + \theta(i).$$

Since the boundary of $\text{im}(C)$ is uncovered, it follows that $\max_{\partial C}(\phi) - \min_{\partial C}(\phi) < 2\pi$. We set the cutting angle ϕ^* as the mid-point of the uncovered part and the corresponding cut point is

$$\phi^* = \frac{1}{2} \left(\max_{\partial C}(\phi) + \min_{\partial C}(\phi) \right) + \pi, \quad c^* = (r \cos \phi^*, r \sin \phi^*),$$

where r is the robustness parameter of the sublevel set (and the radius of the disk S , where $\text{im}(C) \subset S$). By using the phase parameter θ , we do not need to worry about PL effects when computing c^* .

Unwrap operation. If the boundary of $\text{im}(C)$ is covered, we must first **Unwrap** the boundary before we perform the **Cut** procedure. The **Unwrap** operation is divided into the steps illustrated in Figure 4(c)-(e). Similarly to the *cut point*, we determine the optimal *unwrap point*. As before, we traverse ∂C and compute a phase plot $h : \partial C \rightarrow \theta$, unwrapping the phase to obtain a continuous function $\phi : \partial C \rightarrow \phi$ (Figure 4(c)-(d)). The unwrapping point ϕ^* , is

$$\phi^* = \frac{1}{2} \left(\max_{\partial C}(\phi) + \min_{\partial C}(\phi) + 2n\pi \right),$$

where n is the smallest integer such that $|\min(\theta) + 2n\pi - \max(\theta)| < \pi$, and $c^* = (r \cos \phi^*, r \sin \phi^*)$. To **Unwrap** the boundary, let $X \in \partial C$ be the set of points on the boundary such that $\phi(X) > \phi^* - \delta$, and $Y \in \partial C$ be the set of points that $\phi(Y) < \phi^* + \delta - 2n\pi$. As illustrated in Figure 4(e), to **Unwrap** we set

$$\begin{aligned} \phi(X) &= \phi^* - \delta, & \phi(Y) &= \phi^* - 2n\pi + \delta, \\ v_x &= \begin{pmatrix} r \cos(\phi(x)) \\ r \sin(\phi(x)) \end{pmatrix} & x \in X, & v_y &= \begin{pmatrix} r \cos(\phi(y)) \\ r \sin(\phi(y)) \end{pmatrix} & y \in Y, \end{aligned}$$

where r is the magnitude of the vectors on the boundary (e.g., the sublevel set parameter). The final step is to **Restore** the boundary to its original values. As in case (b), the deformation is bounded by $r + \varepsilon$. We omit the proof, but the key observation is that internal nodes move less than $r + \varepsilon$ whereas the boundary has the original values.




Figure 6: SyntheticA. (a) The original vector field: sinks are red, sources are green and saddles are blue. (b) The topological skeleton: saddle-sink separatrices are red and saddle-source separatrices are green. (c)-(d) 1st level simplification: before (c) and after (d) **Smoothing**. (e)-(f) 2nd level simplification: before (e) and after (f) **Smoothing**.

4.4 Synthetic Examples

We illustrate our robustness-based simplification strategy on three PL synthetic examples, highlighting the three different cases.

SyntheticA (Figure 6) corresponds to the example in Figure 2. It involves pairs of critical points connected by separatrices. At r_1 , we have a component that contains critical points x_1 and x_2 and at




Figure 7: SyntheticB. (a) the original vector field with its topological skeleton. (a)-(b): Single level simplification before (a) and after (b) by **Cut** and **Smoothing**. (c) Only applying **Smoothing** does not make the region a critical point free field.




Figure 8: SyntheticC. (a) the original vector field with topological skeleton. (b)-(c) Before (b) and after (c) simplification by combining **Unwrap**, **Cut** and **Smoothing**.

r_3 we have a component that contains all four critical points x_1 to x_4 . The simplification hierarchy involves two steps ranked by robustness values: first x_1 and x_2 are simplified, and then x_3 and x_4 . Since both components (marked by yellow boundary) have a trivial Conley index, this corresponds to case (a), where only **Smoothing** operations are needed. SyntheticB (Figure 7) involves a group of four critical points that are interconnected by separatrices, which could be simplified in a single level using a robustness-based strategy. Since the component of interest has a nontrivial Conley index, directly applying Laplacian smoothing fails (as shown in Figure 7(c)). The component's boundary is uncovered, so we apply case (b) of our simplification by combining **Cut** with **Smoothing**.

SyntheticC (Figure 8) corresponds to case (c) of our algorithm. This is an untypical case involving a pair of critical points not directly connected by a separatrix. In this case, the component of interest C has nontrivial Conley index, and the boundary of its image is covered. The robustness-based strategy cancels the critical point pair without any issue by combining **Unwrap**, **Cut** and **Smoothing** operations. We further focus on this example by illustrating the image space of C , $\text{im}(C)$, during various steps of simplification in Figure 9. In Figure 9(a), the entire boundary and disk are covered. However, from the left phase plot in Figure 10, we can see that the degree is 0. Once the optimal unwrapping point is computed, we perform the **Unwrap** operation, giving the right phase plot in Figure 10 and the image space in Figure 9(b), leaving the boundary S uncovered. The effect of the **Cut** operation in image space is shown in Figure 9(c), creating a void surrounding the origin. Lastly, in Figure 9(d), the boundary is restored for the final output.




Figure 9: SyntheticC. The image space is shown through the different steps: (a) original, (b) after **Unwrap**, (c) after **Cut**, and (d) final output after **Restore**.

5 RESULTS

We demonstrate our robustness-based simplification strategy on a number of real-world datasets. When possible, we compare our method with distance-based simplification. The first real-world dataset we explore is the top layer of a 3D simulation of global oceanic eddies [17] for 350 days of the year 2002. The 2D time-varying vector field has resolution 3600×2400 . We extract tiles




Figure 10: SyntheticC. Left: The phase plot, original version (blue), and the phase-unwrapped version (red). Right: The phase plot with optimal unwrap point (orange) and the modified phase plot with boundary uncovered (purple).

representing the flow in the central Atlantic Ocean (60×60) and construct standard triangulation on the point samples. We select multiple time slices from this data: OceanA contains slices #21217 and #21311; OceanB and OceanC correspond to slices #20904 and #20821, respectively; OceanD includes a time-varying sequence of slices from #20710 to #20715. Our second real-world dataset is a 2D time-varying vector field simulation of homogeneous charge compression ignition (HCCI) engine combustion [13] represented as a 640×640 regular grid with a periodic boundary. The data consists of 299 time-steps at intervals of 10^{-5} seconds. We selected slice #173 from this data, referred to as the Combustion dataset.

5.1 Topologically Equivalent Scenarios

In many scenarios, our approach produces topologically equivalent and visually comparable results to the distance-based approach, such as for the OceanB dataset (Figure 11(a)). The critical point pairs of interest are highlighted by the black dashed boxes in the top row left. The critical points are colored by their robustness values (red—low, blue—high). The upper right pair is more robust than the lower middle pair and is further apart. The simplification results generated by distance-based and the robustness-based approaches are shown in the second and third rows, respectively. The approximated isolating neighborhoods are highlighted by the white boxes (middle row), whereas the sublevel sets the yellow enclosure (bottom row). From the comparison, we observe that, first, both the distance and robustness metrics generate the same pairs of critical points, and second, the simplification orderings determined by these two metrics agree. A subtle difference in the resulting vector fields is visible due to the different local regions determined by the two metrics and different algorithms for modifications.

OceanA dataset (Figure 12 (a)-(b)) shows a more complex scenario where the region encloses more than two critical points. The vector fields in this example are from slices #21217 and #21311. Each of these two clusters (highlighted by the black dashed boxes in Figure 12(a)-(b)) consists of four critical points that are close in distance and have small identical robustness values. The robustness metric groups them as one cluster automatically and computes a region based on their sublevel set. The bottom row of Figure 12(a)-(b) provides the simplification results using the algorithm introduced in Section 4. Although the distance-based method cannot group these four critical points in one simplification, for comparison we compute an isolating neighborhood that encloses them and apply Laplacian smoothing. Both methods are shown to return similar results. Nevertheless, the robustness-based method can handle regions with more complex boundary configurations.

5.2 Inconsistent Hierarchical Scenarios

We also identified a number of scenarios where the distance-based and robustness-based methods disagree. One example is the OceanC dataset (Figure 11(b)). Here, two pairs of critical points are studied (highlighted in the top row of Figure 11(b)). Even though the pairing of these four critical points is consistent with both metrics, their actual simplification orderings are different. The distance-based method cancels the pair in the middle-right of the domain first, while the robustness-based method cancels the lower-middle pair first. Figure 13 provides another example that shows




Figure 11: (a) The OceanB dataset. (b) The OceanC dataset. For each subfigure: Top Row: Left – shows robustness values with the region of interest highlighted (robustness values are colored from red to white, where red means low and white means high robustness); Right – shows the vector field marked by critical point types along with separatrices. Middle Row: the two-step hierarchical simplification based on distance. Bottom Row: the two-step hierarchical simplification based on robustness.




Figure 12: The OceanA dataset: (a) #20311; (b)-(c) #21217. (a)-(b) For each subfigure, Top Row: Left – shows robustness values with region of interest highlighted; Right – shows the vector field marked by critical point types and its topological skeleton. Bottom Row: results after distance-based (left) and robustness-based simplifications (right). (c) A region (yellow boundary) with a nontrivial Conley index and uncovered boundary (top), where smoothing does not remove its critical points (bottom).

the discrepancy of the two approaches in determining the simplification ordering of critical point pairs in the time-varying setting. In this example, we look at consecutive time steps from the OceanD dataset. Figure 13(a) highlights the critical points of interest. The pairings of these four critical points again agree with each other using both topological-skeleton and robustness metrics. We perform a per-slice simplification using the two approaches. The results are shown in the second (distance-based) and third (robustness-based) columns in (b)-(c), respectively. From the results, we see that the cancellation orderings change over time using the distance-based metric. This is due to an increased distance between the two critical points near the upper-right corner, resulting in a change of the simplification order. On the other hand, the robustness for these

two pairs is stable. Therefore, the robustness-based simplification returns a consistent outcome in this example.

5.3 Challenging Scenarios

There are a number of cases where the topological-skeleton-based metric combined with the Laplacian smoothing technique is incapable of simplifying the given vector field. For example, for the SyntheticB dataset shown in Figure 7, it is impossible to find an isolating neighborhood with a trivial Conley index that encloses all the critical points due to the boundary condition. Therefore, even though the obtained local region is guaranteed to be zero degree, Laplacian smoothing fails to solve for a critical point free field. On the other hand, the simplification algorithm introduced in Sec-




Figure 13: The OceanD dataset. (a) A sampled time series with pairs of critical points highlighted, where white numbers indicate time stamps. (b) #21710. (c) #21715. For each subfigure (b)-(c), Top Row: The original vector field (left) and with the separatrices (right). Middle Row: The simplification ordering for the distance-based strategy. Bottom Row: The simplification ordering for the robustness-based strategy. Orderings for distance and robustness-based methods are consistent in (b) and different in (c).




Figure 14: The Combustion dataset. The bottom-up hierarchical simplifications (Top) from the distance-based strategy and (Bottom) from the robustness-based strategy.

tion 4 successfully simplifies the field. A similar situation occurs in Figure 12(c) (OceanA slice #21217). In this example, we try to apply Laplace smoothing in the local region computed based on robustness (top). The boundary configuration of this region is rather complex and does not satisfy a trivial Conley index. The Laplacian smoothing based on this boundary configuration fails (bottom), but the proposed simplification method succeeds. These two examples showcase the utility of the proposed algorithm in solving a critical point free field within any given regions with zero degree. This relieves the requirement of the trivial Conley index whose corresponding isolating neighborhood is sometimes difficult to obtain.

Figure 8 shows a nontypical case that involves the cancellation of a pair of critical points not directly connected by separatrix. It is impossible for the topological-skeleton-based method to compute an isolating neighborhood that encloses two critical points (but not the others) not connected by separatrix [5]. Nonetheless, the ro-

bustness metric derives a local region that encloses only these two critical points with total degree equal to zero under a certain configuration of the flow magnitude. Hence, these two critical points may be cancelled. Whereas this may rarely occur in the real-world data, it illustrates the flexibility and generality of the proposed method. In practice, a simpler but similar situation may occur.

In the slice of the combustion data (Figure 14), the simplification results and hierarchies of the distance-based metric (first row), and the robustness-based metric (second row) do not agree. The corresponding vector field is a high resolution incompressible flow. Conventional topological-skeleton-based methods are potentially difficult to apply, as the separatrices either do not exist or require many integration steps before reaching the sink/source-like critical points, making pairing the critical points challenging. A standard solution in 2D flow is to compute the skeleton of the vector field rotated by 90 degrees [32]. Whereas this approach may work well for 2D

fields, it is not straightforward in the 3D case. In addition, computing a topological-skeleton for the dual vector field and deriving the subsequent isolating neighborhood for a given pair is computationally expensive. In contrast, the robustness-based method does not require the computation of topology, and its computation is fast and parallelizable, making it more practical for large datasets.

6 DISCUSSIONS

We have presented a new and complementary simplification framework that does not depend on the topological skeleton but incorporates topological information through *robustness*. Rather than considering the geometric proximity of critical points, we consider the minimum perturbation required to remove critical points. Our algorithm comes with theoretical guarantees on the amount of perturbation we introduce. The motivation for Laplacian smoothing is to produce more visually appealing results. However, to the best of our knowledge, no nontrivial bounds exist on the amount of perturbation introduced by such a smoothing. In practice, smoothing only marginally increases the amount of perturbation³.

Scalability: Our method should scale to very large datasets. The robustness computation and the simplification steps (e.g., **Cut** and **Unwrap**) run in linear time in the size of the mesh. For example, for a region of 21k vertices and 64k edges, **Cut** required 2 seconds in MATLAB and 0.03 seconds in C++.

Generality: The simplification procedure requires only that the degree of the boundary be zero and so applies to a wide range of cases. It can deal with highly rotational data (e.g., centers) as well as cases where critical points are not connected by separatrices.

Other metrics: We use robustness and the L_∞ norm (the maximum over the domain), but using other metrics such as the L_2 -norm, which incorporates both the magnitude of the vectors and the area to capture a quantity closer to the energy of a perturbation, would be interesting. The simplification requires only degree-zero components and any metric could be used to construct a hierarchy. It is an open question to find degree-zero regions under different metrics.

Time-varying and 3D vector fields: The main challenge in simplifying time-varying 2D vector fields is to achieve consistency across time-slices, e.g., obtaining critical points correspondences at a given simplification level. Finally, the prospect of extending our framework to 3D vector fields is promising. Whereas there remain technical obstacles, certain operations (such as cutting and smoothing) in our pipeline readily extend to higher dimensions. This, however, is beyond the scope of this paper but will be addressed in future work.

ACKNOWLEDGEMENTS

We thank Jackie Chen for the combustion dataset and Mathew Maltese from LANL and the BER Office of Science UV-CDAT team for the ocean datasets. PR was supported by DOE NETL and KAUST award KUS-C1-016-04. PS was supported by TOPOSYS (FP7-ICT-318493). GC was supported by NSF IIS-1352722. BW was supported in part by INL 00115847 DE-AC07051D14517 and DOE NETL.

REFERENCES

- [1] H. Carr, J. Snoeyink, and U. Axen. Computing contour trees in all dimensions. *ACM-SIAM SODA*, pages 918–926, 2000.
- [2] F. Chazal, A. Patel, and P. Skraba. Computing the robustness of roots. Manuscript, http://ailab.ijs.si/primoz_skraba/papers/fp.pdf, 2011.
- [3] F. Chazal, A. Patel, and P. Skraba. Computing well diagrams for vector fields on R^n . *Applied Math. Letters*, 25(11):1725–1728, 2012.
- [4] G. Chen, Q. Deng, A. Szymczak, R. Laramée, and E. Zhang. Morse set classification and hierarchical refinement using Conley index. *IEEE TVCG*, 18(5):767–782, 2012.

³See the supplementary material for the detailed perturbation measurements and performance numbers.

- [5] G. Chen, K. Mischaikow, R. Laramée, P. Pilarczyk, and E. Zhang. Vector field editing and periodic orbit extraction using Morse decomposition. *IEEE TVCG*, 13(4):769–785, 2007.
- [6] G. Chen, K. Mischaikow, R. Laramée, and E. Zhang. Efficient Morse decompositions of vector fields. *IEEE TVCG*, 14(4):848–862, 2008.
- [7] W. de Leeuw and R. van Liere. Collapsing flow topology using area metrics. In *IEEE Vis*, pages 349–354, 1999.
- [8] W. de Leeuw and R. van Liere. Multi-level topology for flow visualization. *Comp. & Graph.*, 24(3):325–331, 2000.
- [9] H. Edelsbrunner, J. Harer, and A. Zomorodian. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete Comp. Geom.*, 30:87–107, 2003.
- [10] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comp. Geom.*, 28:511–533, 2002.
- [11] H. Edelsbrunner, D. Morozov, and A. Patel. The stability of the apparent contour of an orientable 2-manifold. In *Topo. M. in Data Anal. and Vis.*, pages 27–41. Springer, 2010.
- [12] H. Edelsbrunner, D. Morozov, and A. Patel. Quantifying transversality by measuring the robustness of intersections. *Found. of Comp. Math.*, 11:345–361, 2011.
- [13] E. Hawkes, R. Sankaran, P. Pébay, and J. Chen. Direct numerical simulation of ignition front propagation in a constant volume with temperature inhomogeneities. *Combustion and Flame*, 145:145–159, 2006.
- [14] J. Helman and L. Hesselink. Representation and display of vector field topology in fluid flow data sets. *IEEE Computer*, 22(8):27–36, 1989.
- [15] T. Klein and T. Ertl. Scale-space tracking of critical points in 3D vector fields. *Topo. Meth. in Vis.*, pages 35–49, 2007.
- [16] R. Laramée, H. Hauser, L. Zhao, and F. Post. Topology based flow visualization: state of the art. *Topo. Meth. in Vis.*, pages 1–19, 2007.
- [17] M. Maltud, F. Bryan, and S. Peacock. Boundary impulse response functions in a century-long eddying global ocean simulation. *Environmental Fluid Mechanics*, 10:275–295, 2010.
- [18] K. Polthier and E. Preuß. Identifying vector fields singularities using a discrete hodge decomposition. *Vis. and Math. III*, pages 112–134, 2003.
- [19] J. Reininghaus, J. Kasten, T. Weinkauff, and I. Hotz. Efficient computation of combinatorial feature flow fields. *IEEE TVCG*, 2011.
- [20] J. Reininghaus, N. Kotava, D. Guenther, J. Kasten, H. Hagen, and I. Hotz. A scale space based persistence measure for critical points in 2d scalar fields. *IEEE TVCG*, 17(12):2045–2052, 2011.
- [21] J. Reininghaus, C. Lowen, and I. Hotz. Fast combinatorial vector field topology. *IEEE TVCG*, 17(10):1433–1443, 2011.
- [22] L. Sipeki and A. Szymczak. Simplification of Morse decompositions using morse set mergers. In *Topo-In-Vis 2013*, 2013.
- [23] A. Szymczak and E. Zhang. Robust morse decompositions of piecewise constant vector fields. *IEEE TVCG*, 18(6):938–951, 2012.
- [24] H. Theisel, C. Rössl, and H.-P. Seidel. Combining topological simplification and topology preserving compression for 2D vector fields. In *Pacific Graphics*, pages 419–423, 2003.
- [25] Y. Tong, S. Lombeyda, A. Hirani, and M. Desbrun. Discrete multi-scale vector field decomposition. *ACM ToG*, 22:445–452, 2003.
- [26] X. Tricoche, G. Scheuermann, and H. Hagen. A topology simplification method for 2D vector fields. In *IEEE Vis*, pages 359–366, 2000.
- [27] X. Tricoche, G. Scheuermann, and H. Hagen. Continuous topology simplification of planar vector fields. *IEEE Vis*, pages 159–166, 2001.
- [28] X. Tricoche, G. Scheuermann, and H. Hagen. Topological visualization of time-dependent 2D vector fields. *EuroVis*, pages 117–126, 2001.
- [29] B. Wang, P. Rosen, P. Skraba, H. Bhatia, and V. Pascucci. Visualizing robustness of critical points for 2D time-varying vector fields. *CGF*, 32(3):221–230, 2013.
- [30] T. Weinkauff, H. Theisel, K. Shi, H.-C. Hege, and H.-P. Seidel. Extracting higher order critical points and topological simplification of 3d vector fields. In *IEEE Vis*, pages 559–566, 2005.
- [31] R. Westermann, C. Johnson, and T. Ertl. A level-set method for flow visualization. *IEEE Vis*, pages 147–154, 2000.
- [32] E. Zhang, K. Mischaikow, and G. Turk. Vector field design on surfaces. *ACM ToG*, 25:1294–1326, 2006.