

On Vietoris–Rips complexes: the persistent homology of cyclic graphs

Ethan Coldren

Abstract

We give an $O(n^2(k+\log n))$ algorithm for computing the k -dimensional persistent homology of a filtration of clique complexes of circular arc graphs on n vertices. This is nearly quadratic in the number of vertices n , and therefore a large improvement upon the traditional persistent homology algorithm, which is cubic in the number of simplices of dimension at most $k+1$, and hence of running time $O(n^{3(k+2)})$ in the number of vertices n . Our algorithm applies, for example, to Vietoris–Rips complexes of points sampled from a curve in \mathbb{R}^d , so long as the scale parameter is sufficiently small. In the case of the plane \mathbb{R}^2 , we prove that our algorithm applies for all scale parameters if the n vertices are sampled from a convex closed differentiable curve whose convex hull contains its evolute. We ask if there are other geometric settings in which computing persistent homology is (say) quadratic or cubic in the number of vertices, instead of in the number of simplices.

1 Introduction

Given only a finite sample from a metric space, what properties of the space can one recover from the finite sample? Vietoris–Rips complexes are a commonly used tool in applied topology in order to recover the homotopy type, homology groups, or persistent homology of a space from a finite sample [10, 11, 15, 16, 18, 19, 20, 21, 25]. Given a metric space X and a scale parameter $r \geq 0$, the *Vietoris–Rips simplicial complex* $\text{VR}(X; r)$ has a simplex every finite subset of X of diameter at most r .

It can in general be expensive to compute the homotopy type or persistent homology of a Vietoris–Rips complex. Indeed, let $n = |X|$ be the number of points in a finite metric space X . Computing the k -dimensional persistent homology of $\text{VR}(X; r)$ as r increases is cubic in the number of simplices of dimension at most $k+1$, and hence of running time $O\left(\binom{n}{k+2}^3\right) = O(n^{3(k+2)})$ in the number of vertices n .¹ In this paper we show that if X is sampled from a curve in \mathbb{R}^d , and if the scale parameter is small enough such that the underlying

¹For example, computing the 3-dimensional persistent homology of a Vietoris–Rips complex of n points is cubic in the number of simplices, but of order $O(n^{15})$ in the number of vertices n .

1-skeleton is a *cyclic graph*, then the k -dimensional homology of $\text{VR}(X; r)$ (in that range of scales) can be computed in running time $O(n^2(k + \log n))$, which is nearly quadratic in the number of vertices n . A *cyclic graph* is a combinatorial abstraction of the 1-skeleton of a Vietoris–Rips complex built on a subset of the circle; a precise definition is given in Section 3.

For some curves, the restriction that the scale r be sufficiently small is not needed. Indeed, if X is sampled from a convex closed differentiable curve in the plane whose convex hull contains its evolute, then the k -dimensional homology of $\text{VR}(X; r)$ can be computed (for all r) in running time $O(n^2(k + \log n))$, which is nearly quadratic in the number of vertices n . The *evolute* of a curve is the envelope of the normals, or equivalently, the locus of the centers of curvature. We hope this is a first step towards identifying more general geometric settings in which computing persistent homology is (say) quadratic or cubic in the number of vertices, instead of in the number of simplices.

Our main results follow.

Theorem 1. *Let $G_1 \subseteq G_2 \subseteq \dots \subseteq G_M$ be an increasing sequence of cyclic graphs on a final vertex set of size n . Then the k -dimensional persistent homology of the resulting increasing sequence of clique complexes $\text{Cl}(G_1) \subseteq \text{Cl}(G_2) \subseteq \dots \subseteq \text{Cl}(G_M)$ can be computed in running time $O(n^2(k + \log n))$.*

Theorem 2. *Given a strictly convex closed differentiable planar curve C equipped with the Euclidean metric, the 1-skeleton of $\text{VR}(C; r)$ is a cyclic graph for all $r \geq 0$ if and only if the convex hull of C contains the evolute of C .*

Corollary 1. *There is an $O(n^2(k + \log n))$ algorithm for determining the k -dimensional persistent homology of $\text{VR}(X; r)$, where X is a sample of n points from a strictly convex closed differentiable planar curve whose convex hull contains its evolute.*

Corollary 1 follows from the prior two theorems since, as we explain in A, given a sample X of n points from some strictly convex closed differentiable planar curve C whose convex hull contains its evolute, it is easy to determine the cyclic graph structure on the 1-skeleton of $\text{VR}(X; r)$ even *without* knowledge of C .

Corollary 2. *There is an $O(n \log n)$ algorithm for determining the homotopy type of $\text{VR}(X; r)$, where X is a sample of n points from a strictly convex closed differentiable planar curve whose convex hull contains its evolute, and where $r \geq 0$.*

We emphasize that if X is planar, the homotopy type of $\text{VR}(X; r)$ in Corollary 2 and Corollary 1 can be surprising (an odd sphere S^{2k+1} for any $k \geq 0$, or a wedge of even spheres $\bigvee^m S^{2k}$ for any $m \geq 0$ and $k \geq 0$); see Figure 1.

Though Theorem 2, Corollary 1, and Corollary 2 are all about points in \mathbb{R}^2 , we remark that Theorem 1 is not. Indeed, if C is a differentiable curve in \mathbb{R}^d , then there will be some constant $r < r_C$ such that for any finite subset $X \subseteq C$ of size n , the k -dimensional persistent homology of $\text{VR}(X; r)$ in the range $r < r_C$




Figure 1: A possible evolution of the homotopy types of $\text{VR}(X; r)$, for X a subset of a convex differentiable planar curve containing its evolute, as the scale r increases (horizontal axis). The homotopy type is either a single odd sphere S^{2k+1} , or a wedge sum of even spheres $\bigvee^m S^{2k}$ for some $m \geq 0$. The vertical axis gives a cartoon of how m might vary with r : the number m of $2k$ -spheres is a non-decreasing function of r for $2k = 0$, but otherwise m need not be a monotonic function of r for $2k \geq 2$.

can be computed in time $O(n^2(k + \log n))$. In \mathbb{R}^2 it is possible to give precise characterizations of r_C , to describe those curves C for which $r_C = \infty$.

We would like to emphasize that the goal of this paper is not to recover a curve from a finite sample, which is a well-studied problem [13]. Instead, our goal is to better understand the computational complexity of the homology and homotopy types of Vietoris–Rips complexes. Vietoris–Rips complexes are designed to recover not only curves but also arbitrary homotopy types. When given a finite subset $X \subseteq \mathbb{R}^d$ for d small, in order to understand the “shape” of X one would often compute its Čech or *alpha complex* [25] instead of computing its Vietoris–Rips complex. Indeed, by the nerve lemma the Čech and alpha complexes will have milder homotopy types. However, in higher-dimensional Euclidean space, it becomes prohibitively expensive to compute a Čech or alpha complex, and hence Vietoris–Rips complexes are frequently used [13, 25]. The theory of Vietoris–Rips complexes, though more subtle than that for the aforementioned Čech and alpha complexes, is important since Vietoris–Rips complexes are computable in \mathbb{R}^d for d large whereas Čech and alpha complexes are not.

Our work motivates the following question: Are there other geometric contexts besides the plane where computing the persistent homology of the Vietoris–Rips complex of a sample of n points can be similarly improved, from cubic in the number of simplices to a low-degree polynomial in n ?

Question 1. For $X \subseteq \mathbb{R}^2$ arbitrary, is there a cubic or near-quadratic algorithm in the number of vertices $n = |X|$ for determining the k -dimensional persistent homology of $\text{VR}(X; r)$?

Question 2. For $X \subseteq \mathbb{R}^3$ arbitrary, what is the computational complexity of

computing the k -dimensional persistent homology of $\text{VR}(X; r)$ in terms of the number of vertices $n = |X|$?

Question 3. For C a sufficiently nice curve in \mathbb{R}^d , what is the computational complexity of computing the k -dimensional persistent homology of a subset of n points from C ?

See the conclusion section for some initial ideas towards Question 3.

We remark that it is NP-hard to compute the homology of the clique complex of an arbitrary clique complex; see Theorem 7 of [7]. Since every finite graph can be realized as the unit ball graph of a collection of points in \mathbb{R}^d for d sufficiently large, it follows that computing the homology (or persistent homology) of Vietoris–Rips complexes in an Euclidean space is NP-hard. However, to our knowledge this NP-hardness result may not hold in restricted low dimensions, such as \mathbb{R}^2 or \mathbb{R}^3 .

Another motivating question behind this work is the following. Given a planar subset $X \subseteq \mathbb{R}^2$, is the Vietoris–Rips complex $\text{VR}(X; r)$ necessarily homotopy equivalent to a wedge of spheres for all $r \geq 0$? See Problem 7.3 of [6] and Section 2, Question 5 of [29]. Some evidence towards this conjecture is contained in [18] and [6]. Our results show that the conjecture is true in the limited case where X is a subset of a strictly convex closed differentiable curve whose convex hull contains its evolute.

2 Related work

Vietoris–Rips complexes were invented independently by Vietoris for use in algebraic topology [45], and by Rips for use in geometric group theory [31]. Indeed, Rips proved that if a group G equipped with the word metric is δ -hyperbolic, then $\text{VR}(G; r)$ is contractible for $r \geq 4\delta$. An important theorem by Hausmann [33] states that if M is a Riemannian manifold, then $\text{VR}(M; r)$ is homotopy equivalent to M for scale parameters r sufficiently small (depending on the curvature of M). This theorem has been extended by Latschev [35] to state that if X is a (possibly finite) metric space that is sufficiently close to M in the Gromov–Hausdorff distance, then $\text{VR}(X; r)$ is still homotopy equivalent to M .

Hausmann’s and Latschev’s theorems form the theoretical basis for more recent applications of Vietoris–Rips complexes in applied and computational topology [25, 15, 16]. There are by now a wide variety of reconstruction guarantees— one can use Vietoris–Rips complexes to recover a wide variety of topological properties, such as homotopy type, homology, or fundamental group, from a finite subset drawn from some unknown underlying shape [10, 11, 15, 16, 18, 19, 20, 21, 25]. Several algorithms exist in order for approximating the persistent homology of a Vietoris–Rips complex filtration in a more computationally efficient manner [14, 22, 23, 39, 44].

This paper relies upon and builds upon cyclic graphs, and on the known homotopy types of the Vietoris–Rips complex of the circle [1, 2, 3, 4, 5, 8, 43].

Section 5 of our paper can be viewed as a generalization of [5] from ellipses to a much broader class of planar curves.

3 Preliminaries

We set notation for topological and metric spaces, simplicial complexes, persistent homology, cyclic graphs, convex curves, and evolutes. See [9, 32, 34] for background on topological spaces, simplicial complexes, and homology, and [25] for background on Vietoris–Rips complexes and persistent homology.

Topological spaces

A *topological space* is a set X equipped with a collection of subsets of X , called open sets, such that any union of open sets is open, any finite intersection of open sets is open, and both X and the empty set are open. For X a topological space and $Y \subseteq X$ a subset, we denote the *interior* of Y by $\text{int}(Y)$ and the *boundary* of Y by ∂Y . Let $I = [0, 1]$ denote the unit interval. We let S^k denote the k -dimensional sphere and $\bigvee^m S^k$ denote the m -fold wedge sum of S^k with itself. We write $X \simeq Y$ to denote that spaces X and Y are *homotopy equivalent*, which roughly speaking means that “they have the same shape up to bending and stretching”.

Metric spaces

A *metric space* is a set X equipped with a distance function $d: X \times X \rightarrow \mathbb{R}$ satisfying certain properties: nonnegativity, symmetry, the triangle inequality, and the identity of indiscernibles ($d(x, x') = 0$ if and only if $x = x'$). Given a point $x \in X$ and a radius $r > 0$, we let $B_X(x, r) = \{y \in X \mid d(x, y) < r\}$ denote the open ball with center x and radius r . Given a metric space X , a point $x \in X$, and a set $Y \subseteq X$, we define $d(x, Y) = \inf\{d(x, y) \mid y \in Y\}$.

Simplicial complexes

A simplex is a generalization of the notion of a vertex, edge, triangle, or tetrahedron to arbitrary dimensions. Formally, given $k + 1$ points x_0, x_1, \dots, x_k in general position, a simplex of dimension k (a k -simplex) is the smallest convex set containing them. A simplicial complex K on a vertex set X is a collection of subsets (simplices) of X , including each element of X as a singleton, such that if $\sigma \in K$ is a simplex and $\tau \subseteq \sigma$ is a face of σ , then also $\tau \in K$. We do not distinguish between abstract simplicial complexes (which are combinatorial) and their geometric realizations (which are topological spaces).

Vietoris–Rips complexes

A *Vietoris–Rips complex* $\text{VR}(X; r)$ is a simplicial complex, defined from a metric space X and distance $r \geq 0$, by including as a simplex every finite set of points in

X that has a diameter at most r [33]. Said differently, the vertex set of $\text{VR}(X; r)$ is X , and $\{x_0, x_1, \dots, x_k\}$ is a simplex when $d(x_i, x_j) \leq r$ for all $0 \leq i, j \leq k$.




Figure 2: Vietoris–Rips complexes, on the same vertex set X at two different choices of scale r .

Homology and persistent homology

Given a topological space Y and an integer $k \geq 0$, the homology group $H_k(Y)$ measures the independent “ k -dimensional holes” in Y (roughly speaking). For example, $H_0(Y)$ measures the number of connected components, $H_1(Y)$ measures the loops, and $H_2(Y)$ measures the “2-dimensional voids” in Y .

Given an increasing sequence of spaces $Y_1 \subseteq Y_2 \subseteq \dots \subseteq Y_M$, persistent homology is a way to “track the holes” as the spaces get larger. A common choice for applications is $Y_i = \text{VR}(X; r_i)$, for X a metric space and for $r_1 < r_2 < \dots < r_M$ an increasing sequence of scale parameters. We apply the homology functor (with coefficients in a field) in order to get a sequence of vector spaces $H_k(Y_1) \rightarrow H_k(Y_2) \rightarrow \dots \rightarrow H_k(Y_M)$, which decomposes into a collection of 1-dimensional interval summands. Each interval corresponds to a k -dimensional topological feature that is born and dies at the start and endpoints of the interval. Our algorithm for persistent homology in Section 4 works simultaneously for all choices of field coefficients. It also works for integer coefficients (which are much more subtle [40]), or for persistent homotopy [27, 36], since all of the spaces that appear in our context are homotopy equivalent to wedges of spheres, with controllable maps in-between.

Cyclic graphs and clique complexes

A *directed graph* $G = (X, E)$ consists of a set of vertices X and edges $E \subseteq X \times X$, where no loops, multiple edges, or edges oriented in opposite directions are allowed. A *cyclic graph* [2, 5] is a directed graph in which the vertex set is equipped with a counterclockwise cyclic order, such that whenever we have three cyclically ordered vertices $x \prec y \prec z \prec x$ and a directed edge $x \rightarrow z$, then we also have the directed edges $x \rightarrow y$ and $y \rightarrow z$. We say a cyclic graph G is *finite* if its vertex set is finite, and otherwise G is *infinite*. For example, in the proof of Theorem 2, when $r > 0$ we will show that the 1-skeleton of $\text{VR}(C; r)$

is an infinite cyclic graph (recall C is a curve). We give more preliminaries on cyclic graphs in Section 4, for use in the proof of Theorem 1.




Figure 3: Four example cyclic graphs. The homotopy types of their clique complexes, from left to right, are S^1 , S^2 , $\vee^2 S^2$, and S^3 (since, as will be explained in Section 4, the *winding fractions* are $\frac{1}{4}$, $\frac{1}{3}$, $\frac{1}{3}$, and $\frac{3}{8}$, with the graphs having 1, 2, 3, and 1 *periodic orbit(s)*).

The *clique complex* of a (directed or undirected) graph G (with no loops or multiple edges), denoted $\text{Cl}(G)$, is the largest simplicial complex that contains G as its 1-skeleton. Note, for example, that a Vietoris–Rips complex is the clique complex of its 1-skeleton.

Strictly convex curves

A set $Y \subseteq \mathbb{R}^2$ is *strictly convex* if for all $y, y' \in Y$ and $t \in (0, 1)$, the point $ty + (1 - t)y'$ is in the interior of Y . We say a curve $C \subseteq \mathbb{R}^2$ is *strictly convex* if $C = \partial Y$ for some strictly convex set $Y \subseteq \mathbb{R}^2$. If L is a line and C is strictly convex curve in \mathbb{R}^2 that intersect transversely, then L and C intersect in either zero or two points; see for example [12, 24, 30, 42].

Tangent vectors, normal vectors, and evolutes

Let $\alpha: I \rightarrow \mathbb{R}^2$ be a differentiable curve in the plane. Then $\alpha'(t)$ is the tangent vector to α at time t , and we denote the corresponding *unit tangent vector* by $T(\alpha(t)) = \frac{\alpha'(t)}{\|\alpha'(t)\|}$. The *unit normal vector*, which is perpendicular to $T(\alpha(t))$ and points in the direction the curve is turning, is given by $n(\alpha(t)) = \frac{T'(\alpha(t))}{\|T'(\alpha(t))\|}$.

The *curvature* of α at time t is $\kappa(t) = \frac{\|T'(t)\|}{\|\alpha'(t)\|}$, with corresponding *radius of curvature* $r(t) = \frac{1}{\kappa(t)}$. The *center of curvature* is the point on the inner normal line to $\alpha(t)$ at distance equal to the radius of curvature away, given by $x_\alpha(t) = \alpha(t) + \frac{1}{\kappa(t)}n(\alpha(t))$. The *evolute* of a curve is the envelope of the normals, or equivalently, the set of all centers of curvature.

Critical Points, Extrema, and Monotonicity

Let $f: Y \rightarrow \mathbb{R}$ be a differentiable real-valued function. A point $x \in \text{int}(Y)$ is a *critical point* of f if $f'(x) = 0$. The function f has an *absolute maximum*

at a point x if $f(x) \geq f(y)$ for all $y \in Y$, and an *absolute minimum* at x if $f(x) \leq f(y)$ for all $y \in Y$. The function f has a *local maximum* at a point $x \in Y$ if $f(x) \geq f(y)$ for all y in some open set containing x , and similarly for a *local minimum*. For $Y \subseteq \mathbb{R}$, a function $f: Y \rightarrow \mathbb{R}$ is said to be *monotonic* if we have $f(x) \leq f(y)$ for all $x, y \in Y$ with $x \leq y$, or alternatively, $f(x) \geq f(y)$ for all $x \leq y$.

4 Proof of Theorem 1 on an algorithm for persistent homology

Theorem 1 states that given any increasing sequence of cyclic graphs $G_1 \subseteq G_2 \subseteq \dots \subseteq G_M$, we can compute the k -dimensional persistent homology of the resulting increasing sequence of clique complexes $\text{Cl}(G_1) \subseteq \text{Cl}(G_2) \subseteq \dots \subseteq \text{Cl}(G_M)$ in running time $O(n^2(k + \log n))$. Before proving Theorem 1, we provide some further background on cyclic graphs and their associated dynamical systems. We then give the algorithm for computing even-dimensional persistent homology, followed by the algorithm for odd homological dimensions.

4.1 Cyclic dynamical systems and winding fractions

Let G be a finite cyclic graph with vertex set X . The associated *cyclic dynamical system* is generated by the dynamics $f: X \rightarrow X$, where we assign $f(x)$ to be the vertex y with a directed edge $x \rightarrow y$ that is counterclockwise furthest from x (else $f(x) = x$ if x is not the source vertex of any directed edges in G). Since X is finite, the dynamical system $f: X \rightarrow X$ necessarily has at least one periodic orbit. By Lemma 2.3 of [5] or Lemma 3.4 of [8], every periodic orbit of f has the same length ℓ and winding number ω (the number of times a periodic orbit $x \rightarrow f(x) \rightarrow f^2(x) \rightarrow \dots \rightarrow f^\ell(x) = x$ wraps around the cyclic ordering on X). We define the *winding fraction* of G to be $\text{wf}(G) = \frac{\omega}{\ell}$.




Figure 4: Two example cyclic graphs and their associated dynamical systems $f: X \rightarrow X$. (Left) This cyclic graph G has a dynamical system with 3 periodic orbits, each of winding fraction $\frac{1}{3}$, giving $\text{Cl}(G) \simeq \bigvee^2 S^2$. (Right) This cyclic graph G has a dynamical system with a single periodic orbit of winding fraction $\frac{3}{8}$, giving $\text{Cl}(G) \simeq S^3$ since $\frac{1}{3} < \frac{3}{8} < \frac{2}{5}$.

Let P be the number of periodic orbits in a finite cyclic graph G . By Proposition 4.1 of [5], we have

$$\text{Cl}(G) \simeq \begin{cases} S^{2k+1} & \text{if } \frac{k}{2k+1} < \text{wf}(G) < \frac{k+1}{2k+3} \\ \bigvee^{P-1} S^{2k} & \text{if } \text{wf}(G) = \frac{k}{2k+1} \end{cases} \text{ for some } k \in \mathbb{N}.$$

Furthermore, as described in Section 4.6, given an inclusion of cyclic graphs $G \hookrightarrow G'$, we have strong results determining the topology of the map $\text{Cl}(G) \hookrightarrow \text{Cl}(G')$.

4.2 Even-dimensional homology

Let G be a finite cyclic graph with vertex set X of size n . Let E be a list of all directed edges, in sorted order from first to last, such that the subgraph of G formed by including any initial segment of edges in E is also a cyclic graph. We show how to compute the persistent homology of the clique complexes of the subgraphs of G formed by starting with vertex set X , and then adding edges one at a time in sorted order. To compute $2k$ -dimensional homology, we must count the number of periodic orbits with winding fraction $\frac{k}{2k+1}$. Let B be an array of n booleans storing whether each vertex is periodic or not (true means periodic, false means nonperiodic). Initialize each entry of B to be false (unless $k = 0$, in which case every vertex is periodic with winding fraction 0 prior to adding any edges). Let P be the number of vertices that are periodic with winding fraction $\frac{k}{2k+1}$, initialized to be 0 (again unless $k = 0$). Let f be an array of n integers, where $f[i] = j$ when the cyclic dynamical system maps the i -th vertex to the j -th vertex. Initialize f so that $f[i] = i$ for all i . For each new edge, in order of appearance, we must update whether some vertices are periodic with winding fraction $\frac{k}{2k+1}$. If the source vertex of the new edge is periodic (look it up in B), walk along f , marking every vertex along this (old) periodic orbit as nonperiodic in B . Decrement P . Now, update f to add the new edge. Walk $2k + 1$ steps along f starting from the source vertex of the new edge. If we get back to where we started, after looping k times around, then re-walk along this newly found periodic orbit, marking all the vertices as periodic in B , and increment P . We're now done updating f , B , and P for the new edge. Then, $P - 1$ is the number of $2k$ -dimensional spheres in the homotopy type of $\text{Cl}(G)$. Thus, we now know what the homology groups are. See Section 4.6 for an explanation of how to recover not only the homology at each stage, but also the persistent homology (i.e., the maps between homology groups induced by inclusions).

We now explain why this algorithm works. The algorithm has a loop invariant that f , B , and P are correct. When we add a new larger step $x \mapsto f(x)$ to the cyclic dynamics f , we are implicitly removing a prior smaller step from x . This prior step was either part of a periodic orbit or not. If it was on a periodic orbit, we update f , B , and P to account for destroying this periodic orbit. When we add the new step $x \mapsto f(x)$, it either creates a new periodic orbit or not. We check to see if it does, and, if so, we update f , P , and B accordingly. Thus, the loop invariant is maintained throughout the execution of

the algorithm. Since the $2k$ -dimensional homology of the clique complex of G is determined purely by the number of periodic orbits of winding fraction $\frac{k}{2k+1}$, this algorithm produces the correct homology at each stage.

4.3 Pseudocode for even-dimensional homology

The following pseudocode for the even-dimensional persistent homology algorithm described above accepts as inputs the vertex set X , the sorted list E of directed edges, and k . It computes the $2k$ -dimensional persistent homology of the increasing sequence of clique complexes. We let $n = |X|$.

```
function computeEvenDimensionalPersistentHomology(X, E, k):
  set numPeriodicOrbits to 0 (unless k=0)
  set isPeriodic to an array of length n, filled with 0's (unless k=0)
  set f to an array of length n, where the i-th entry is i
  set edges to a sorted list of all edges between points in X
  for edge in edges:
    if isPeriodic[edge.sourceVertex]:
      walk along the periodic orbit, marking each vertex nonperiodic
      numPeriodicOrbits -= 1
    edit f to add the new edge
    walk around f, starting at edge.sourceVertex, taking 2k+1 steps
    if we returned to edge.sourceVertex, and looped around k times:
      walk along the new periodic orbit, marking each vertex periodic
      numPeriodicOrbits += 1
  print the number (numPeriodicOrbits-1) of 2k-spheres
```

See Section 4.6 for an explanation of how to recover not only the homology at each stage, but also the persistent homology.

4.4 Odd-dimensional homology

If we're interested in $2k + 1$ dimensional homology, run the algorithm for even-dimensional homology in dimensions $2k$ and $2k + 2$. There is only ever at most one $2k + 1$ -dimensional sphere. This sphere is born when the last periodic orbit with winding fraction $\frac{k}{2k+1}$ is destroyed (at which point the winding fraction first exceeds $\frac{k}{2k+1}$), and this sphere dies when the first periodic orbit with winding fraction $\frac{k+1}{2k+3}$ is created.

4.5 Computational complexity

Computing the list of edges in sorted order takes $O(n^2 \log(n^2)) = O(n^2 \log n)$ time. Walking along the length of a periodic orbit takes $O(k)$ time. We walk along the length of a periodic orbit a constant number of times for each edge. Since there are $O(n^2)$ edges, this takes $O(n^2 k)$ time. Thus, the total runtime is $O(n^2(k + \log n))$.

4.6 Determining the persistent homology maps

In Section 4.2-4.4 we provide an algorithm to compute the homology groups (and indeed the homotopy types) of an increasing sequence of clique complexes of cyclic graphs $\text{Cl}(G_1) \subseteq \text{Cl}(G_2) \subseteq \dots \subseteq \text{Cl}(G_M)$. From these computations, it is not difficult to determine the persistent homology information, i.e., the maps on homology induced by inclusions. Indeed, if $\text{Cl}(G_i)$ and $\text{Cl}(G_{i+1})$ are not (wedge sums of) spheres of the same dimension, then necessarily the induced map on homology is trivial in all homological dimensions. Furthermore, if $\text{Cl}(G_i)$ and $\text{Cl}(G_{i+1})$ are each homotopy equivalent to S^{2k+1} , then it follows from Proposition 4.9 of [2] that the induced map $\text{Cl}(G_i) \hookrightarrow \text{Cl}(G_{i+1})$ is a homotopy equivalence.

For the case of even-dimensional homology, we rely on Proposition 4.2 of [5]. Each interval in the $2k$ -dimensional persistent homology barcode will be labeled with a periodic orbit of winding fraction $\frac{k}{2k+1}$. Upon adding a new directed edge with source vertex x , there are four cases in the algorithm for even-dimensional homology in Section 4.2.

1. If x was previously nonperiodic but became periodic, then unless this is the very first periodic orbit of winding fraction $\frac{k}{2k+1}$ to appear², begin a new persistent homology interval in homological dimension $2k$. Label this interval with the periodic orbit containing x .
2. If x was previously nonperiodic and remains nonperiodic, then no updates are needed. All persistent homology intervals continue.
3. If x was previously periodic and remains periodic (necessarily with a new periodic orbit), then simply update the periodic orbit label of this persistent homology interval. All persistent homology intervals continue.
4. If x was in a periodic orbit \mathcal{O} and becomes nonperiodic, then under repeated applications of f , the vertex x now maps into a different periodic orbit \mathcal{O}' . End the persistent homology interval corresponding to exactly one of the periodic orbits \mathcal{O} or \mathcal{O}' according to the elder rule: end the persistent homology interval that was born more recently, and (re)label the persistent homology interval that continues with the periodic orbit \mathcal{O}' . All other persistent homology intervals continue unchanged.

4.7 Adding vertices

If the sequence of cyclic graphs do not all have the same vertex set, we can use a small modification to the algorithm above to still compute the winding fractions and periodic orbits. We break the sequence of cyclic graphs into a sequence of operations, either adding a new edge, or adding a new vertex (and all of the corresponding edges for the new vertex which may be required to maintain cyclicity). We already know how to add a new edge. When we add

²In homological dimension 0 we are using reduced homology.

a new vertex, all we must do is add the outgoing edges, which can be done in $O(n)$ time by walking around the convex hull. Furthermore, we must add any new outgoing edges that end in the new vertex, which can also be done in $O(n)$ time by walking around the convex hull.

5 Proof of Theorem 2 on cyclic graphs and evolutes

Throughout this section, let C be a strictly convex closed differentiable planar curve equipped with the Euclidean metric. We provide an example before giving the proof of Theorem 2.

Example 1. Let $C = \{(a \cos(t), b \sin(t)) \mid t \in \mathbb{R}\} \subseteq \mathbb{R}^2$ be an ellipse, where we assume $a > b > 0$. The evolute of C is given by the points $\left\{ \left(\frac{a^2 - b^2}{a} \cos^3 t, \frac{b^2 - a^2}{b} \sin^3 t \right) \right\} \subseteq \mathbb{R}^2$. One can show that the convex hull of C contains its evolute if and only if $\frac{a}{b} \leq \sqrt{2}$. Indeed, for one direction, note that the point $(0, \frac{b^2 - a^2}{b})$ on the evolute (corresponding to $t = \frac{\pi}{2}$) is contained in the convex hull of C if and only if $-b \leq \frac{b^2 - a^2}{b}$, meaning $a^2 \leq 2b^2$, or equivalently $\frac{a}{b} \leq \sqrt{2}$. Hence the convex hull of C contains its evolute if and only if C is an “ellipse of small eccentricity” [5], showing that our work generalizes [5] to a broader class of curves.




Figure 5: Evolutes of ellipses $C = \{(a \cos(t), b \sin(t)) \mid t \in \mathbb{R}\} \subseteq \mathbb{R}^2$. On the left we have $\frac{a}{b} < \sqrt{2}$, and the convex hull of C contains the evolute of C . On the right we have $\frac{a}{b} > \sqrt{2}$.

Let C be a strictly convex closed differentiable curve in the plane.

Definition 1. Define the continuous function $h: C \rightarrow C$ which maps a point $p \in C$ to the unique point in the intersection of the normal line to C at p with $C \setminus \{p\}$.

Lemma 1. The map $h: C \rightarrow C$ is of degree one, which implies that h is surjective.

Proof. Suppose that C is parametrized by $\alpha: S^1 \rightarrow C$. Write $h(\alpha(e^{2\pi it})) = \alpha(e^{2\pi i(t+f(e^{2\pi it})))}$ where $f: S^1 \rightarrow [0, 2\pi)$ is a continuous function, which is possible since $h \circ \alpha$ and α never intersect. Then we can do a “straight-line” homotopy from f down to the zero function to show that $h \circ \alpha$ and α are homotopy equivalent. Indeed, consider the homotopy $H: S^1 \times I \rightarrow C$ defined by $H(e^{2\pi it}, s) = \alpha(e^{2\pi i(t+sf(e^{2\pi it})))}$, in which $H(\cdot, 0) = \alpha$ and $H(\cdot, 1) = h \circ \alpha$. Since these maps are homotopy equivalent, it follows³ that the winding number of $h \circ \alpha$ is equal to the winding number of α , which is 1. Hence h is surjective. \square

For $p \in C$, define the function $d_p: C \rightarrow \mathbb{R}$ by $d_p(q) = d(p, q)$, where $d(p, q)$ is the Euclidean distance between p, q .

Lemma 2. *Let $p \in C$. Then a point $q \in C$ is a critical point of the function $d_p: C \rightarrow \mathbb{R}$ if and only if $q = p$ or $h(q) = p$.*

Proof. It is clear that p is a global minimum of d_p , and therefore we may restrict attention to $q \neq p$. Consider an arbitrary point $p = (p_1, p_2) \in C \subseteq \mathbb{R}^2$. Let $\alpha: I \rightarrow C$ be a parametrized curve in C . Note that $d_p(\alpha(t)) = \sqrt{(p_1 - \alpha_1(t))^2 + (p_2 - \alpha_2(t))^2}$, for all $t \in I$. This gives us

$$\frac{d}{dt}d_p(\alpha(t)) = \frac{-\alpha'_1(t)(p_1 - \alpha_1(t)) - \alpha'_2(t)(p_2 - \alpha_2(t))}{\sqrt{(p_1 - \alpha_1(t))^2 + (p_2 - \alpha_2(t))^2}} = \frac{-\alpha'(t) \cdot (p - \alpha(t))}{\|p - \alpha(t)\|}. \quad (1)$$

Therefore, the only critical points of d_p away from p are when $\frac{d}{dt}d_p(\alpha(t)) = 0$, i.e. $\alpha'(t) \cdot (p - \alpha(t)) = 0$. These are precisely the points where the tangent line to $\alpha(t)$ is perpendicular to the line between $\alpha(t)$ and p . Therefore, the critical points of d_p occur at p and all $q \in C$ such that $h(q) = p$. \square

We are now ready to prove Theorem 2.

Proof of Theorem 2. Let C be a strictly convex closed differentiable planar curve C equipped with the Euclidean metric. We must show that the 1-skeleton of $\text{VR}(C; r)$ is a cyclic graph for all $r \geq 0$ if and only if the convex hull of C contains the evolute of C . We will show that the following are equivalent.

- (i) The evolute of C is contained in the convex hull $\text{conv}(C)$.
- (ii) Function $h: C \rightarrow C$ is injective.
- (iii) For all points $p \in C$, the distance function d_p has exactly two critical points.
- (iv) For all points $p \in C$, the distance function d_p is monotonic along the two arcs in C from p to the global maximum of d_p .
- (v) The intersections $B_{\mathbb{R}^2}(p, r) \cap C$ are connected for all $p \in C$ and $r > 0$.

³ A more general statement, which follows from the same proof, is that if two maps $\alpha, \tilde{\alpha}: S^1 \rightarrow S^1$ satisfy $\alpha(p) \neq \tilde{\alpha}(p)$ for all $p \in S^1$, then the winding numbers of α and $\tilde{\alpha}$ are equal.

(vi) The 1-skeleton of $\text{VR}(C; r)$ is a cyclic graph for all $r \geq 0$.

(i) \Leftrightarrow (ii). The intuition (before the proof) is as follows. For an example, see Figure 5. If $\alpha: I \rightarrow C$ is a curve moving in the counterclockwise direction, then note that $h(\alpha(t))$ will also be moving in the counterclockwise direction at $t \in I$ if and only if the center of curvature $x_\alpha(t)$ is in $\text{conv}(C)$. For the proof, we first note that a continuous map $h: C \rightarrow C$ of degree one is injective if and only if for any continuous function $\alpha: I \rightarrow C$, the orientations of $\alpha(t)$ and $h(\alpha(t))$ match for all times t . The result then follows from Lemma 6, which says that $\alpha(t)$ and $h(\alpha(t))$ have matching orientations for all t if and only if the evolute of C is contained in the convex hull $\text{conv}(C)$.

(ii) \Leftrightarrow (iii). Note that h is injective if and only if for each $p \in C$, there is a unique point $q \in C$ with $h(q) = p$, which by Lemma 2 is equivalent to d_p having exactly two critical points (a global minimum p , and the point q satisfying $h(q) = p$ as a global maximum).

(iii) \Leftrightarrow (iv). For (iii) \Rightarrow (iv), note that for all $p \in C$, compactness implies that d_p has at least two extrema (a global minima at p , and a global maxima). If each d_p has exactly two critical points, then each d_p must have exactly two extrema (as every extrema is a critical point). Therefore (iv) is satisfied.

We remark that the proof of (iv) \Rightarrow (iii) is subtle (see Figure 6), but we proceed regardless. Suppose that d_p has more than two critical points for some $p \in C$; we must find some point $\tilde{p} \in C$ such that $d_{\tilde{p}}$ does not satisfy the monotonicity property in (iv). The assumption on p means that there exists some critical point q of d_p that is neither the global minimum nor maximum of d_p . Since q is a critical point we can find a curve $\alpha: (-\delta, \delta) \rightarrow C$ with $\alpha(0) = q$ and $\frac{d}{dt}d_p(\alpha(t)) = 0$ for $t = 0$.




Figure 6: For a *single* point p , it may be that d_p has more critical points than extrema. Indeed, see the half-circle-half-ellipse example for C above, with the specific point p as pictured. We have three critical points and only two extrema of d_p (note that q is a critical point of d_p that is not an extremum). Nevertheless, (iv) \Rightarrow (iii) still holds since for *other* points $\tilde{p} \in C$, we have more than two extrema of $d_{\tilde{p}}$.

We claim that there is a point $\tilde{p} \in C$ arbitrarily close to p such that

- (a) $\frac{d}{dt}d_{\tilde{p}}(\alpha(t)) < 0$ for $t = 0$, and
- (b) the global maximum of $d_{\tilde{p}}$ is arbitrarily close to the global maximum of d_p .

We first show (a). By (1), the sign of this derivative depends only on whether the angle between $\alpha'(0)$ and $\tilde{p}-q$ is larger or smaller than $\frac{\pi}{2}$ radians. Since the angle between $\alpha'(0)$ and $p-q$ is exactly equal to $\frac{\pi}{2}$ radians, there is some satisfactory point \tilde{p} arbitrarily close to p . We next show that we can additionally satisfy (b). Let $p^+ \in C$ be the unique global maximum of $d_p: C \rightarrow \mathbb{R}$; the value of this global maximum is $d_p(p^+)$. Since d_p is a continuous function on a compact domain with a unique global maximum, given any $\delta > 0$, there exists some $\epsilon > 0$ such that every point $x \in C$ with $d_p(x) \geq d_p(p^+) - \epsilon$ satisfies $\|x - p^+\| < \delta$. Now, choose \tilde{p} sufficiently close to p so that the sup norm of the difference between the functions d_p and $d_{\tilde{p}}$ is less than $\frac{\epsilon}{2}$. It then follows that any global maximum of $d_{\tilde{p}}$ is within δ of p^+ .

Note that \tilde{p} is the global minimum of \tilde{p} , and that as we wrap around C in one direction towards the global maximum of $d_{\tilde{p}}$, we have found a point $\alpha(0)$ such that $\frac{d}{dt}d_{\tilde{p}}(\alpha(t)) < 0$ for $t = 0$. It follows that $d_{\tilde{p}}$ does not satisfy the monotonicity property in (iv), and hence we have completed the proof of (iii) \Leftrightarrow (iv).

(iv) \Leftrightarrow (v). Given a fixed point $p \in C$, note that $B(p, r) \cap C$ is connected for all $r > 0$ if and only if d_p satisfies the monotonicity property in (iv).

(v) \Leftrightarrow (vi) The final equivalence follows from Definition 4.10 and Lemma 4.11 of [8]. \square

We can now also prove Corollary 2.

Proof of Corollary 2. We provide an $O(n \log n)$ algorithm for determining the homotopy type of $\text{VR}(X; r)$, where X is a sample of n points from a strictly convex differentiable planar curve C whose convex hull contains its evolute, and where $r \geq 0$. By Theorem 2 we know that the 1-skeleton of $\text{VR}(C; r)$ is a cyclic graph, and it follows that the same is true of $\text{VR}(X; r)$ for any $X \subseteq C$. We use the $O(n \log n)$ algorithm in A to determine the cyclic ordering of the points in X (along C) from their coordinates in \mathbb{R}^2 ; this algorithm does not require knowledge of C . The result then follows since given a cyclic graph G with n vertices, there exists an $O(n \log n)$ algorithm for determining the homotopy type of the clique complex of G . Indeed, this is contained in Theorem 5.7 and Corollary 5.9 of [3], in which the algorithm is stated for a Vietoris–Rips complex of points on the circle, though it holds more generally for the clique complex of any finite cyclic graph. The algorithm proceeds by removing *dominated vertices*, without changing the homotopy type of the clique complex, until arising at a minimal regular configuration (in which each vertex has the same number of outgoing neighbors, such as the two cyclic graphs in the middle of Figure 3). One can read off the homotopy type from this regular configuration. \square

Remark 1. *The algorithm in Corollary 2 is furthermore of linear running time $O(n)$ if the vertices are provided in cyclic order.*

6 Conclusion

We have shown that the convex hull of a strictly convex closed differentiable planar curve C contains the evolute of C if and only if the 1-skeleton of $\text{VR}(C; r)$ is a cyclic graph for all $r \geq 0$. As a consequence, if X is any finite set of n points from C , then the homotopy type of $\text{VR}(X; r)$ can be computed in time $O(n \log n)$. Furthermore, we give an $O(n^2(k + \log n))$ time algorithm for computing the k -dimensional persistent homology of $\text{VR}(X; r)$ (or more generally, the persistent homology of an increasing sequence of cyclic graphs). This is significantly faster than the traditional persistent homology algorithm, which is cubic in the number of simplices of dimension at most $k + 1$, and hence of running time $O(n^{3(k+2)})$ in the number of vertices n .

Important questions motivated by our work are as follows.

- For $X \subseteq \mathbb{R}^2$, is there a cubic or near-quadratic algorithm in the number of vertices $n = |X|$ for determining the k -dimensional persistent homology of $\text{VR}(X; r)$? This would be more likely if the conjecture that the Vietoris–Rips complex $\text{VR}(X; r)$ of any planar subset $X \subseteq \mathbb{R}^2$ is homotopy equivalent to a wedge sum of spheres were true (see Problem 7.3 of [6] and Section 2, Question 5 of [29]).
- As for one dimension higher, what is the computational complexity of computing the k -dimensional persistent homology of $\text{VR}(X; r)$ in terms of the number of vertices $n = |X|$ when $X \subseteq \mathbb{R}^3$?
- Given a set of points $X \subseteq \mathbb{R}^2$ that all lie on their convex hull (e.g. $\text{conv}(X) = X$), and such that the distance from a point $x \in X$ to each of its neighbors is monotonically increasing and then monotonically decreasing as one wraps around the convex hull in a cyclically ordered fashion, does there exist a curve that passes through all of the points that does not intersect its evolute? Our algorithm does not require that we know what this curve is, but only whether it exists.

We end with some initial thoughts on Question 3. Let C be a closed curve in \mathbb{R}^d . If the intersections $B_{\mathbb{R}^d}(p, r) \cap C$ are connected for all $p \in C$ and $r > 0$, then the 1-skeleton of $\text{VR}(C; r)$ will be cyclic for all r . The ideas in Section 4 would then provide fast algorithms for computing the homotopy type and persistent homology of $\text{VR}(X; r)$ whenever X is a finite subset of C . Furthermore, even if the 1-skeleton of $\text{VR}(C; r)$ is not cyclic for all $r \geq 0$, it will be cyclic for sufficiently small $r \leq r_C$, where r_C is some constant depending on C . This will allow the persistent homology algorithm in Section 4 to be applied in the regime when the scale $r \leq r_C$ is small.

7 Acknowledgements

We would like to thank Sean Willmot for his joint work on this project, and Bei Wang for asking about the computational complexity of the persistent homology

of points from the circle.

References

- [1] Michał Adamaszek. Clique complexes and graph powers. *Israel Journal of Mathematics*, 196(1):295–319, 2013.
- [2] Michał Adamaszek and Henry Adams. The Vietoris–Rips complexes of a circle. *Pacific Journal of Mathematics*, 290:1–40, 2017.
- [3] Michał Adamaszek, Henry Adams, Florian Frick, Chris Peterson, and Corrine Previte-Johnson. Nerve complexes of circular arcs. *Discrete & Computational Geometry*, 56:251–273, 2016.
- [4] Michał Adamaszek, Henry Adams, and Francis Motta. Random cyclic dynamical systems. *Advances in Applied Mathematics*, 83:1–23, 2017.
- [5] Michał Adamaszek, Henry Adams, and Samadwara Reddy. On Vietoris–Rips complexes of ellipses. *Journal of Topology and Analysis*, pages 1–30, 2017.
- [6] Michał Adamaszek, Florian Frick, and Adrien Vakili. On homotopy types of Euclidean Rips complexes. *Discrete & Computational Geometry*, 58(3):526–542, 2017.
- [7] Michał Adamaszek and Juraj Stacho. Algorithmic complexity of finding cross-cycles in flag complexes. In *Proceedings of the Twenty-Eighth Annual Symposium on Computational Geometry*, pages 51–60. ACM, 2012.
- [8] Henry Adams, Samir Chowdhury, Adam Jaffe, and Bonginkosi Sibanda. Vietoris–Rips complexes of regular polygons. Preprint, arXiv:1807.10971, 2018.
- [9] Mark A Armstrong. *Basic Topology*. Springer, 2013.
- [10] Dominique Attali and André Lieutier. Geometry driven collapses for converting a Čech complex into a triangulation of a shape. *Discrete & Computational Geometry*, 2014.
- [11] Dominique Attali, André Lieutier, and David Salinas. Vietoris–Rips complexes also provide topologically correct reconstructions of sampled shapes. *Computational Geometry*, 46(4):448–465, 2013.
- [12] Christian Bär. *Elementary differential geometry*. Cambridge University Press, 2010.
- [13] Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and topological inference*, volume 57. Cambridge University Press, 2018.

- [14] Magnus Bakke Botnan and Gard Spreemann. Approximating persistent homology in Euclidean space through collapses. *Applicable Algebra in Engineering, Communication and Computing*, 26(1-2):73–101, 2015.
- [15] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.
- [16] Gunnar Carlsson, Tigran Ishkhanov, Vin de Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76:1–12, 2008.
- [17] Marco Castelpietra and Ludovic Rifford. Regularity properties of the distance functions to conjugate and cut loci for viscosity solutions of hamilton-jacobi equations and applications in riemannian geometry. *ESAIM: Control, Optimisation and Calculus of Variations*, 16(3):695–718, 2010.
- [18] Erin W. Chambers, Vin de Silva, Jeff Erickson, and Robert Ghrist. Vietoris–Rips complexes of planar point sets. *Discrete & Computational Geometry*, 44(1):75–90, 2010.
- [19] Frédéric Chazal, David Cohen-Steiner, Leonidas J Guibas, Facundo Mémoli, and Steve Y Oudot. Gromov–Hausdorff stable signatures for shapes using persistence. In *Computer Graphics Forum*, volume 28, pages 1393–1403, 2009.
- [20] Frédéric Chazal, Vin de Silva, and Steve Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, pages 1–22, 2013.
- [21] Frédéric Chazal and Steve Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proceedings of the 24th Annual Symposium on Computational Geometry*, pages 232–241. ACM, 2008.
- [22] Tamal K Dey, Fengtao Fan, and Yusu Wang. Computing topological persistence for simplicial maps. In *Proceedings of the thirtieth annual symposium on Computational geometry*, page 345. ACM, 2014.
- [23] Tamal K Dey, Dayu Shi, and Yusu Wang. Simba: An efficient tool for approximating Rips-filtration persistence via simplicial batch-collapse. *arXiv preprint arXiv:1609.07517*, 2016.
- [24] Manfredo P Do Carmo. *Differential Geometry of Curves and Surfaces: Revised and Updated Second Edition*. Courier Dover Publications, 2016.
- [25] Herbert Edelsbrunner and John L Harer. *Computational Topology: An Introduction*. American Mathematical Society, Providence, 2010.
- [26] Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.
- [27] Patrizio Frosini, Claudia Landi, and Facundo Mémoli. The persistent homotopy type distance. *arXiv preprint arXiv:1702.07893*, 2017.

- [28] Joseph Howland Guthrie Fu. Tubular neighborhoods in Euclidean spaces. *Duke Mathematical Journal*, 52(4):1025–1046, 1985.
- [29] William Gasarch, Brittany Terese Fasy, and Bei Wang. Open problems in computational topology. *ACM SIGACT News*, 48(3):32–36, 2017.
- [30] Vyacheslav L Girko. *Spectral Theory of Random Matrices*, volume 10. Academic Press, 2016.
- [31] Mikhael Gromov. Hyperbolic groups. In Stephen M Gersten, editor, *Essays in Group Theory*. Springer, 1987.
- [32] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
- [33] Jean-Claude Hausmann. On the Vietoris–Rips complexes and a cohomology theory for metric spaces. *Annals of Mathematics Studies*, 138:175–188, 1995.
- [34] Dmitry N Kozlov. *Combinatorial Algebraic Topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, 2008.
- [35] Janko Latschev. Vietoris–Rips complexes of metric spaces near a closed Riemannian manifold. *Archiv der Mathematik*, 77(6):522–528, 2001.
- [36] David Letscher. On persistent homotopy, knotted complexes and the Alexander module. In *Proceedings of the 3rd Innovations in Theoretical computer Science Conference*, pages 428–441. ACM, 2012.
- [37] Carlo Mantegazza. Notes on the distance function from a submanifold–v3. *CvGmt Preprint Server–<http://cvgmt.sns.it>*, 2010. <http://cvgmt.sns.it/media/doc/paper/1182/distancenotes.pdf>.
- [38] Carlo Mantegazza and Andrea Carlo Mennucci. Hamilton–Jacobi equations and distance functions on Riemannian manifolds. *Applied Mathematics & Optimization*, 47(1), 2003.
- [39] Facundo Mémoli and Osman Berat Okutan. Quantitative simplification of filtered simplicial complexes. *arXiv preprint [arXiv:1801.02812](https://arxiv.org/abs/1801.02812)*, 2018.
- [40] Amit Patel. Generalized persistence diagrams. *Journal of Applied and Computational Topology*, pages 1–23, 2018.
- [41] Peter Petersen. Classical differential geometry. .
- [42] Peter Petersen. Classical differential geometry, Apr 2016.
- [43] Samadwara Reddy. The Vietoris–Rips complexes of finite subsets of an ellipse of small eccentricity. Bachelor’s thesis, Duke University, April 2017.
- [44] Donald R Sheehy. Linear-size approximations to the Vietoris–Rips filtration. *Discrete & Computational Geometry*, 49(4):778–796, 2013.

- [45] Leopold Vietoris. Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen. *Mathematische Annalen*, 97(1):454–472, 1927.
- [46] Cédric Villani. Regularity of optimal transport and cut locus: From non-smooth analysis to geometry to smooth analysis. *Discrete Contin. Dyn. Sys.*, 30(2):559–571, 2011.

A From Euclidean points to a cyclic graph

Given a sample X of n points from a strictly convex differentiable planar curve C whose convex hull contains its evolute, it is easy to determine the cyclic graph structure on the 1-skeleton of $\text{VR}(X; r)$ even *without* knowledge of C . We can place the points in X in cyclically sorted order by running a convex hull algorithm, which takes $O(n \log n)$.

Determining the direction of each edge Given the n points X on C , suppose we are adding the next shortest undirected edge $\{x, y\}$ and need to determine whether this edge is oriented $x \rightarrow y$ or $y \rightarrow x$. First, we add another array to the algorithm in Section 4, to keep track of the degree of each vertex. Indeed, if any vertex v has full degree $n - 1$, then the Vietoris-Rips complex is contractible (it is a cone with v as its apex), and so we’re trivially done. So, consider the case where no vertex has degree $n - 1$. Going counterclockwise from x , all the other vertices in X fall into three categories: first the vertices v with an edge $x \rightarrow v$, then the vertices that are not connected to x , then the vertices v with an edge $v \rightarrow x$, before finally getting back to x (see Figure 7(right)). To determine the direction on the new edge $\{x, y\}$, note that since the graph is cyclic, y will be adjacent (in the cyclic order) to either a vertex v from the first category ($x \rightarrow v$) or to a vertex v from the third category ($v \rightarrow x$). If y is adjacent exclusively to a vertex v with $x \rightarrow v$, then the direction on our new edge is $x \rightarrow y$. If y is adjacent exclusively to a vertex v with $v \rightarrow x$, then the direction on our new edge is $y \rightarrow x$. Finally, if y is adjacent to a vertex of each type, then after adding the new edge $\{x, y\}$ necessarily x has degree $n - 1$, meaning that the Vietoris–Rips complex is contractible.

B Evolutes and injectivity

The goal of this section is to prove Lemma 6. Lemma 6 is used in the proof of Theorem 2 in order to prove (i) \iff (ii), namely that the evolute of C is contained in the convex hull of C if and only if the function $h: C \rightarrow C$ is injective.

We begin with some background on orientations. A differentiable closed simple curve $\alpha: I \rightarrow C \subseteq \mathbb{R}^2$ with α injective is said to be positively (resp. negatively) oriented if it is moving in the counterclockwise (resp. clockwise) direction around C . It follows from Lemma 3 that $(\alpha'(t), n(\alpha(t)))$ is a positive




Figure 7: We identify that the new edge, the dashed edge, must be oriented from bottom-right to top-left.

(resp. negative) basis for \mathbb{R}^2 , for all $t \in I$. The curve $\beta: I \rightarrow C$ is said to have a *matching orientation* to α when the basis $(\beta'(t), n(\beta(t)))$ has the same sign as $(\alpha'(t), n(\alpha(t)))$ [24].

Recall from Section 3 that for $\alpha: I \rightarrow \mathbb{R}^2$ a differentiable curve in the plane, we define the inner normal vector $n(t)$, the curvature $\kappa(t)$, and the center of curvature $x_\alpha(t) = \alpha(t) + \frac{1}{\kappa(t)}n(t)$.

Lemma 3. *Let C be a convex curve in \mathbb{R}^2 , and let $\alpha: I \rightarrow C$ be a differentiable curve moving in the counterclockwise (resp. clockwise) direction around C . Then $(\alpha'(t), n(\alpha(t)))$ is a positive (resp. negative) basis for \mathbb{R}^2 .*

Proof. Definition 2.1.2 and Theorem 2.4.2 of [41] show that the *signed curvature* $n(\alpha(t)) \cdot \alpha'(t)$ never changes sign for C convex, which implies that the orientation on the basis $(\alpha'(t), n(\alpha(t)))$ never changes signs. \square

Lemma 4. *Let C be a differentiable convex curve in \mathbb{R}^2 , and let $\alpha: I \rightarrow C$ be differentiable. The line through $\alpha(t)$ and $x_\alpha(t)$ intersects C at a unique other point, which we denote by $\beta(t) = h(\alpha(t))$. If we define $s: I \rightarrow \mathbb{R}$ to satisfy $\beta(t) = \alpha(t) + s(t)n(\alpha(t))$, then the function s is differentiable.*

Proof. We employ the implicit function theorem. Define $d^\pm: \mathbb{R}^2 \rightarrow \mathbb{R}$ to be the signed distance to the curve C , namely

$$d^\pm(y) = \begin{cases} d(x, C) & \text{if } x \in \text{conv}(C) \\ -d(x, C) & \text{otherwise.} \end{cases}$$

Since C is a smooth and complete manifold, it follows from Section 3 of [37] that d^\pm is differentiable on an open neighborhood of C , with derivative

$$\nabla d^\pm(x) = \begin{cases} \frac{x-y}{\|x-y\|} & \text{if } x \in \text{int}(\text{conv}(C)) \\ \frac{y-x}{\|x-y\|} & \text{if } x \notin \text{conv}(C) \\ n(x) & \text{if } x \in C, \end{cases}$$

where $y \in C$ is the unique closest point on C to x . Related references include [17, 26, 28, 38, 46]. Furthermore, define $g: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ via $g(t, s) = \alpha(t) + sn(\alpha(t))$, and define $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ by $f = d^\pm \circ g$. Note that g is differentiable since C is, and hence f is differentiable as the composition of d^\pm with g .

Pick t_0, s_0 such that $f((t_0, s_0)) = 0$; hence $g(s_0, t_0) \in C$. In order to apply the implicit function theorem we need to show that the Jacobian of f with respect to s is invertible at (t_0, s_0) ; this is equivalent to showing that $\frac{\partial f}{\partial s}(t_0, s_0) \neq 0$. Using the chain rule for $f = d^\pm \circ g$, we compute

$$\frac{\partial f}{\partial s}(t_0, s_0) = \nabla d^\pm(\alpha(g(t_0, s_0)))^T \cdot \frac{\partial g}{\partial s}(s_0, t_0) = n(g(t_0, s_0))^T \cdot n(\alpha(t_0)).$$

Suppose for a contradiction that the vectors $\alpha'(g(t_0, s_0))$ and $n(\alpha(t_0))$ were parallel. Then the normal to C at $\alpha(t_0)$ and the tangent to C at $g(t_0, s_0)$ would be the same line (they have the same direction vectors, and both pass through $g(t_0, s_0)$). This would mean that $\alpha(t_0)$ lives on the tangent line to C at $g(t_0, s_0)$ and that $\alpha'(t_0)$ is perpendicular to $\alpha'(g(t_0, s_0))$, contradicting convexity. Hence it must be that $\alpha'(g(t_0, s_0))$ and $n(\alpha(t_0))$ are not parallel, and therefore $\frac{\partial f}{\partial s}(t_0, s_0) \neq 0$.

It then follows from the implicit function theorem that there exists an open set U about $t \in \mathbb{R}$, and a differentiable function $s: U \rightarrow \mathbb{R}$, such that $s(t_0) = s_0$ and $f(t, s(t)) = 0$ for all $t \in U$. This gives the differentiability of $s: I \rightarrow \mathbb{R}$, as desired. \square

Lemma 5. *Let C be a differentiable convex curve in \mathbb{R}^2 , and let $\alpha: I \rightarrow C$ be a differentiable curve moving in the counterclockwise direction about C . Let $L(t)$ be the line through $\alpha(t)$ and $x_\alpha(t)$, namely $L(t) = \{\alpha(t) + sn(t) \mid s \in \mathbb{R}\}$. Suppose that $\beta(t) = \alpha(t) + s(t)n(\alpha(t))$ is an arbitrary⁴ differentiable curve with $\beta(t) \in L(t) \setminus \{x_\alpha(t)\}$ for all $t \in I$. Then $(\beta'(t), n(\alpha(t)))$ is a positive basis for \mathbb{R}^2 if and only $\beta(t)$ is in the same connected component of $L(t) \setminus \{x_\alpha(t)\}$ as $\alpha(t)$.*

Proof. We claim that when $s(t) = \frac{1}{\kappa(t)}$, we have $\alpha'(t) \cdot \beta'(t) = 0$. From the Frenet-Serret formulas [24], we have

$$n'(\alpha(t)) = \|\alpha'(t)\| \left(-\kappa(t) \left(\frac{\alpha'(t)}{\|\alpha'(t)\|} \right) + \tau(t)B(t) \right) = -\kappa(t)\alpha'(t), \quad (2)$$

where the last equality follows since the torsion term is $\tau(t) = 0$ for all curves

⁴In Lemma 4 we assume that β has image in C ; that is not necessary here.

in \mathbb{R}^2 . Note that when $s(t) = \frac{1}{\kappa(t)}$, we have

$$\begin{aligned}
\alpha'(t) \cdot \beta'(t) &= \alpha'(t) \cdot \left(\alpha'(t) + s'(t)n(\alpha(t)) + s(t)n'(\alpha(t)) \right) \\
&= \|\alpha'(t)\|^2 + \frac{1}{\kappa(t)} \alpha'(t) \cdot n'(\alpha(t)) \quad \text{since } \alpha'(t) \text{ and } n(\alpha(t)) \text{ are orthogonal} \\
&= \|\alpha'(t)\|^2 - \frac{1}{\kappa(t)} \|\alpha'(t)\| \|n'(\alpha(t))\| \quad \text{by (2)} \\
&= \|\alpha'(t)\| \left(\|\alpha'(t)\| - \frac{1}{\kappa(t)} \|n'(\alpha(t))\| \right) \\
&= 0 \quad \text{by (2)}.
\end{aligned}$$

Now, let's consider the case where $\alpha(t)$ and $\beta(t)$ are on the same connected component of $L(t) \setminus \{x_\alpha(t)\}$. Since $\alpha'(t)$ is perpendicular to $n(\alpha(t))$, it suffices to show that the dot product $\alpha'(t) \cdot \beta'(t)$ is positive. Since $\alpha(t)$ and $\beta(t)$ are on the same connected component, we know that $s(t) < \frac{1}{\kappa(t)}$. Since

$$\alpha'(t) \cdot \beta'(t) = \|\alpha'(t)\|^2 - \frac{1}{\kappa(t)} \|\alpha'(t)\| \|n'(\alpha(t))\| = 0,$$

it must be the case that for $s(t) < \frac{1}{\kappa(t)}$ we have

$$\alpha'(t) \cdot \beta'(t) = \|\alpha'(t)\|^2 - s(t) \|\alpha'(t)\| \|n'(\alpha(t))\| > 0.$$

Finally, consider the case where $\alpha(t)$ and $\beta(t)$ are not on the same connected component of $L(t) \setminus \{x_\alpha(t)\}$. So $s(t) > \frac{1}{\kappa(t)}$. This gives us that

$$\alpha'(t) \cdot \beta'(t) = \|\alpha'(t)\|^2 - s(t) \|\alpha'(t)\| \|n'(\alpha(t))\| < 0.$$

□

Lemma 6. *Let C be a convex curve in \mathbb{R}^2 , and let $\alpha: I \rightarrow C$ be a differentiable curve moving in the counterclockwise direction about C . The line through $\alpha(t)$ and $x_\alpha(t)$ intersects C at a unique point of $C \setminus \{\alpha(t)\}$, which we denote by $\beta(t) = h(\alpha(t))$. Then α and β have matching orientations at time t if and only if $x_\alpha(t) \in \text{conv}(C)$.*

Proof. Note that, by definition of convexity, C lies completely on one side of its tangent lines. A vector $v \in \mathbb{R}^2$ with its tail placed at $p \in C$ points toward the interior of C if it is completely contained on the same side of the tangent line to C at p as C . For such a vector, there exists some constant $c > 0$ such that cv intersects $C \setminus \{p\}$ at a unique point \tilde{p} . If we instead place the tail of v at \tilde{p} , then v is on the side of the tangent line to C at \tilde{p} that does not contain C . Thus, v points toward the exterior of C at \tilde{p} . By definition, the unit normal vector to C at each $p \in C$ points to the interior of C . This means that $n(\alpha(t))$ points to the interior of C when its tail is placed at $\alpha(t)$, and to the exterior of C when

its tail is placed at $\beta(t)$. In summary, when their tails are placed at $\beta(t)$, both of the vectors $-n(\alpha(t))$ and $n(\beta(t))$ point toward the interior of C .

Define $s: I \rightarrow \mathbb{R}$ to satisfy $\beta(t) = \alpha(t) + s(t)n(\alpha(t))$; note that s is differentiable by Lemma 4. Suppose $x_\alpha(t) \in \text{conv}(C)$. So $0 < \frac{1}{\kappa(t)} < s(t)$. From Lemma 5, $(\beta'(t), n(\alpha(t)))$ is a negative basis for \mathbb{R}^2 , and hence $(\beta'(t), -n(\alpha(t)))$ is a positive basis for \mathbb{R}^2 . So it must be the case that $(\beta'(t), -n(\alpha(t)))$ and $(\beta'(t), n(\beta(t)))$ are bases of the same sign. Thus, α and β have matching orientations.

Next, suppose that $x_\alpha \notin \text{conv}(C)$. So $s(t) < \frac{1}{\kappa(t)}$. From Lemma 5, $(\beta'(t), n(\alpha(t)))$ is a positive basis. However, at $\beta(t)$, the vector $n(\alpha(t))$ points outward while $n(\beta(t))$ must point inward. Thus, $(\beta'(t), n(\beta(t)))$ is a negative basis, and so α and β do not have matching orientations. \square