



OPEN ACCESS

EDITED BY

Quoc Thong Le Gia,
University of New South Wales, Australia

REVIEWED BY

Sabah Jassim,
University of Buckingham, United Kingdom
Francesco Marchetti,
University of Padua, Italy

*CORRESPONDENCE

Amish Mishra
✉ amishra2019@fau.edu

RECEIVED 03 March 2023

ACCEPTED 21 June 2023

PUBLISHED 06 July 2023

CITATION

Mishra A and Motta FC (2023) Stability and machine learning applications of persistent homology using the Delaunay-Rips complex. *Front. Appl. Math. Stat.* 9:1179301. doi: 10.3389/fams.2023.1179301

COPYRIGHT

© 2023 Mishra and Motta. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Stability and machine learning applications of persistent homology using the Delaunay-Rips complex

Amish Mishra* and Francis C. Motta

Department of Mathematical Sciences, Florida Atlantic University, Boca Raton, FL, United States

Persistent homology (PH) is a robust method to compute multi-dimensional geometric and topological features of a dataset. Because these features are often stable under certain perturbations of the underlying data, are often discriminating, and can be used for visualization of structure in high-dimensional data and in statistical and machine learning modeling, PH has attracted the interest of researchers across scientific disciplines and in many industry applications. However, computational costs may present challenges to effectively using PH in certain data contexts, and theoretical stability results may not hold in practice. In this paper, we define, implement, and investigate a simplicial complex construction for computing persistent homology of Euclidean point cloud data, which we call the Delaunay-Rips complex (DR). By only considering simplices that appear in the Delaunay triangulation of the point cloud and assigning the Vietoris-Rips weights to simplices, DR avoids potentially costly computations in the persistence calculations. We document and compare a Python implementation of DR with other simplicial complex constructions for generating persistence diagrams. By imposing sufficient conditions on point cloud data, we are able to theoretically justify the stability of the persistence diagrams produced using DR. When the Delaunay triangulation of the point cloud changes under perturbations of the points, we prove that DR-produced persistence diagrams exhibit instability. Since we cannot guarantee that real-world data will satisfy our stability conditions, we demonstrate the practical robustness of DR for persistent homology in comparison with other simplicial complexes in machine learning applications. We find in our experiments that using DR in an ML-TDA pipeline performs comparatively well as using other simplicial complex constructions.

KEYWORDS

persistent homology, topological data analysis, machine learning, persistence diagram, stability, Delaunay triangulation, Vietoris-Rips, simplicial complex

1. Introduction

As the volume and complexity of data collected in the experimental sciences and in industry applications continues to grow, scientists and engineers often turn to methods that transform data into more compact and manageable representations, while retaining crucial characteristics that enable meaningful data analysis, visualization, and support the development of high-performing predictive models. Tools from the applied computational topology subfield Topological Data Analysis (TDA) offer several solutions to address some of the challenges with processing and analyzing complex or high-dimensional data. TDA leverages results from algebraic topology to measure and quantify qualitative, shape-based features of data and continues to attract the interest of researchers across computational, mathematical, and experimental disciplines including financial networks [1], signals in images [2], cancer detection [3], forensics [4], and material science [5].

Persistent homology (PH) is the flagship method within TDA and provides a robust method to encode multi-dimensional geometric/topological features of a dataset in a compact representation known as a persistence diagram (PD). PH is often computed from point cloud data, i.e., a finite collection of points in a metric space. One begins by associating to the cloud a nested family of topological spaces—namely a collection of simplicial complexes consisting of simplices of various dimension (e.g., points, edges, triangles, tetrahedra, etc.)—that is parameterized by a (real) scale parameter. This filtration of simplicial complexes is meant to capture geometric/topological structures in the point cloud across scales. From the filtration, a PD, consisting of a collection of ordered pairs above the diagonal in \mathbb{R}^2 . Each (birth, death)-pair represents a topological feature of a fixed dimension (connected component, hole, void, etc.) which appears at the birth scale and disappears at the death scale.

Construction of a filtration depends on two factors: which simplices to include in the family of complexes, and determination of the scales at which each simplex should appear. Both factors contribute to the overall computational burden to compute PH on point cloud data, and may represent barriers to deploying PH on large or high-dimensional data sets. For instance, an appealing and often used construction known as the Rips filtration requires only knowledge of the pairwise dissimilarities between points. This makes computing the scales at which all simplices should appear quite straightforward. However, the size of the Rips complex grows exponentially in the number of points. An alternative to the Rips filtration is the Alpha complex filtration [6], which restricts the complexes to simplices in the Delaunay triangulation of the cloud. However, this method assigns scales to simplices according to the neighborhood of the Voronoi cells associated with each data point, which can be an expensive computation. Some other solutions [7–9] have more recently been proposed to reduce the size of the complexes ingested by the PH algorithm, which has been implemented in a variety of languages (C++, Python, Julia, etc.) and optimized to further improve computational efficiencies [10–12].

In this work, we take inspiration and combine the computational efficiencies of the Rips and Alpha filtrations to define a new filtration on Euclidean point cloud data. This new construction is built on what we refer to as the Delaunay-Rips (DR) complex. The DR complex may be viewed as a special case of the lazy-witness complex [9, 13], in that the simplices in the DR complex and DR filtration (by extension) are those in the Delaunay triangulation of the point cloud. We also establish that, generically, this filtration will inherit previously established stability results which guarantee that, roughly, certain small changes in the underlying data will result in small changes in the PD [14, 15]. That said, we further prove that, for certain configurations of points, the DR filtration suffers from an instability in which an arbitrarily small perturbation of the underlying data can result in a non-infinitesimal change in the PD. We provide an unconventional proof using an exact reduction of the boundary matrix constructed by the DR filtration and show the explicit persistence calculation to demonstrate the instability. This instability is due to the instability in the Delaunay triangulation itself at degenerate configurations of points.

Much of the fundamental appeal of PH technologies is due to the numerous stability results that show the transformations sending data to diagrams and/or diagrams to vectorized topological representations are Lipschitz continuous transformations [16]. However, it is often the case in practice that the topological representations derived from PDs—and used in data analysis and statistical model development—are not stable, despite the purported success of the models on which they're based. Thus, we are further motivated to ask, to what extent does the instability observed in the DR filtration matter in practice? To the best of our knowledge, the effects of persistence diagram instabilities on trained model performance has not been done, although this investigation is important in assessing model fidelity. We interrogate this question empirically by performing systematic comparisons of the performance and costs of the Alpha, Rips, and DR filtrations on synthetic and real datasets and machine learning (ML) modeling tasks. We find the effect of the instability is negligible in these cases while the practical computational advantages enjoyed by DR can be significant. Thus, the contributions of this paper are:

- We define the DR complex, which is a refinement of the Vietoris-Rips simplices on the data using the Delaunay triangulation as a backbone.
- We offer an algorithm and a Python implementation of the DR filtration.
- We document some empirical runtime comparisons for constructing the PD using DR with popular implementations of the Rips [11] and Alpha [17] filtrations.
- We provide a straightforward proof of the stability of the Delaunay triangulation in some neighborhood of a generic point cloud, which establishes stability of the DR filtration on an appropriately chosen neighborhood of the data.
- We provide a rare proof of an instability in a filtration used for computing PH.
- We examine the practical impact of the instability on synthetic and real-world classification tasks and find the instability has limited impact on model performance.

This paper is organized as follows: In Section 2 we briefly review the necessary mathematical preliminaries. Section 3 formally introduces the DR filtration, provides pseudo-code of the implementation we used to compute the DR filtration on point cloud data, and compares the empirical runtime of this algorithm to implementations of Rips and Alpha constructions across increasing point cloud size and dimension. In Section 3.3 we discuss the stability properties of the DR filtration and demonstrate—through a by-hand calculation of a family of PDs—how a discontinuity in the transformation from point cloud to PD can arise given a perturbation in the underlying cloud. In Section 4 we report the result of several systematic comparisons of ML model performance trained on topological features derived from Alpha, Rips, and DR filtrations on synthetic and experimental data, including using random forest classifiers trained on persistence image (PI) vectorizations [18] of persistence diagrams, and support vector machines trained on persistence statistics feature vectors derived from ECG time-series data [19].

2. Background

To extract topological features from a point cloud, one often begins by treating the points as the *vertices* of a so-called *simplicial complex* that consists of vertices, edges (pairs of vertices), triangles (sets of three vertices), tetrahedra (sets of four vertices), and higher-dimensional analog (sets of $n > 4$ vertices), to construct an object with well-defined notions of “shape.” Our interest is in computing algebraic objects, namely persistent homology groups, that characterize topological (homological) invariants (e.g., connected components, holes, voids, etc.) across scales. We briefly review these notions here. For a deeper treatment of the mathematical foundations of these sections, we refer the reader to Edelsbrunner and Harer [20] and Hatcher [21]. For a more computational treatment with many practical considerations and examples, we recommend Otter et al. [22].

2.1. Simplicial homology

Definition 2.1. Let K_0 be a finite set and $\mathcal{P}(K_0)$ the powerset (i.e., the set of all subsets) of K_0 . An **abstract simplicial complex built on K_0** is a collection, $K \subset \mathcal{P}(K_0)$, of non-empty subsets of K_0 with the properties that $\{v\} \in K$ for all $v \in K_0$, and if $\sigma \in K$ then $\tau \in K$ for all $\tau \subseteq \sigma$.

In this paper we are concerned with simplicial complexes built from point clouds such that the points are identified with the singleton sets in the complex. In general, the elements of a simplicial complex, K , will be called **simplices**, while we may refer to sufficiently small subsets by other names. For example, we will refer to singleton sets as **vertices** of K , size-two sets as **edges**, etc. We say that a simplex has **dimension p** or is a **p -simplex** if it has size $p + 1$; so vertices are dimension 0 simplices, edges are dimension 1, triangles are 2-simplices, etc. We denote a p -simplex by $[v_0 v_1 \dots v_p]$ if it contains the 0-simplices $v_i, i = 0, \dots, p$. K_p denotes the collection of all p -simplices and the **k -skeleton** of K is the union of the sets K_p for all $p \in \{0, 1, \dots, k\}$. If τ and σ are simplices such that $\tau \subset \sigma$, then we call τ a **face** of σ . We say that τ is a face of σ of **codimension q** if the dimensions of τ and σ differ by q . The **dimension** of K is defined as the maximum of the dimensions of any of its simplices.

Definition 2.2. Let K be a simplicial complex and $p \geq 0$. A **p -chain** of K is a formal sum of p -simplices in K written as

$$c = \sum a_i \sigma_i$$

where $a_i \in \mathbb{F}$ is a field, and σ_i are p -simplices. The p -chains then form a **vector space of p -chains** over a field \mathbb{F} , which we denoted $C_p(K)$.

$C_p(K)$ and $C_{p-1}(K)$ are naturally related by a linear map called the **boundary map** that sends each p -chain to its boundary ($p - 1$)-chain.

Definition 2.3. Let K be an n -dimensional simplicial complex and $\sigma = [u_0 u_1 \dots u_p] \in C_p(K)$. The **p -boundary map**, ∂_p is defined on p -simplices to be

$$\partial_p \sigma = \sum_{j=0}^p (-1)^j [u_0 \dots \hat{u}_j \dots u_p]$$

where the hat indicates that u_j is omitted. Extending via linearity, $\partial_p : C_p(K) \rightarrow C_{p-1}(K)$ is further defined on any p -chain, $c = \sum_i a_i \sigma_i$, by

$$\partial_p c = \sum_i a_i \partial_p \sigma_i.$$

Intuitively the chain of codimension-1 faces of a simplex encode the boundary of that simplex. Let’s take a look at some examples. The boundary of the triangle $[abc]$ is the chain $[bc] - [ac] + [ab]$, consisting of the three edges that form its boundary. Extending to chains, the boundary of the two edges that meet at the vertex b (namely $[ab] + [bc]$) would be $[b] - [a] + [c] - [b] = [c] - [a]$, reflecting the fact that the boundary of this 1-dimensional object consists only of its two terminal vertices. Connecting the chain groups via their boundary maps forms a **chain complex**

$$0 \rightarrow C_n(K) \xrightarrow{\partial_n} \dots C_{p+1}(K) \xrightarrow{\partial_{p+1}} C_p(K) \xrightarrow{\partial_p} C_{p-1} \xrightarrow{\partial_{p-1}} \dots C_0(K) \rightarrow 0,$$

where the first map is the trivial linear map that sends the 0-vector in the trivial vector space, 0, to the 0-vector in $C_n(K)$.

Informally, a p -dimensional hole in a simplicial complex will be represented by a $(p - 1)$ -chain that could be (but is not) the boundary of a p -chain. For example, as we have seen the boundaryless 1-chain $[bc] + [ac] + [ab]$ is the boundary of a 2-chain (the 2-simplex $[abc]$). However, if $[abc]$ were not in the simplicial complex, we would be justified in saying the complex contained a hole enclosed by the edges of the missing triangle. To make these notions precise we introduce two subspaces of $C_p(K)$ that respectively encode all the p -chains that are without boundary, and those which are actually boundaries of $p + 1$ chains.

Definition 2.4. A **p -cycle**, $\gamma \in C_p(K)$, is a p -chain with no boundary, i.e., $\partial_p \gamma = 0$. The collection of all p -cycles in a simplicial complex K is denoted $Z_p(K)$ and forms a subspace of $C_p(k)$ because $Z_p(K) = \ker \partial_p$.

Definition 2.5. A **p -boundary**, $\beta \in C_p(K)$, is a p -chain that is the boundary of a $(p + 1)$ -chain, i.e., $\partial_{p+1} \sigma = \beta$ for some $\sigma \in C_{p+1}(K)$. The collection of p -boundaries in the simplicial complex K , denoted $B_p(K)$, is also a subspace of $C_p(K)$ since $B_p = \text{im } \partial_{p+1}$.

With p -cycles and p -boundaries defined, we can formally define the object which captures representatives of holes in a simplicial complex: those cycles which are not boundaries.

Definition 2.6. The **p -th homology group** is $H_p = Z_p(K)/B_p(K) = \ker \partial_p / \text{im } \partial_{p+1}$. The **p -th Betti number** is the dimension of the quotient vector space, $\beta_p = \text{rank } H_p$.

H_p is then the vector space of equivalence classes of p -cycles, where two p -cycles are equivalent if they differ by a p -boundary. The fact that H_p is well-defined relies on the fundamental result that ensures the boundary of a boundary chain must be empty, so that B_p is actually a subspace of Z_p .

Lemma 2.7. (*Fundamental lemma of homology* [20]) The composition of any two consecutive boundary maps in the chain complex is trivial. That is,

$$\partial_p \partial_{p+1}(\sigma) = 0$$

for every integer p and every $(p + 1)$ -chain σ .

As suggested, the p -th Betti number corresponds to the number of p -dimensional “holes” in the corresponding simplicial complex. In fact, the 0-th Betti number counts the number of connected components, the 1-st Betti number counts the number of (independent) loops, the 2-nd Betti number counts the number of (independent) voids, etc.

2.2. Persistent homology

Although simplicial homology is sufficient for capturing some intrinsic shape characteristics of a fixed simplicial complex, the natural question that arises with point cloud data is how to construct a simplicial complex from the points to provide a meaningful representation of latent structure in the cloud. In fact, there may be topological features of interest for one complex built on the data (for instance at some fine scale), and another set of topological features (at a larger scale) of interest from another complex. A promising approach to deal with this question is persistent homology, an extension of homology that considers a parameterized family of simplicial complexes, rather than just one.

Definition 2.8. Let K_s be a simplicial complex for each $s \in \mathbb{R}$ such that $K_s \subset K_t$ for all $s \leq t$. We refer to such a collection of complexes as a **filtration**. When each $K_s \subseteq K$, for some fixed finite simplicial complex K , and $K_t = K$ for some $t \in \mathbb{R}$, we’ll refer to the parameterized collection as a **filtration of K** .

Note that a filtration on a finite simplicial complex K necessarily only contains finitely many distinct sub-complexes, which we can relabel

$$\emptyset := K_0 \subseteq K_1 \subseteq K_2 \subseteq \dots \subseteq K_n = K.$$

Here we have relabelled $K_i := K_{s_i}$ for the scales $s_i \leq s_{i+1}$, at which the sub-complexes change. Each simplex $\sigma \in K$ may also be assigned the minimum scale at which it appears in the filtration. In other words, one may define $f: K \rightarrow \mathbb{R}$ by $f(\sigma) = t$ if $\sigma \in K_t$ and $\sigma \notin K_s$ for any $s < t$. Thus, a filtration induces a so-called *monotonic function* on the simplices of K , where $f(\sigma) \leq f(\tau)$ if σ is a face of τ . Conversely, any monotonic function $f: K \rightarrow \mathbb{R}$ with the property that $f(\sigma) \leq f(\tau)$ if $\sigma \subseteq \tau$, defines a filtration on K by taking $K_s = f^{-1}(\infty, s]$.

Definition 2.9. Let $\emptyset \subseteq K_1 \dots \subseteq K_n = K$ be a filtration of K . The **p -th persistent homology groups** of the filtration are then formally defined as

$$H_p^{i,j} = Z_p(K_i) / (B_p(K_j) \cap Z_p(K_i)).$$

for $0 \leq i \leq j \leq n$. The corresponding **p -th persistent Betti numbers** are the ranks of these groups,

$$\beta_p^{i,j} = \text{rank } H_p^{i,j}.$$

An element of $H_p^{i,j}$ corresponds to a cycle in the filtration that persisted from K_i to K_j (i.e., a cycle in K_i that did not become a boundary in K_j) and the content of all the p -th persistent homology groups of a filtration can be summarized

in a dimension- p persistence diagram which tracks the pairs of indices in the filtration at which homological features first appear and later disappear. It is common practice to say a feature is “born” in complex K_i and “dies” in complex K_j if the scale it became a cycle (without being a boundary) is in K_i and earliest complex it becomes a boundary is K_j .

Definition 2.10. Let $\mu_p^{i,j}$ be the number of p -dimensional classes born in complex K_i that die entering complex K_j . Then

$$\mu_p^{i,j} = (\beta_p^{i,j-1} - \beta_p^{i,j}) - (\beta_p^{i-1,j-1} - \beta_p^{i-1,j}).$$

The **p -persistence diagram** of the filtration given by $f: K \rightarrow \mathbb{R}$, denoted $\text{Dgm}_p(f)$, is a multiset of pairs (s_i, s_j) in the extended real plane \mathbb{R}^2 with multiplicity $\mu_p^{i,j}$. Each pair (s_i, s_j) represents a nontrivial persistent homology class that is born in complex $K_i = K_{s_i}$ and dies upon entering complex $K_j = K_{s_j}$ because it merges with a homology class that was born before K_i .

Now, we can define a metric we can use to find the distance between two p -persistence diagrams.

Definition 2.11. Let X and Y be two p -persistence diagrams. Define $\|x - y\|_\infty := \max\{|x_1 - y_1|, |x_2 - y_2|\}$ for $x = (x_1, x_2) \in X, y = (y_1, y_2) \in Y$. We define the **bottleneck distance** between the diagrams as

$$W_\infty(X, Y) = \inf_{\eta: X \rightarrow Y} \sup_{x \in X} \|x - \eta(x)\|_\infty$$

where the infimum is taken over all bijections η . Note that X and Y have countably infinitely many copies of the diagonal (points with no persistence), thus allowing η to be a bijection that may pair non-zero persistence points to the diagonal.

A valuable property of persistence diagrams built from filtrations on a fixed complex K is their stability with respect to changes in the monotonic function determining filtrations on K :

Theorem 2.12. [20] Let K be a finite simplicial complex and $f, g: K \rightarrow \mathbb{R}$ two monotonic functions. For each dimension p , the bottleneck distance between the diagrams $\text{Dgm}_p(f)$ and $\text{Dgm}_p(g)$ satisfies

$$W_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) \leq \|f - g\|_\infty,$$

where $\|f - g\|_\infty = \max_{\sigma \in K} |f(\sigma) - g(\sigma)|$.

We will also use a notion of distance between subsets of a metric space to control the distance between point clouds living in \mathbb{R}^D .

Definition 2.13. Let X and Y be two non-empty subsets of a metric space (M, d) . We define their **Hausdorff distance** $d_H(X, Y)$ by

$$d_H(X, Y) := \inf\{\varepsilon > 0; X \subseteq Y_\varepsilon \text{ and } Y \subseteq X_\varepsilon\},$$

where

$$X_\varepsilon := \bigcup_{x \in X} \{m \in M; d(x, m) \leq \varepsilon\}.$$

For finite subsets $X, Y \subset (M, d)$ (i.e., finite point clouds), the Hausdorff distance reduces to the maximum distance from a point in one set to the closest point in the other set. We prove a related statement in Lemma 3.3.

The formal notion of a filtration on a simplicial complex and its persistent homology groups defined in this section provide a means to extract multiscale structure from point cloud data, and thereby alleviate some of the concern about which complex may best capture structure in the cloud. Still, the methods of constructing a filtration from a point cloud are numerous and come with advantages and disadvantages that depend on the nature of the data.

2.3. Vietoris-Rips and Alpha complexes

One of the simplest and most commonly used methods to build a filtration on a finite point cloud, $X \subset (M, d)$, is to treat the points as vertices and add simplices at a scale determined by their diameter in the metric space.

Definition 2.14. Let $X \subset \mathbb{R}^D$ be a point cloud and $\varepsilon \geq 0$. The **Vietoris-Rips complex** at scale ε is defined as

$$VR_\varepsilon(X) = \{\sigma \subseteq X \mid d(x, x') \leq 2\varepsilon, \forall x, x' \in \sigma\}.$$

In other words, for a given scale $\varepsilon \geq 0$, if $d(x, x') \leq 2\varepsilon$ for $x, x' \in X$, one adds the p -simplex $\sigma = [x_0 x_1 \dots x_k]$ to the complex at the largest scale of any of its edges.

Although an algorithm to compute the Rips complex at any scale is simple to implement, constructing it on point cloud, X , with large numbers of points results in a computational challenge: eventually (at scales at and beyond half the diameter of the point cloud) the Rips complex will be equal to the powerset of X and so will contain $2^{|X|}$ simplices. Moreover, the Rips complex will eventually contain simplices of all dimensions (up to the size of the point cloud minus 1), and so will contain homological information even beyond the dimension of the cloud (assuming the data lives in a finite-dimensional vector space). In practice this can be mitigated by imposing a restriction on the maximum dimension of simplices included in any complex in the filtration. Moreover, it is often the case in practice that, as the scale increases, additional simplices may appear in the complex that do not affect its homology [23, 24]. Filtration methods which avoid inclusion of “extraneous” simplices may be preferable for large point clouds [7–9]. Before defining examples of such methods, it will be helpful for the remainder of the paper to define when a Euclidean point cloud is in “general position”.

Definition 2.15. A set of points in a d -dimensional Euclidean space is in **general position** if no $d + 2$ of them lie on a common $(d - 1)$ -sphere.

For example, this means for a set of points to be in general position in \mathbb{R}^2 , no 4 of them can be co-circular. This condition ensures that each subset of $d + 1$ points lie on a unique d -dimensional sphere which will ensure a unique Delaunay triangulation, defined below.

Let $X \subset \mathbb{R}^D$ be a finite point cloud. Let $x \in X$ and define

$$V_x := \{p \in \mathbb{R}^D \mid d(p, x) \leq d(p, x') \forall x' \in X\}.$$

Each V_x is called a **Voronoi cell** of X and captures all points which are not closer to any other point in X than x . Note that $\{V_x\}_{x \in X}$ forms a cover of \mathbb{R}^D . This cover is known as the Voronoi decomposition of \mathbb{R}^D with respect to X . To construct the Delaunay triangulation from this cover, we define

$$\text{Del}(X) := \{\sigma \subseteq X \mid \bigcap_{x \in \sigma} V_x \neq \emptyset\}.$$

It is known that $\text{Del}(X)$ is itself a simplicial complex [20]. An n -simplex $\sigma \in \text{Del}(X)$ will be referred to as a **Delaunay simplex**. We will use $\text{Del}(X)$ as the underlying structure when defining the Delaunay-Rips complex in Section 3.1.

First we recall another commonly used filtration construction, known as the Alpha filtration, well studied for point clouds $X \subset \mathbb{R}^D$. We recall the definition given in III.4 of Edelsbrunner and Harer [20].

Definition 2.16. Let $\varepsilon \geq 0$ and let $S_x(\varepsilon) := V_x \cap B_x(\varepsilon)$, where $B_x(\varepsilon)$ is the d -dimensional ball of radius ε centered on $x \in X$. The **Alpha complex** at scale $\varepsilon \geq 0$ is

$$\text{Alpha}_\varepsilon(X) = \{\sigma \subseteq X \mid \bigcap_{x \in \sigma} S_x(\varepsilon) \neq \emptyset\}.$$

Note that since $S_x(\varepsilon) \subseteq V_x$, the set of 1-simplices of the Alpha_ε complex form a subcomplex of the 1-skeleton of the Delaunay triangulation.

By construction, $VR_\varepsilon(X)$ and $\text{Alpha}_\varepsilon(X)$ are simplicial complexes for all $\varepsilon \in \mathbb{R}$, and if $s \leq t$, both $\text{Alpha}_s \subseteq \text{Alpha}_t$ and $VR_s \subseteq VR_t$. Thus, each filtration construction yields an ordering on a set of simplices in a simplicial complex built on the point cloud X . For Vietoris-Rips, the scale of each simplex is determined by the distance between the farthest two vertices that define the simplex. However, Vietoris-Rips also assigns a non-zero weight to every subset of a set of vertices which is an exponentially slow computation in the number of vertices. Alpha, on the other hand, does not compute scales for every subset of the set of vertices. Rather, the scales assigned to simplices are determined by when the restricted epsilon balls on the Voronoi cells intersect. This additional computation is what we seek to avoid in our construction of the Delaunay-Rips complex in the subsequent section.

3. The Delaunay-Rips complex and stability

3.1. Definition and construction

Combining the Alpha and Rips constructions provides an alternative method of building a family of complexes on a point cloud $X \subset \mathbb{R}^d$. The idea is similar to the construction of the Delaunay-Čech complex defined in Bauer and Edelsbrunner [25] and can be seen as a special case of a lazy weak witness complex [9]. Delaunay-Rips utilizes the conceptual simplicity of the Vietoris-Rips complex while cutting down on the number of high dimensional and potentially extraneous simplices. The idea is that we build the Vietoris-Rips complex on X but only add edges if the edges occur in the Delaunay 1-skeleton of the point cloud. The

higher dimensional p -simplices are then added in the traditional Vietoris-Rips manner, i.e., if and only if their 1-skeletons appear.

Definition 3.1. The **Delaunay-Rips (DR) complex** for a given scale $\varepsilon \geq 0$ is defined,

$$\text{DR}_\varepsilon(X) = \{\sigma \subseteq \text{Del}(X) \mid d(x, x') \leq 2\varepsilon, \forall x, x' \in \sigma\}.$$

Just as previously with Vietoris-Rips and Alpha filtrations, the Delaunay-Rips filtration can be thought of as a method for building a complex on a point cloud and assigning a scale, and thus a monotonic ordering, to the simplices in our complex built on X .

Figure 1 illustrates and compares examples of Rips, Alpha, and Delaunay-Rips filtrations, built on a cartoon point cloud in \mathbb{R}^2 , and their associated H_0 and H_1 persistence diagrams. The Rips filtration produces a homology class that is born at scale value 7.26 and persists until dying at scale value 10.21. This is represented by the birth and death coordinates of the single orange point in the PD. Notice that although there are two loops in the topological surface of the overlapping circles at scale value 8.55, Rips does not capture the second loop in the upper portion because a triangle is added as soon as its 3 boundary edges appear. This is in contrast to the Alpha filtration which produces a homology class at scale value 7.26 that persists until scale value 10.88 and another H_1 class born at 8.55 and dies at scale value 9.23. Recall that the Alpha complex at a particular scale value requires the balls centered at each 0-simplex of a d -simplex in \mathbb{R}^D to be restricted to the Voronoi cells of the respective 0-simplex. Hence, the d -simplex is only added to the complex when all of the restricted balls intersect. In our example, the dashed red edge—which appears at scale 8.55, along with the higher dimensional triangles it creates, in the Rips and DR filtrations—does not appear until scale value 9.23 in the Alpha filtration due to the aforementioned property of the Alpha complex. Particularly, that edge needs to wait until the radii of the balls equals the radius of the circle containing the 3 points that make up the simplex that edge lies on.

The DR filtration produces an H_1 homology class that persists from scale value 7.26 to 10.86. In the DR construction, we add triangles as soon as their 3 edges appear in the complex: this is illustrated as a shaded triangle in the figure. Notice how at $r = 18$ although many balls overlap, we have only added the simplices that appear in the Delaunay triangulation of the point cloud. Observe that for a fixed point cloud, $X \subset \mathbb{R}^D$, the abstract simplicial complex built on X (Definition 2.1) according to both the Alpha and Del-Rips filtrations is the Delaunay triangulation of X , $\text{Del}(X)$. Of course, the weights assigned to the simplices of $\text{Del}(X)$ by Del-Rips may differ from the weights assigned by the Alpha filtration. A notable comparison here is that the Del-Rips PD shows an H_1 class born at the same scale value as it was in the Rips filtration, but it persists for longer (due to the combinatorial cut down in simplices). Another notable comparison is that Del-Rips does not capture a second H_1 class in the PD like Alpha and the only H_1 class that is captured does not persist for quite as long as the corresponding H_1 class in Alpha.

As a remark, the reader may wonder if that second loop should be captured in the PD or not. Although there is a scale value at which the balls overlap in such a way as to reveal two loops in the data, it is not clear whether capturing the less persistent

```

Input:  $P = (p_1, \dots, p_n)$ ,  $\text{dim}$  = maximum homology
        dimension to compute
Output: Delaunay Rips Filtration
triangulation  $\leftarrow$  Delaunay( $P$ )
filtration  $\leftarrow$   $[[[i], 0]]$  for  $1 \leq i \leq n$  /* Add the
        0-simplices into the filtration */
for each simplex  $\in$  triangulation do
    for  $d \leftarrow 1$  to  $\text{dim} + 1$  do
        faces  $\leftarrow$  all  $d$ -simplex subsets of current
        simplex
        for each face  $\in$  faces do
            if face  $\notin$  filtration and  $d == 1$  then
                value  $\leftarrow$  distance(face[0], face[1])
                /* Calculate the euclidean distance
                */
                filtration.append((face, value)) /* Append
                1-simplices */
            end
            if face  $\notin$  filtration and  $\text{dim} > 1$  then
                find subface of face with greatest value
                filtration.append((subface, value)) /* Add
                higher order simplices */
            end
        end
    end
end

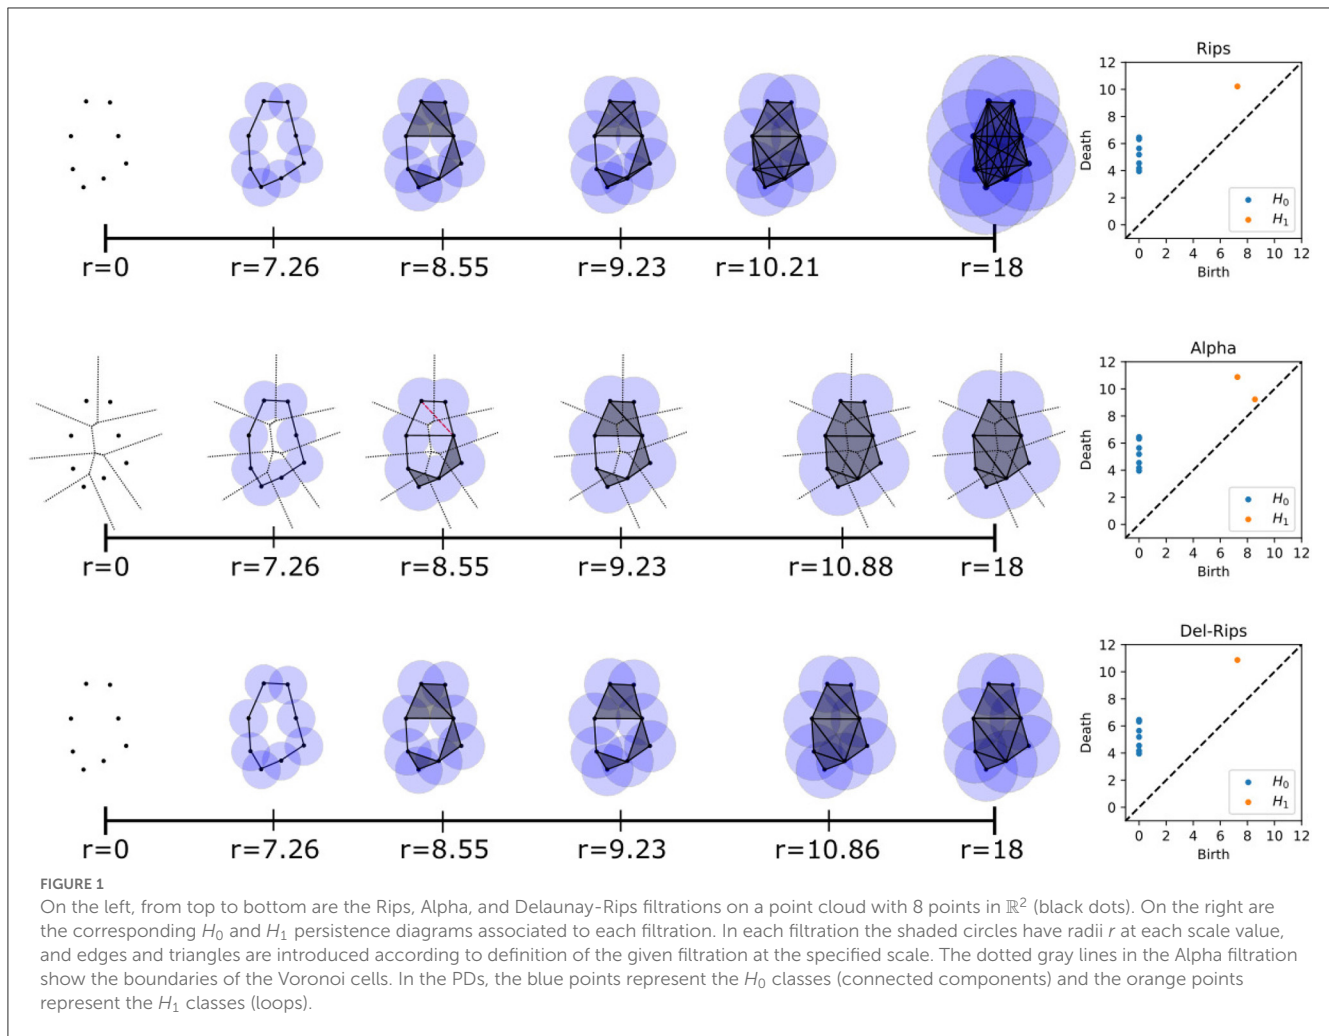
```

Algorithm 1. An algorithm to compute Delaunay-Rips filtration

loop would be valuable in any particular application. Employing Alpha certainly captures that other loop well, but at the cost of computational efficiency as it seeks to balance the weights on the simplices across varying dimensions appropriately. Section 4 partially addresses whether we can sacrifice the fidelity of Alpha to the underlying topological structure of the data for the advantage of a computational speed-up.

3.2. Implementation and runtime analyses

Here we show empirical results of the performance of computing persistence diagrams using the Delaunay-Rips filtration on varying datasets. We fix our field of coefficients to be \mathbb{Z}_2 when computing persistent homology groups. **Algorithm 1**, which we have implemented in Python [26], constructs the DR filtration across scales. This filtration then gets passed to the Persistent Homology Algorithms Toolbox (PHAT) [27] to construct the boundary matrix, reduce the boundary matrix, and extract the persistent pairs. Experiments were run on a computer with an Intel i7-10875H processor running at 2.3 GHz, 64 GB of RAM, and running Ubuntu 20.04.5 LTS. The runtimes are measured as the time between the data set being input into the corresponding algorithm and the persistence diagram being produced. For comparison, we choose the Python module Ripser [11] and a Python Alpha implementation from the Python package Cechmate [17]. **Figure 2** shows how runtimes vary with the number of points



being sampled from a noisy 2-sphere. Each data point on the plot is the median of 10 trials. The box-and-whisker plots on each data point show the max, min, and interquartile ranges of the runtimes. Similarly, Figure 3 shows the runtimes as we vary the dimension of the sphere.

Notice that the Rips computation slows down in higher dimensions because we insist Ripser compute simplices up to the ambient dimension of the data set as is computed for Alpha and DR filtrations. For example, for data points on a 3-sphere (in \mathbb{R}^4), Rips, Alpha, and DR compute $H_0, H_1, H_2,$ and H_3 classes. Particularly, for Rips, this increase in dimension exponentially increases the number of simplices and thus increases the size of the boundary matrix which causes computational slowdown.

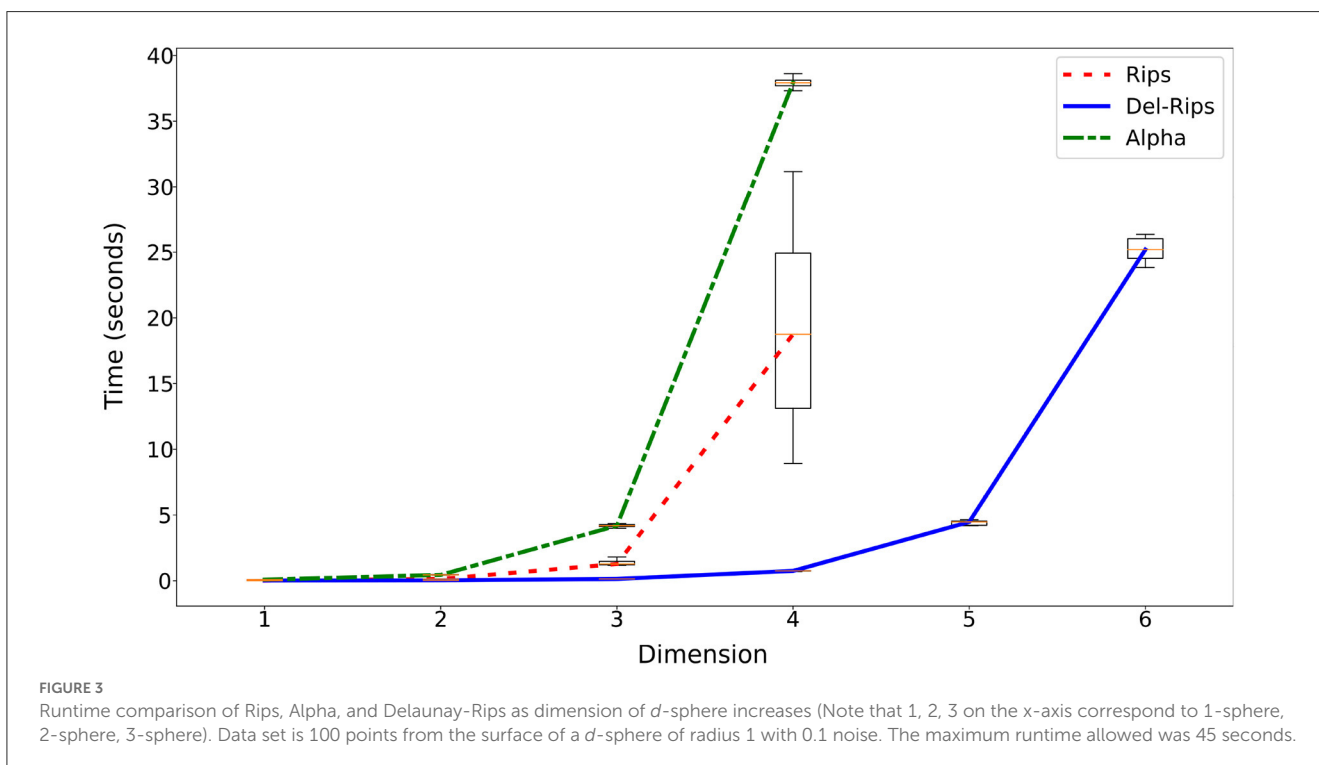
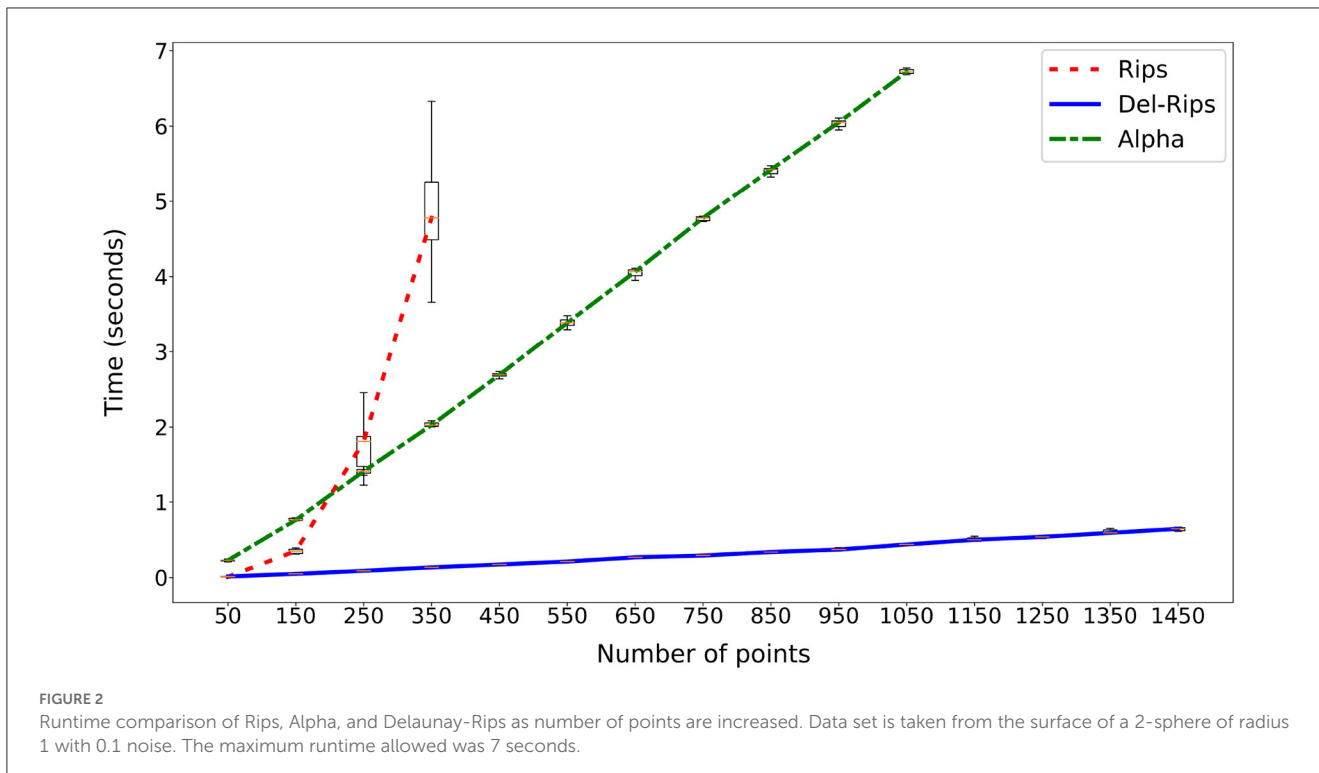
Notice that the Alpha computation slows down dramatically compared to DR even though both rely on computing the Delaunay triangulation. This is due to the computational overhead computing the scales at which multiple Voronoi cells intersect to assign simplex weights. Since DR avoids this computation, which is expensive in the Python implementation of Alpha in Cechmate, the practical speedup is significant. It is important to note that alternative implementations of Alpha [such as GUDHI [28]] may show speedups over the Python implementation of DR due to implementation details (e.g., the choice of implementation language), and while the filtered complexes built by Alpha and

Del-Rips are the same, we expect DR to enjoy some advantage in practice due to fast(er) calculation of simplex scales in optimized versions of these methods.

To further illustrate the increased severity of the impact on time complexity with increased data and (maximum homological) dimension, we compute persistence pairs using Rips, Alpha, and DR and document the runtimes in Figure 4. Figure 4A indicates the expected exponential increase of the Rips runtime as the number of points increases within a fixed dimension, with the degree of efficiency gained by DR increasing dramatically with increasing ambient and homological dimension. Comparing Alpha to DR in Figure 4B, we observe a roughly constant ratio of runtimes over the number of points (consistent with Figure 2) that decreases by a factor of approximately 1.9 and 1.6 when increasing the maximum homological dimension from 2 to 3 and 3 to 4, respectively.

3.3. Stability properties of the Delaunay-Rips complex

Toward understanding the impact of a perturbation of the underlying point cloud data has on the resulting DR persistence



diagram, we adopt the notion of an ε -perturbation of a point cloud from Weller [29]. If $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^D$ denotes a point cloud we let $P' = \{p'_1, \dots, p'_n\}$ denote a perturbation of P , where we imagine placing $\varepsilon < \left(\min_{p_i, p_j \in P} \{d(p_i, p_j)\}\right) / 2$ balls around

each point $p_i \in P$ and selecting from each ball a perturbed point $p'_i \in P'$.

Definition 3.2. We call $P' = \{p'_1, \dots, p'_n\}$ an ε -perturbation of $P = \{p_1, \dots, p_n\}$ if for each $p \in P$ there is exactly one $p' \in P'$ such

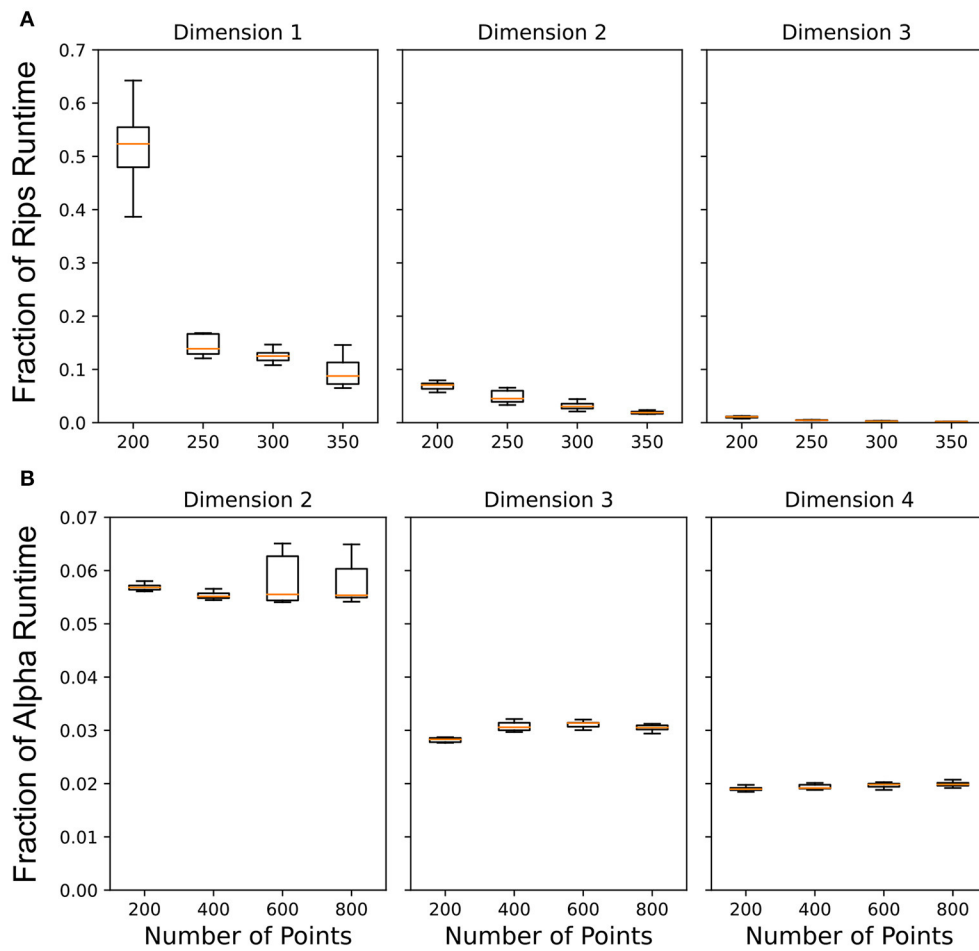


FIGURE 4 Boxplots of ratio of runtimes when computing persistence pairs using the Delaunay-Rips filtration to runtimes on the same data using (A) Rips and (B) Alpha filtrations for increasing number of points sampled from a noisy d -sphere. Each boxplot summarizes 10 instantiations with outliers removed. Data dimensions $(d + 1)$ and maximum homological dimension (d) increase from left to right.

that $d(p, p') < \varepsilon$ and, conversely, p is the only point in P within ε of p' . We refer to each p, p' as a perturbation pair.

In other words, there is a bijection $\rho : P \rightarrow P'$ associating to each point $p_i \in P$ a point $p'_i \in P'$, so that $d(p_i, \rho(p_i)) < \varepsilon$.

Lemma 3.3. Given that P' is an ε -perturbation of a $P \subset \mathbb{R}^D$, there exists a perturbation pair $p_i \in P$ and $p'_i \in P'$ such that

$$d_H(P, P') = d(p_i, p'_i).$$

Proof. First, identify the perturbation pair, $x \in P$ and $x' \in P'$ which are farthest from one another among all perturbation pairs and let

$$\delta := d(x, x') = \max_{p \in P, p' \in P'} (d(p, p')). \tag{1}$$

By such a choice, we ensure that for any $p' \in P'$, it must be that $p' \in P_\delta$ since $d(p, p') \leq d(x, x') = \delta$. Similarly, for any $p \in P$, $d(p, p') \leq d(x, x') = \delta$ and so $P \subseteq P'_\delta$. Therefore $d_H(P, P') \leq \delta$ by Definition 2.13.

Now, assume $\lambda < \delta$. For any $p' \in P'$

$$d(x, p') \geq \min_{p \in P} (d(x, p)) = d(x, x') = \delta > \lambda.$$

Therefore $x \notin P'_\lambda$ since it is not in any λ -ball around the points of P' , therefore $d_H(P, P') > \lambda$ for every $\lambda < \delta$. Thus, $d_H(P, P') = \delta = d(x, x')$. That is to say, the Hausdorff distance between P and P' is exactly equal to the largest distance a point in P was perturbed within an ε -perturbation.

To leverage Theorem 2.12 we will pursue conditions that ensure the underlying Delaunay triangulation does not change under a perturbation of the points. We say that a finite $P \subset \mathbb{R}^D$ and an ε -perturbation, P' have the same Delaunay triangulation with respect to the perturbation pairing if $\sigma = [p_{i_0} \dots p_{i_k}] \in \text{Del}(P)$ if and only if $\sigma' = [p'_{i_0} \dots p'_{i_k}] \in \text{Del}(P')$ for all $0 \leq k \leq D$, where p'_{i_i} is the unique point in P' uniquely associated to p_{i_i} , with $d(p_{i_i}, p'_{i_i}) < \varepsilon$. As abstract simplicial complexes, $\text{Del}(P)$ and $\text{Del}(P')$ are indistinguishable, since the association of perturbation pairs induces a bijection between simplices. Thus we will not distinguish the associated simplices in these complexes.

Theorem 3.4. Let $P \subset \mathbb{R}^D$ be a point cloud and P' an ε -perturbation of P with the same Delaunay triangulation with respect to the perturbation pairing; call it K . Let $f_P, f_{P'} : K \rightarrow \mathbb{R}$ be monotonic functions defined by assigning the Delaunay-Rips scales to the simplices of K as viewed in P and P' , respectively. For

each dimension p , the bottleneck distance between the persistence diagrams is bounded from above by twice the Hausdorff distance between the point clouds P and P' :

$$W_\infty \left(\text{Dgm}_p(f_P), \text{Dgm}_p(f_{P'}) \right) \leq 2d_H(P, P').$$

Proof. We will leverage the result of Theorem 2.12. Since P is a finite point cloud, there will be a simplex $\sigma \in K$ such that

$$\|f_P - f_{P'}\|_\infty = |f_P(\sigma) - f_{P'}(\sigma)|.$$

Since the Delaunay-Rips scale of a simplex is determined by the two points that are the farthest apart in the simplex, there exists $p, q \in P$ and $r', s' \in P'$, all of which are 0-simplices in σ , such that $f_P(\sigma) = d(p, q)$ and $f_{P'}(\sigma) = d(r', s')$. It could be that p, q and r', s' form two perturbation pairs (e.g., $p' = r'$ and $q' = s'$) but this need not hold in every case. For instance, the perturbation of the points p and q that determine the scale of σ in $\text{Del}(P)$ may have moved to p' and q' (both of which are necessarily in σ in $\text{Del}(P')$) closer together than the points r' and s' that determine the scale of σ in $\text{Del}(P')$. We consider these two cases separately.

Case 1: $p' = r'$ and $q' = s'$. Without loss of generality, assume

$$d(p, q) > d(p', q').$$

Then

$$\begin{aligned} |f_P(\sigma) - f_{P'}(\sigma)| &= |d(p, q) - d(r', s')| \\ &= |d(p, q) - d(p', q')| \\ &= d(p, q) - d(p', q') \\ &\leq d(p, p') + d(p', q') + d(q', q) - d(q', p') \\ &= d(p, p') + d(q', q) \\ &\leq d(x, x') + d(x, x') \\ &= 2d_H(P, P'), \end{aligned}$$

where $x \in P$ and $x' \in P'$ are the paired points farthest from one another in P and P' as in Lemma 3.3. Case 2: $p' \neq r'$ or $q' \neq s'$. As mentioned,

$$d(r', s') \geq d(p', q') \tag{2}$$

$$d(p, q) \geq d(r, s), \tag{3}$$

by assumption of which pair of vertices in σ determine the scales of in $\text{Del}(P)$ and $\text{Del}(P')$. Now, there are two sub-cases to consider.

Case 2a: $d(p, q) \geq d(r', s')$. Using Equation 2 we have

$$\begin{aligned} |f_P(\sigma) - f_{P'}(\sigma)| &= |d(p, q) - d(r', s')| = \\ &= d(p, q) - d(r', s') \leq d(p, q) - d(p', q') \end{aligned}$$

and the rest follows by the same argument as in Case 1.

Case 2b: $d(p, q) < d(r', s')$. Using Equation 3 we have

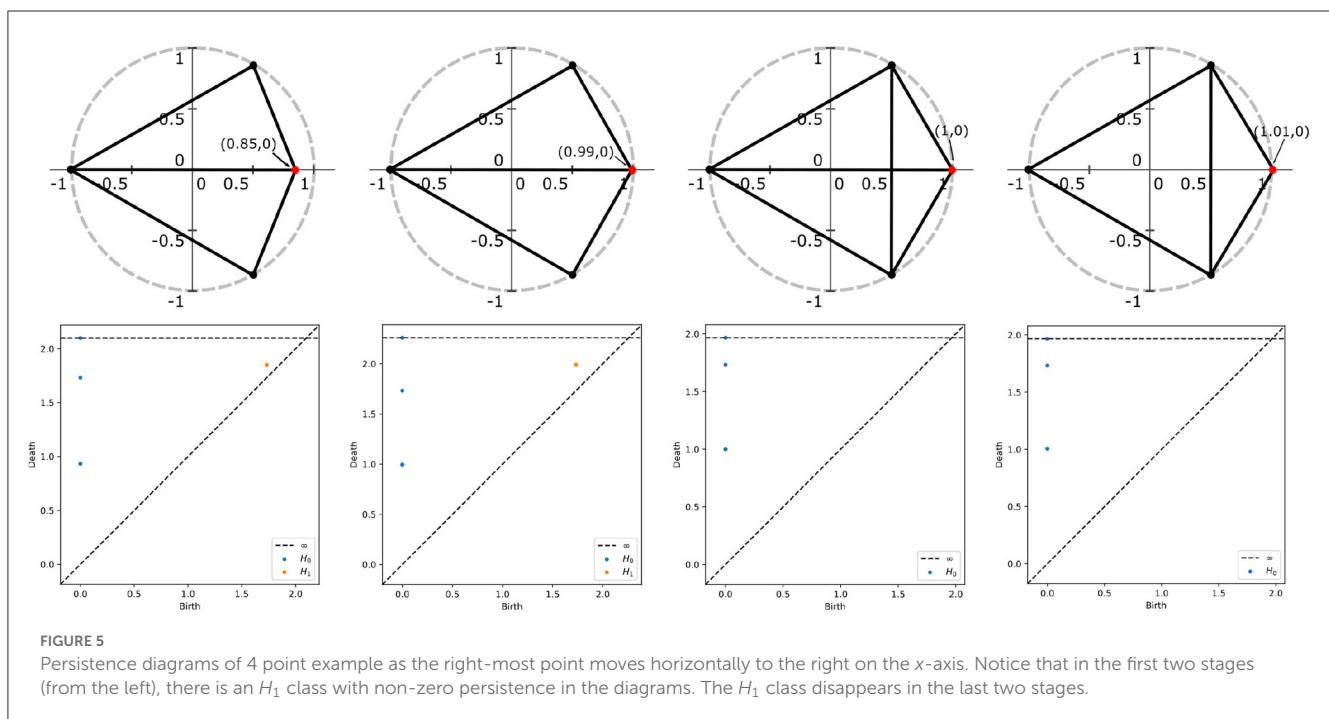
$$\begin{aligned} |f_P(\sigma) - f_{P'}(\sigma)| &= |d(p, q) - d(r', s')| = \\ &= d(r', s') - d(p, q) \leq d(r', s') - d(r, s) \end{aligned}$$

and the rest follows as in Case 1 with appropriate relabeling.

Finally, by applying Theorem 2.12, we conclude that the bottleneck distance between the persistence diagrams associated with P and P' , respectively, is bounded from above by twice the Hausdorff distance between the two point clouds:

$$W_\infty \left(\text{Dgm}_p(f_P), \text{Dgm}_p(f_{P'}) \right) \leq \|f_P - f_{P'}\|_\infty \leq 2d_H(P, P').$$

We see that when the underlying Delaunay triangulation on our point clouds is fixed (with respect to the perturbation pairing), assigning scales to simplices using the DR algorithm guarantees stability of the corresponding persistence diagram. Although it is known that for a point cloud,



P in general position, there exists a sufficiently small $\varepsilon > 0$ such that every ε -perturbation P' will have the same Delaunay triangulation [30], the size of ε may be quite small, which casts doubt on the practical utility of the above stability result for real applications in which there is measurement uncertainties. In the next section, we explore what can happen when the underlying Delaunay triangulation does change as a result of perturbing data points.

3.4. Persistence diagram instability

The DR construction gains computational efficiency at the cost of stability. We demonstrate a simple, yet clear example of how a discontinuity in the transformation from data to diagram can arise under a perturbation of the underlying data. In Figure 5, moving from left to right, we imagine moving the right-most point (in red) to the right toward the unique inscribing circle for the other three points. When the right-most point is inside the circle, there is an H_1 class with non-zero persistence. This class disappears immediately when the 4 points lie on the same circle. Informally, this means that DR sees the 4 point form a loop in the first two stages and then immediately loses sight of the loop when the points become cocircular. In the figure, we have marked the Delaunay triangulation of the points to showcase the position of the right-most (red) point at which an edge flip occurs (namely when all four points lie on the same circle). What we are seeing is point clouds that are very similar in structure visually, but are producing very different persistence diagrams. An arbitrarily small perturbation of the right-most point to inside the circle gives a very different persistence diagram from an arbitrarily small perturbation to outside the circle. We now proceed to formally prove the instability of the persistence diagrams associated with this particular configuration of points.

Lemma 3.5. Let $P' = \{(-1, 0), (\frac{1}{2}, \frac{\sqrt{3}}{2}), (\frac{1}{2}, -\frac{\sqrt{3}}{2}), (1-x, 0)\}$ with $0 < x < \delta < 2 - \sqrt{3}$. Using the Delaunay-Rips complex to construct a filtration on this point cloud, there is only one H_1 homology class with non-zero persistence.

Proof. We have $P' = \{(-1, 0), (\frac{1}{2}, \frac{\sqrt{3}}{2}), (\frac{1}{2}, -\frac{\sqrt{3}}{2}), (1-x, 0)\}$ with $0 < x < \delta < 2 - \sqrt{3}$. Our filtration has 4 key scale values, $t = 0 < \sqrt{1-x+x^2} < \sqrt{3} < 2 - \delta$ as shown in Figure 6.

We construct a boundary matrix B with entries from the field \mathbb{Z}_2 and reduce it to \bar{B} using the standard reduction algorithm found in Chapter VII of Edelsbrunner and Harer [20], which computes the pairing of simplices which respectively give rise to and kill off persistent homology classes in the H_0 and H_1 persistence diagrams. The ordering of the columns of B is determined by the ordering of the simplices given by the DR filtration, while ensuring each simplex appears after its faces.

$$\begin{matrix}
 & a & b & c & d & bd & cd & ab & ac & ad & abd & acd \\
 B = & \begin{pmatrix}
 a & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\
 b & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\
 c & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\
 d & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 \\
 bd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
 cd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
 ab & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
 ac & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
 ad & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
 abd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 acd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
 \end{pmatrix} \\
 & a & b & c & d & bd & cd & ab & ac & ad & abd & acd \\
 \bar{B} = & \begin{pmatrix}
 a & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
 b & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
 c & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
 d & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
 bd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
 cd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
 ab & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
 ac & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
 ad & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
 abd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 acd & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
 \end{pmatrix}
 \end{matrix}$$

By computing the scales of each simplex, the persistence pairs for the H_0 class with their persistence diagram coordinate (birth/death pair) are found to be

$$(a, N/A) : (0, \infty)$$

$$(b, ab) : (0, \sqrt{3})$$

$$(c, cd) : (0, \sqrt{1-x+x^2})$$

$$(d, bd) : (0, \sqrt{1-x+x^2}).$$

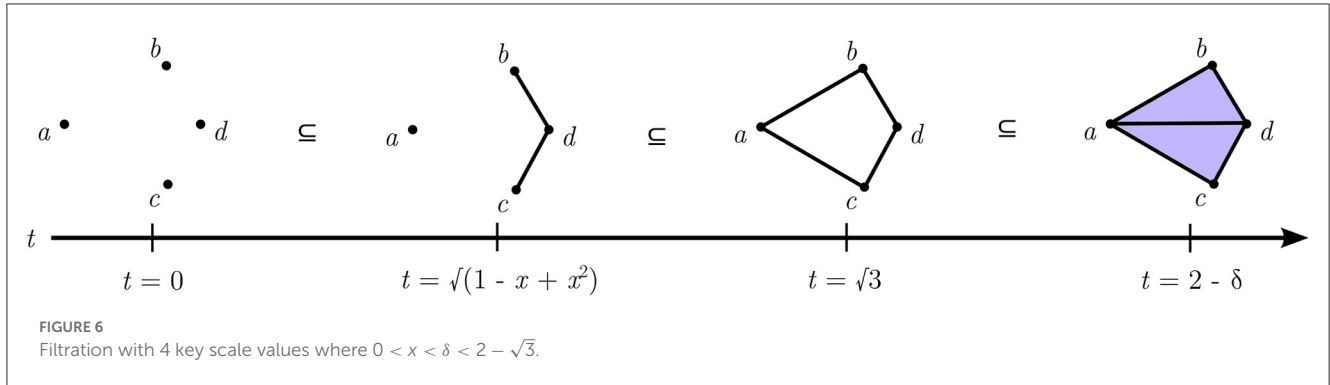
Likewise the H_1 persistence pairs are

$$(ad, abd) : (2-x, 2-x)$$

and

$$(ac, acd) : (\sqrt{3}, 2-x).$$

The only H_1 class with non-zero persistence is $(\sqrt{3}, 2-x)$.



Theorem 3.6. Let (\mathcal{P}, d_H) be the space of point clouds equipped with the Hausdorff metric and let (\mathcal{D}, W_∞) be the space of persistence diagrams equipped with the bottleneck metric. Let

$$\text{Pers}_1 : \mathcal{P} \rightarrow \mathcal{D}$$

where $\text{Pers}_1(P)$ is the persistence diagram of the H_1 classes of a point cloud $P \in \mathcal{P}$ constructed using the Delaunay-Rips complex. This map is discontinuous.

Proof. Let $P \in \mathcal{P}$ be $P = \{(-1, 0), (\frac{1}{2}, \frac{\sqrt{3}}{2}), (\frac{1}{2}, -\frac{\sqrt{3}}{2}), (1, 0)\}$. Note that the points all lie on the unit circle, so the Delaunay 1-skeleton has an edge between every pair of points (See the third frame in Figure 5). For this configuration of points, the vertical edge between the points $(\frac{1}{2}, \frac{\sqrt{3}}{2})$ and $(\frac{1}{2}, -\frac{\sqrt{3}}{2})$ appears at the exact same scale value as the cycle formed by all four points. Therefore in the DR filtration, the H_1 class whose boundary is the four outer edges dies at the same time as it is born.

Fix $\varepsilon = 0.1$. We now show that for any $\delta > 0$, there exists $P' \in \mathcal{P}$ such that $d_H(P, P') < \delta$, but $W_\infty(\text{Pers}_1(P), \text{Pers}_1(P')) \geq \varepsilon$. Take $P' = \{(-1, 0), (\frac{1}{2}, \frac{\sqrt{3}}{2}), (\frac{1}{2}, -\frac{\sqrt{3}}{2}), (1-x, 0)\}$ with $0 < x < \min\{\delta, \frac{2-\sqrt{3}}{2}\}$. This is a small perturbation of P gotten by pushing the point $(1, 0)$ inside the unit circle, thereby putting the points in general position (See the second frame in Figure 5 for an example). It is straightforward to compute the Hausdorff distance d_H between P and P' as

$$d_H(P, P') = x < \delta.$$

Recall that $\text{Pers}_1(P)$ has no H_1 class with non-zero persistence. Thus, to compute $W_\infty(\text{Pers}_1(P), \text{Pers}_1(P'))$, we must match the H_1 class of $\text{Pers}_1(P')$ with the diagonal. The H_1 class of $\text{Pers}_1(P')$ has birth $\sqrt{3}$ and death $2-x$ as calculated in Lemma 3.5. Using the max norm, we find

$$\begin{aligned} W_\infty(\text{Pers}_1(P), \text{Pers}_1(P')) &= (2-x) - (\sqrt{3}) \geq \\ &2 - \frac{2-\sqrt{3}}{2} - \sqrt{3} \geq 0.1 = \varepsilon. \end{aligned}$$

So Pers_1 is discontinuous at P .

As a remark, note that any metric on the space of point clouds that is bounded above by the Hausdorff distance will produce this discontinuity (for example, the Gromov-Hausdorff distance). This gives us insight into when the DR construction of the persistence

diagram may experience an instability—namely when points are not in general position.

Thus far, we have mathematically shown theoretical stability of the persistence diagram of a data set whose Delaunay triangulation does not change under the influence of a perturbation of the point locations. However, it should not be expected that in real-world applications the conditions guaranteeing stability will be met, and we have shown in this section that when the underlying Delaunay triangulation does change, the degree of change between diagrams may not be controlled by the degree of change in the underlying data. We are thus led to ask, to what extent does such an instability matter in practice?

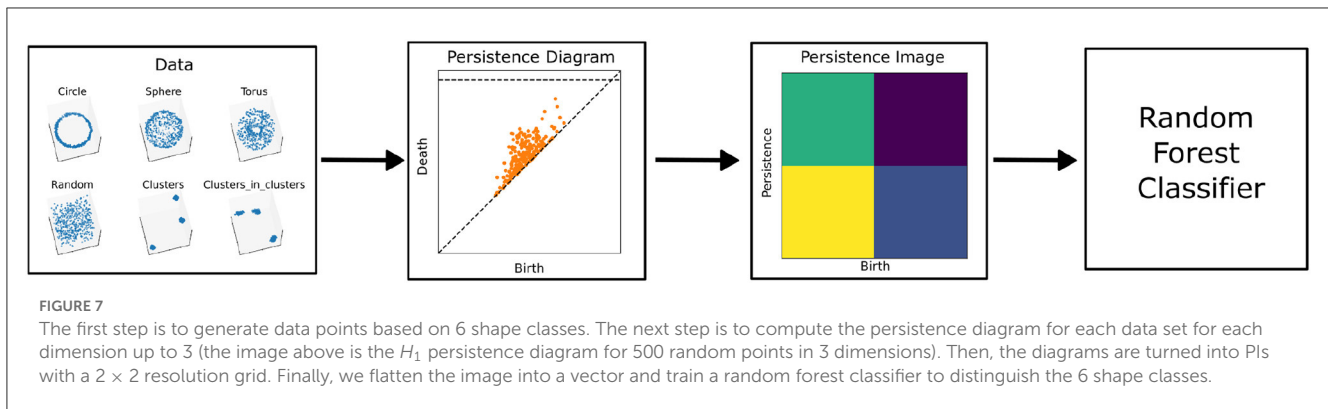
4. Machine learning model performance using Rips, Alpha, and Delaunay-Rips filtrations

Although we have special cases where instability in the PD may arise (as shown in Section 3.4), we show here that this instability may have little impact in applications. We demonstrate the robustness of DR for machine learning in a synthetic data context and a real data context. For the synthetic data, we work with random forest classifiers for a multi-class classification task. For the real data, we train support vector machines with linear kernel for a binary classification task.

4.1. Classification of synthetic shape data

To test the robustness of DR to changes in the degree of perturbation to the locations of points in point cloud data, we develop ML classification models using Rips, Alpha, and DR filtrations on point clouds generated by randomly sampling various manifold and adding random noise to perturb the points [31]. Although the persistence diagrams produced using the DR filtration enjoy stability when the underlying Delaunay triangulation is unchanged (see Section 3.3), in reality, the underlying Delaunay triangulation of the point is expected to change even for modest levels of noise.

Figure 7 shows the general pipeline for our experiment. We generated 100 point clouds consisting of 500 data points for each of 6 shape classes (circle, sphere, torus, random, three clusters, and



three clusters within three clusters) as subsets of \mathbb{R}^3 . For a fixed level of noise ν , points in each cloud were perturbed by randomly chosen vectors from the ambient space with maximum magnitude equal to ν . For each cloud, we computed the 0, 1, and 2-persistence diagrams using Rips, Alpha, and DR filtrations and then vectorized the resulting diagrams using PIs from the Python package “persim” in the scikit-tda library [32]. The resulting feature vectors were then used to train a random forest classifier using the implementation in scikit-learn [33]. To evaluate and compare the different filtrations, we computed the median classification accuracy of the trained models on held out data using 10-fold cross validation.

Fixing a homology class and fixing a noise level ν for the perturbation of the point cloud, we determine the birth and persistence ranges of persistence pairs produced over all filtration methods and over all samples, and we segment this region into 2×2 resolution persistence images. For example, we iterate through all H_1 persistence diagrams for all three of our filtration methods that were produced using $\nu = 0.20$. Then, we find the maximum birth range and persistence range and use those values to set a 2×2 resolution grid to produce the PIs. The purpose of doing this was to ensure that if we compared pixels of the PIs corresponding to different filtration methods, we would make a fair comparison because corresponding persistence pairs would land in corresponding pixels (we leverage this design when assessing feature importance).

Figure 8 shows the median accuracy of our model for increasing noise levels (we plotted box-whisker plots to show the spread of the accuracy from 10-fold cross validation). As a baseline comparison, note that if our ML model was randomly classifying the test data, we expect to see accuracy of $1/6 \approx 17\%$. Since we are seeing over 70% median accuracy for each noise level, our model truly is finding distinguishing features between the 6 shape classes.

Notably, as is evident in Figure 8, the degradation of ML model performance with increasing noise levels is very similar between DR, Alpha, and Rips filtrations. We do note that the median classification accuracy using Alpha filtrations to generate persistence diagrams is slightly better than the other methods across all noise levels. However, except for a single noise level (0.15), all median accuracies are within the ranges of the other two methods.

In addition to model accuracy, an important consideration is which features are being used to achieve said accuracy. Identifying salient features provides a level of model interpretability [34], and

may guide the modeler to methods to reduce the dimension of the input, which may further improve model performance. To further compare the three filtration methods, we compared the most important topological features learned during training as determined by a random forest classifier. We began with the PDs produced from the $\nu = 0.20$ noisy data. We generated PIs with the following resolutions for each persistence diagram:

- H_0 diagram resolution: 5×1
- H_1 diagram resolution: 5×5
- H_2 diagram resolution: 5×5

As a result, we produced 55-dimensional feature vectors for each of the 600 samples (100 samples each of 6 classes). We trained a random forest classifier with the same parameters as before, once using a train-test split of 70–30, and used the built-in assessment [33] of the Gini importance to quantify feature importance. The Gini importance of a feature is calculated as the amount of reduction to the Gini index [35] brought by that feature. Thus, the higher the Gini importance, the more important the feature is for making classification decisions. Figure 9 shows heatmaps indicating feature importance. Notice how for the different filtrations we used to obtain the PDs, the PIs generated have similar corresponding pixels that the ML model found important. For example, among the $5+25+25$ pixel feature space for the DR-based random forest classifier, the most important features were the bottom-left (dark blue) H_2 pixel with Gini importance 0.19 and the bottom H_0 pixel with Gini importance 0.17. These were in similar regions to the most important features for Alpha and Rips based classifiers. Hence, we have reason to believe that regardless of the filtration method used, the ML model learned the importance of similar features which are distinguishing between the shapes.

4.2. Classification of sleep state

We next investigate the applicability of DR to a classification problem involving biophysical data. The goal is to showcase the comparable effectiveness of DR with Alpha and Rips filtrations in terms of model performance metrics on a problem in which computational efficiency may be a relevant constraint. Our application is a reanalysis of the data and methodology employed in Chung et al. [19], in which the authors develop an ML

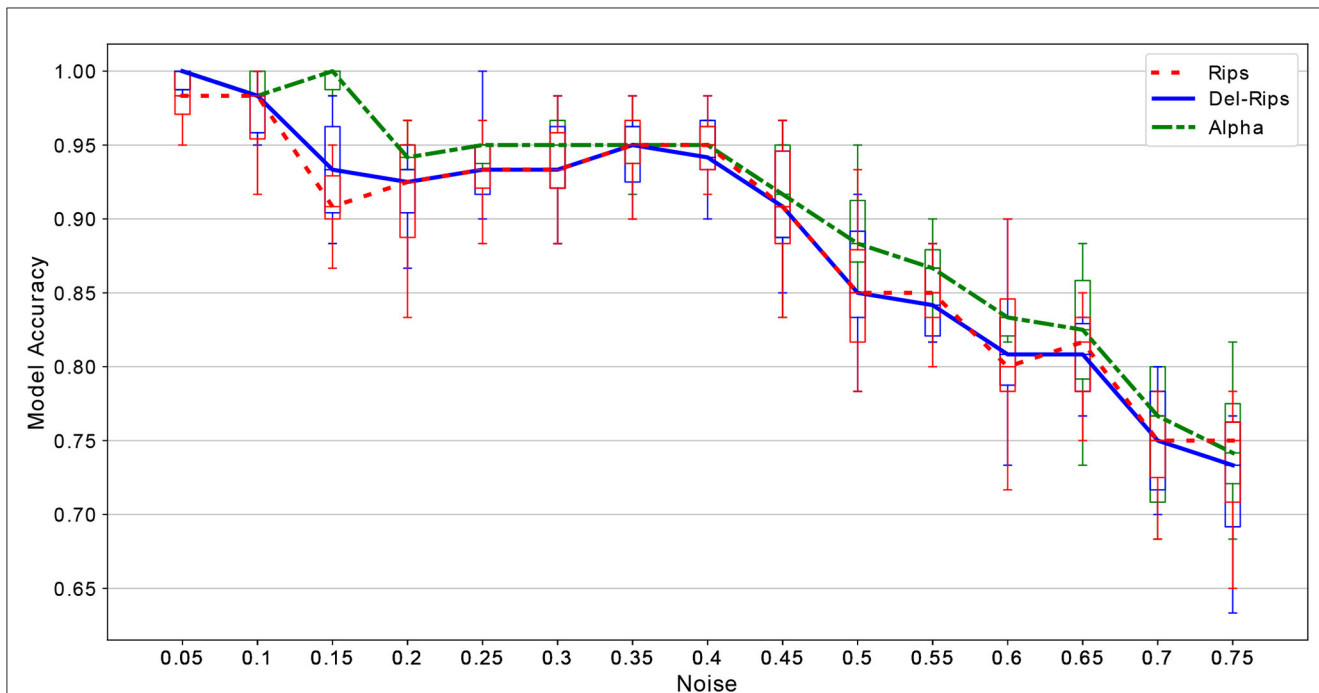


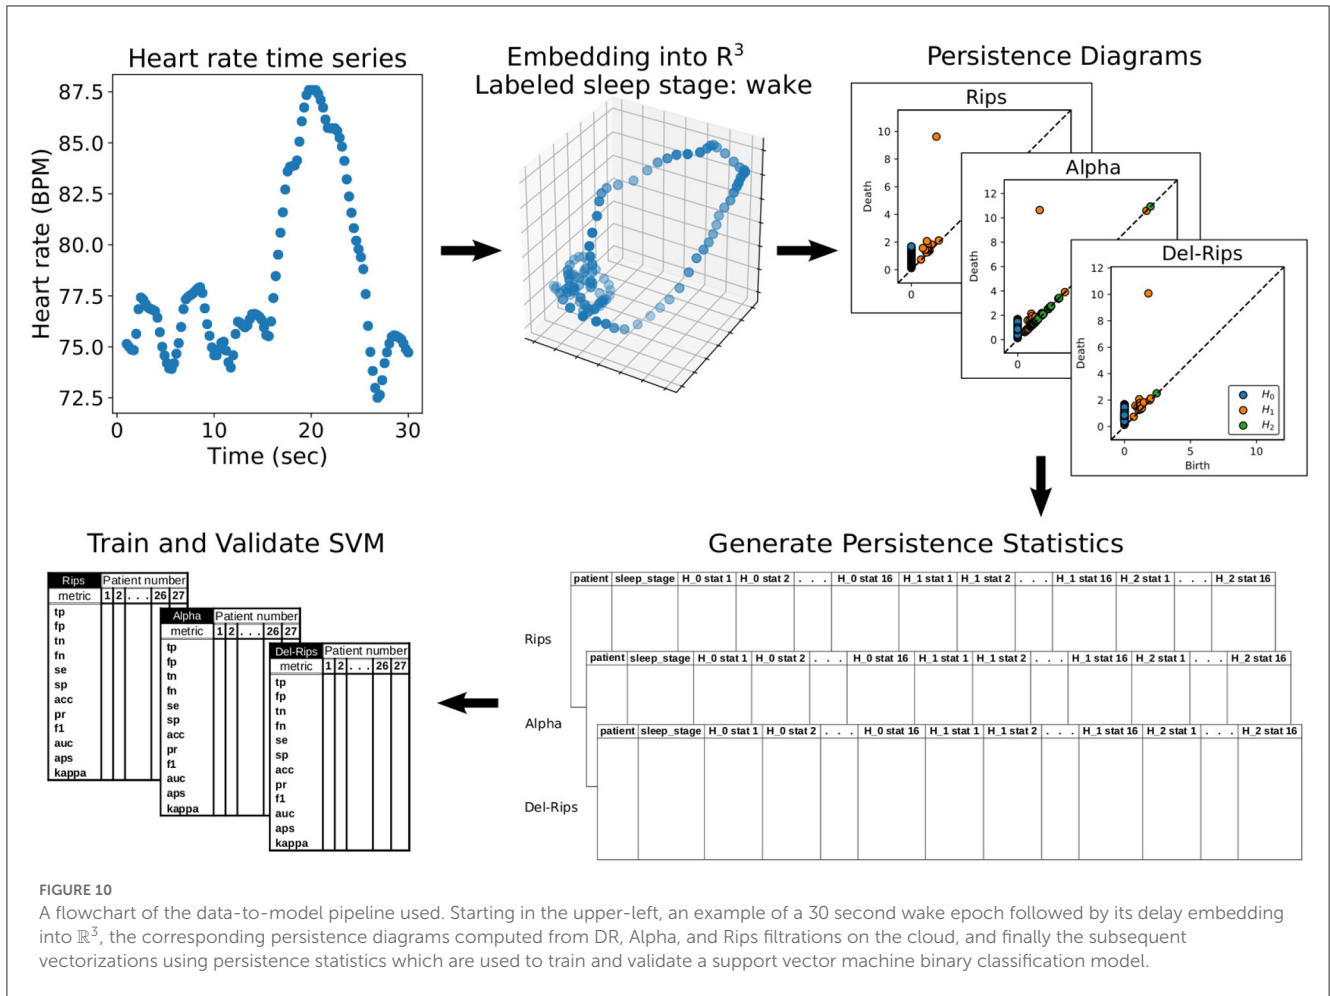
FIGURE 8
A plot of the median random forest classification accuracy computed on Rips, DR, and Alpha based persistence diagrams across varying noise levels for the original point clouds. The box-and-whisker plots for each noise level show the range of the accuracy across 10-fold cross validation.



FIGURE 9
Heatmaps of H_0 , H_1 , and H_2 persistence diagram PI vectorizations. We took the 55 features for training the random forest classifier based on each filtration function and found their importance in their corresponding classifier models. The number in each pixel is the Gini importance of that pixel in our random forest's classification decisions. The higher the number, the more important the pixel.

model to classify sleep stages of a participant using observed instantaneous heart rate (IHR) time series collected using ECG. A high-level visualization of the data-to-model development pipeline

we implement [36] is provided in Figure 10. As a note, the authors of this manuscript (that investigates the properties and usage of the DR filtration) did not work directly with the data; instead, the



processed data was obtained with permission from the authors of Chung et al. [19]. For details on the original study that produced the original data, the study participants, data collection protocols, board approvals, and lab equipment see Malik et al. [37].

The authors of Chung et al. [19] trained a machine learning model to predict sleep state, trained on statistics of topological features of high dimensional point clouds that were built from delay embeddings of IHR time series. This approach is premised on the idea that delay embeddings of observed time series can recover dynamical features of underlying attractors in an unknown state space as implied by the well known theorem of Takens [38]. The topological/geometric features of these attractors (as represented by PDs) may be discriminating between dynamic states. Applications based on similar rationale have also been found in contexts such as 3D motion capture data [39].

The data on which models were trained and tested comprised 90 ECG recordings from Chang Gung Memorial Hospital (CGMH) as the CGMH-training database. Each recording was sampled at 200 Hz and each 30s epoch was manually annotated into either “awake” or “asleep” states. Sleep states were further categorized into one of 5 different sleep stages: stage 1 through 4 or REM sleep, giving a potential for a 6-class classification problem. To each recording a standard R-peak detection algorithm with a 5-beat median filter to remove artifacts was applied [similarly done

in Xu and Schuckers [40]]. From this preprocessed data, IHR time series were computed at 4 Hz. For efficiency, we performed delay embeddings of the IHR time series into \mathbb{R}^3 , choosing a delay parameter between consecutive coordinates of 5 (i.e., our delay vectors consisted of points $(x_i, x_{i+5}, x_{i+10}) \in \mathbb{R}^3$, taking every 5th point in the time series as the coordinate for our embedded points.) The number of indices between the starting coordinate of consecutive delay vectors (i.e., the stride) was set to 1 so as to include all timepoints in at least one delay vector.

After embedding each epoch into \mathbb{R}^3 , we computed H_0, H_1 , and H_2 persistence diagrams using the DR, Rips, and Alpha filtrations of the embedded point clouds. Following Chung et al. [19], the resulting PDs were converted to fixed-length feature vectors by computing sample statistics of the persistence pairs in the diagrams. In particular, for each homological dimension $p = 0, 1, 2$, we construct two sets: M_p , the set of means of all persistence pairs and L_p , the persistence of each pair (the death minus the birth). For each set, we calculate the mean, standard deviation, skew, kurtosis, 25th percentile, 50th percentile, 75th percentile, and persistent entropy [41]. The result was 16 persistence statistics corresponding to each dimension- p persistence diagram. NaN was assigned to empty H_p diagrams and these samples were later dropped from the training set. In total, from each sample a total of 48 persistence statistics (16 for each dimension $p = 0, 1, 2$) were computed,

TABLE 1 (a) The median and interquartile ranges across 27 validation participants of each classification model performance metric obtained from 3 classifiers trained either on topological features determined by DR, Rips, or Alpha filtrations.

(a)	Rips		Alpha		Del-Rips	
	median	iqr	median	iqr	median	iqr
se	0.565217	0.428794	0.569892	0.380411	0.555556	0.433198
sp	0.925065	0.271714	0.917511	0.270439	0.917271	0.272863
acc	0.852941	0.161052	0.848276	0.154095	0.852573	0.161564
pr	0.439636	0.316970	0.448718	0.344648	0.424307	0.301726
f1	0.517986	0.140499	0.489655	0.136024	0.472727	0.170183
auc	0.866199	0.083241	0.864330	0.079358	0.866646	0.096021
aps	0.594109	0.180105	0.555916	0.160430	0.552859	0.159507
kappa	0.393895	0.178491	0.368380	0.197346	0.356587	0.192436
(b)	<i>p</i> -value for Rips vs. Del-Rips			<i>p</i> -value for Alpha vs. Del-Rips		
se	1.000000			1.0		
sp	1.000000			1.0		
acc	1.000000			1.0		
pr	1.000000			1.0		
f1	0.276303			1.0		
auc	1.000000			1.0		
aps	0.102470			1.0		
kappa	1.000000			1.0		

(b) The Mood's median test *p*-values after comparing DR model performance metrics against Rips and Alpha model performance metrics.

resulting in a 48-dimensional feature space on which an SVM model was trained.

A total of 67,188 48-dimensional topological feature vectors (one for each 30-second epoch) were used to train a support vector machine with a linear kernel and balanced class weights to account for imbalance in the numbers of sleep and wake samples. In total 3 models were trained, one for each filtration method. The trained models were each validated on 27 held-out participants in the CGMH-validation by computing sensitivity (*se*), specificity (*sp*), accuracy (*acc*), precision (*pr*), F1 score (*f1*), AUC-score (*auc*), average precision score (*aps*), and kappa coefficient (*kappa*) for the epochs of each validation participant. To assess the statistical significance of any differences in performance metrics between DR and either Rips or Alpha filtrations, we applied Mood's median test from SciPy [42] to the distributions of each performance metric across the 27 validation participants, comparing separately Rips to DR and Alpha to DR filtrations. In this context, Mood's median test tests the null hypothesis that the median of each performance metric is the same.

Our model construction differs from the one in the original paper in several key ways. First, the authors of Chung et al. [19] derived topological features from considerably high-dimensional (120) delay embeddings of the IHR time series. Such high-dimensional embeddings pose a challenge for Alpha and DR filtrations since they both require computing Delaunay triangulations, which do not scale efficiently to high dimensions. Furthermore, the original model was trained on statistics derived only from H_0 and H_1 diagrams, while we included H_2 diagrams as well. We note that these changes come at a modest reduction in

model performance compared to what is reported in Chung et al. [19], although our goal was to assess differences between filtrations and not to improve over previously published classification models.

The median and interquartile ranges of each model performance metric across the validation participants as well as the *p*-values of Mood's median test are shown in Table 1. We observe very similar performance metrics for all three filtration methods, with either DR, Alpha, or Rips exhibiting the best median performance, depending on the metric. For Mood's median test, a small *p*-value indicates that there is a notable difference between the medians of the corresponding performance metric between the two classification models. Notice that the *p*-values in Table 1 are all very large, with the exception of the average precision score (*aps*) when comparing Rips vs. DR. This suggests that, given the spread of model performance on different validation participants, we do not have enough evidence to reject the null hypothesis that the median performance metrics are the same.

5. Conclusion

In this paper, we have defined and implemented the Delaunay-Rips (DR) filtration for point cloud data, compared its computational efficiency in practice to other standard filtrations built on point cloud data, characterized some of its stability properties, and demonstrated its application on various datasets in an ML pipeline. We proved that under sufficiently small perturbations of the locations of points in the cloud, the persistence diagrams vary continuously with the data, provided the underlying

simplicial complex on the data set remains fixed. This was done by bounding from above the bottleneck distance between the persistence diagrams by twice the Hausdorff distance between the original data and its perturbed counterpart. In the case when the Delaunay triangulation changes as a result of a larger perturbation, we examined the instability of the DR filtration by carefully proving a discontinuity of the map between the data metric space and the diagram metric space. As far as we know, this result is a first of its kind to use the standard reduction algorithm to compute persistence diagrams on symbolic variables in service of a formal proof. Since the theoretical condition on stability may not always be met in practical applications, we investigated whether the instability in the diagrams generated using DR poses a significant problem when used in an ML pipeline and found that it needs not.

There are several limitations of this study and avenues of further inquiry. For one, we expect the relative performance of each filtration method used to derive topological features for an ML modeling task to be problem specific. Thus we cannot be certain the insensitivity of model performance to filtration method, and the comparable performance of DR to Rips and Alpha we observed will generalize to other contexts.

Our stability results were simplified by insisting we maintain the underlying Delaunay triangulation to keep the space of simplices fixed. While our results in Section 3.4 indicate that, in general, a change in the underlying Delaunay triangulation can cause a discontinuity in the transformation sending data to diagram, a more precise characterization of the degree of the discontinuities is not known, and so we are limited to an empirical evaluation.

In Section 3.2, we compared a Python implementation of DR with other filtration methods. Implementation details including the choice of language may have implications for the relative performance gains of DR over other filtration methods. With a C/C++ implementation of DR, how does the runtime of computing persistence diagrams based on DR compare with runtimes of computing persistence diagrams using other complexes (Cech, Witness, etc.)? Along the same lines, how does our implementation (or a C/C++ implementation) of DR compare in runtime with implementations of other TDA methods in other software packages (e.g., Javaplex [43], Perseus [44], Dionysus [45], DIPA [46], Gudhi [47])? Finally, during the drafting of this paper, an updated version of the software package Ripser was released named Ripser++ [48]. The new approach makes use of GPU acceleration to parallelize processing simplices and finding persistence pairs. How does the runtime of DR for persistent diagram computation compare with that of Ripser++? Is there a way to harness parallelization to computationally benefit DR? These comparisons are outside the scope of this work.

References

- Leibon G, Pauls S, Rockmore D, Savell R. Topological structures in the equities market network. *Proc Nat Acad Sci.* (2008) 105:20589–94. doi: 10.1073/pnas.0802806106
- Chung MK, Bubenik P, Kim PT. Persistence diagrams of cortical surface data. In: *International Conference on Information Processing in Medical Imaging*. Springer (2009). p. 386–397. doi: 10.1007/978-3-642-02498-6_32
- Qaiser T, Sirinukunwattana K, Nakane K, Tsang YW, Epstein D, Rajpoot N. Persistent homology for fast tumor segmentation in whole slide histology images. *Procedia Comput Sci.* (2016) 90:119–24. doi: 10.1016/j.procs.2016.07.033
- Asaad A, Jassim S. Topological Data Analysis for Image Tampering Detection. In: Kraetzer C, Shi YQ, Dittmann J, Kim HJ, editors. *Digital Forensics and*

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

FM contributed the conceptualization of the ideas of the paper, formal analysis, methodology, project administration, computing resources, supervision, validation of results, writing of the original draft, and review and editing of the final draft. AM contributed the conceptualization, data curation, formal analysis, investigation, methodology, software design, validation of results, creation of visualizations, writing of the original draft, and review and editing of the final draft. FM and AM adopt the Contributor Roles Taxonomy (CRediT, <https://credit.niso.org/>) for author attributions. All authors contributed to the article and approved the submitted version.

Acknowledgments

The contents of this manuscript have appeared online in a preprint at arXiv:2303.01501 [49]. The authors would like to acknowledge and thank Yu-Min Chung and Hau-Tieng Wu for their help in obtaining the processed ECG heart rate data. The authors would like to express their sincere gratitude to the reviewers for their valuable feedback, which enhanced the quality and clarity of this manuscript.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Watermarking. Cham: Springer International Publishing (2017). p. 136–46. doi: 10.1007/978-3-319-64185-0_11
5. Kramár M, Goullet A, Kondic L, Mischaikow K. Quantifying force networks in particulate systems. *Physica D*. (2014) 283:37–55. doi: 10.1016/j.physd.2014.05.009
 6. Edelsbrunner H. The union of balls and its dual shape. *Discr Comput Geom*. (1993) 13:415–40. doi: 10.1007/BF02574053
 7. Sheehy DR. Linear-size approximations to the Vietoris-Rips filtration. In: *Proceedings of the twenty-eighth annual symposium on Computational geometry*. (2012). p. 239–248. doi: 10.1145/2261250.2261286
 8. Guibas LJ, Oudot S. Reconstruction using witness complexes. *Discr Comput Geom*. (2007) 40:325–56. doi: 10.1007/s00454-008-9094-6
 9. de Silva V, Carlsson GE. Topological estimation using witness complexes. In: *Symposium on Point Based Graphics*. (2004).
 10. Bauer U. Ripser: efficient computation of Vietoris-Rips persistence barcodes. *J Appl Comput Topol*. (2021) 5:391–423. doi: 10.1007/s41468-021-00071-5
 11. Tralie C, Saul N, Bar-On R. Ripserpy: A lean persistent homology library for python. *J Open Source Softw*. (2018) 3:925. doi: 10.21105/joss.00925
 12. Cufar M. Ripserer jl: flexible and efficient persistent homology computation in Julia. *J Open Source Software*. (2020) 5:2614. doi: 10.21105/joss.02614
 13. Silva VD. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*. (2008) 135:39–64. doi: 10.1007/s10711-008-9261-1
 14. Chazal F, Cohen-Steiner D, Guibas LJ, Mamoli F, Oudot SY. Gromov-hausdorff stable signatures for shapes using persistence. *Comput Graph Forum*. (2009) 28:1393–403. doi: 10.1111/j.1467-8659.2009.01516.x
 15. Cohen-Steiner D, Edelsbrunner H, Harer J. Stability of persistence diagrams. *Discr Comput Geom*. (2005) 37:103–20. doi: 10.1007/s00454-006-1276-5
 16. Skraba P, Turner K. Wasserstein stability for persistence diagrams. *arXiv preprint arXiv:2006.16824*. (2020).
 17. Tralie C, Saul N. Cechmate. (2021). Available online at: <https://github.com/scikit-tda/cechmate> (accessed June 26, 2023).
 18. Adams H, Emerson T, Kirby M, Neville R, Peterson C, Shipman P, et al. Persistence images: A stable vector representation of persistent homology. *J Mach Learn Res*. (2017) 18:1–35.
 19. Chung YM, Hu CS, Lo YL, Wu HT. A persistent homology approach to heart rate variability analysis with an application to sleep-wake classification. *Front Physiol*. (2021) 12:637684. doi: 10.3389/fphys.2021.637684
 20. Edelsbrunner H, Harer J. *Computational Topology: An Introduction*. New York: American Mathematical Society. (2010). doi: 10.1090/mbk/069
 21. Hatcher A. *Algebraic Topology*. Cambridge: Cambridge University Press. (2002).
 22. Otter N, Porter MA, Tillmann U, Grindrod P, Harrington HA. A roadmap for the computation of persistent homology. *EPJ Data Sci*. (2017) 6:1–38. doi: 10.1140/epjds/s13688-017-0109-5
 23. Matouek J, LC reductions yield isomorphic simplicial complexes. *Contrib Discr Mathem*. (2008) 1:3. doi: 10.11575/cdm.v3i2.61933
 24. Adamaszek M, Adams H, Motta F. Random cyclic dynamical systems. *Adv Appl Math*. (2017) 83:1–23. doi: 10.1016/j.aam.2016.08.007
 25. Bauer U, Edelsbrunner H. The Morse theory of Tech and Delaunay complexes. *Trans Am Mathem Soc*. (2016) 369:3741–62. doi: 10.1090/tran/6991
 26. Mishra A. *Delaunay-Rips*. (2022). Available online at: <https://github.com/amish-mishra/cechmate-DR> (accessed June 26, 2023).
 27. Bauer U, Kerber M, Reininghaus J, Wagner H. Phat persistent homology algorithms toolbox. *J Symb Comput*. (2017) 78:76–90. doi: 10.1016/j.jsc.2016.03.008
 28. Rouvreau V. Alpha complex. In: *GUDHI User and Reference Manual 3rd ed GUDHI Editorial Board*. (2023).
 29. Weller FU. Stability of voronoi neighborhood under perturbations of the sites. In: *CCCG*. (1997).
 30. Boissonnat JD, Dyer R, Ghosh A. The stability of Delaunay triangulations. *Int J Comput Geom Appl*. (2013) 23:303–33. doi: 10.1142/S0218195913600078
 31. Mishra A. *Classification of synthetic data into shape classes*. (2022). Available online at: <https://github.com/amish-mishra/TDA-shape-classification-using-DR> (accessed June 26, 2023).
 32. Saul N, Tralie C, Motta F, Catanzaro M, Angeloro G, Sheagren C. *Persim*. (2021). Available online at: <https://persim.scikit-tda.org/en/latest/> (accessed June 26, 2023).
 33. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. *J Mach Learn Res*. (2011) 12:2825–30.
 34. Ribeiro MT, Singh S, Guestrin C. “Why should i trust you?” Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. (2016). p. 1135–1144. doi: 10.1145/2939672.2939778
 35. Krzywinski M, Altman N. Points of Significance: Classification and regression trees. *Nat Methods*. (2017) 14:757–8. doi: 10.1038/nmeth.4370
 36. Mishra A. *ML-Del-Rips-sleep-wake-classification*. (2022). Available online at: <https://github.com/amish-mishra/ML-Del-Rips-sleep-wake-classification> (accessed June 26, 2023).
 37. Malik J, Lo YL, Wu HT. Sleep-wake classification via quantifying heart rate variability by convolutional neural network. *Physiol Measur*. (2018) 39:085004. doi: 10.1088/1361-6579/aad5a9
 38. Takens F. Detecting strange attractors in turbulence. In: Rand D, Young LS, editors. *Dynamical Systems and Turbulence, Warwick 1980*. Berlin, Heidelberg: Springer Berlin Heidelberg (1981). p. 366–81. doi: 10.1007/BFb0091924
 39. Venkataraman V, Ramamurthy KN, Turaga P. Persistent homology of attractors for action recognition. In: *2016 IEEE International Conference on Image Processing (ICIP)*. (2016). p. 4150–4154. doi: 10.1109/ICIP.2016.7533141
 40. Xu X, Schuckers S. Automatic detection of artifacts in heart period data. *J Electrocardiol*. (2001) 34:205–10. doi: 10.1054/jelc.2001.28876
 41. Chintakunta H, Gentimis T, Gonzalez-Diaz R, Jimenez MJ, Krim H. An entropy-based persistence barcode. *Patt Recognit*. (2015) 48:391–401. doi: 10.1016/j.patcog.2014.06.023
 42. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 10: fundamental algorithms for scientific computing in python. *Nature Methods*. (2020) 17:261–72. doi: 10.1038/s41592-019-0686-2
 43. Tausz A, Vejdemo-Johansson M, Adams H. JavaPlex: A research software package for persistent (co)homology. In: Hong H, Yap C, editors. *Proceedings of ICMS 2014 Lecture Notes in Computer Science*. (2014). p. 129–136. Available online at: <http://appliedtopology.github.io/javaplex/> (accessed June 26, 2023). doi: 10.1007/978-3-662-44199-2_23
 44. Nanda V. *Perseus, the Persistent Homology Software*. (2013). Available online at: <http://www.sasupennedu.vnanda/perseus> (accessed March 1, 2023)
 45. Morozov D. *Dionysus*. (2023). Available online at: <https://pypi.org/project/dionysus/> (accessed June 26, 2023).
 46. Reininghaus J. *DIPHA (A Distributed Persistent Homology Algorithm)*. (2017). Available online at: <https://github.com/DIPHA/dipha> (accessed June 26, 2023).
 47. Maria C, Boissonnat JD, Glisse M, Yvinec M. The gudhi library: Simplicial complexes and persistent homology. In: *Mathematical Software ICMS 2014: 4th International Congress, Seoul, South Korea*. Berlin Heidelberg: Springer (2014). p. 167–174. doi: 10.1007/978-3-662-44199-2_28
 48. Simon Zhang MX, Wang H. GPU-accelerated computation of Vietoris-Rips persistence barcodes. *arXiv preprint arXiv:2003.07989* (2020).
 49. Mishra A, Motta FC. Stability and machine learning applications of persistent homology using the Delaunay-Rips complex. *arXiv preprint arXiv:2303.01501* (2023).