



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/191250/>

Version: Published Version

Article:

Gowdrige, T., Dervilis, N. and Worden, K. (2022) On topological data analysis for structural dynamics: an introduction to persistent homology. ASME Open Journal of Engineering, 1. 011038. ISSN: 2770-3495

<https://doi.org/10.1115/1.4055184>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

On Topological Data Analysis for Structural Dynamics: An Introduction to Persistent Homology

T. Gowdrige

Dynamics Research Group,
Department of Mechanical Engineering,
University of Sheffield,
Mappin Street,
Sheffield S1 3JD, UK
e-mail: tgowdrige1@sheffield.ac.uk

N. Dervilis

Dynamics Research Group,
Department of Mechanical Engineering,
University of Sheffield,
Mappin Street,
Sheffield S1 3JD, UK
e-mail: N.Dervilis@sheffield.ac.uk

K. Worden

Dynamics Research Group,
Department of Mechanical Engineering,
University of Sheffield,
Mappin Street,
Sheffield S1 3JD, UK
e-mail: k.worden@sheffield.ac.uk

Topological methods can provide a way of proposing new metrics and methods of scrutinizing data, that otherwise may be overlooked. A method of quantifying the shape of data, via a topic called topological data analysis (TDA) will be introduced. The main tool of TDA is persistent homology. Persistent homology is a method of quantifying the shape of data over a range of length scales. The required background and a method of computing persistent homology are briefly discussed in this work. Ideas from topological data analysis are then used for nonlinear dynamics to analyze some common attractors, by calculating their embedding dimension, and then to assess their general topologies. A method will also be proposed, that uses topological data analysis to determine the optimal delay for a time-delay embedding. TDA will also be applied to a Z24 bridge case study in structural health monitoring, where it will be used to scrutinize different data partitions, classified by the conditions at which the data were collected. A metric, from topological data analysis, is used to compare data between the partitions. The results presented demonstrate that the presence of damage alters the manifold shape more significantly than the effects present from temperature. [DOI: 10.1115/1.4055184]

Keywords: chaotic system, data analytics, nonlinear dynamical systems, nonlinear systems, structural dynamics, structural health monitoring, structures

1 Introduction

Topological methods are very rarely used in structural dynamics generally, although considering the structure and topology of observed data may be useful. Topological methods can provide new metrics and methods of scrutinizing data; the most rudimentary and powerful of which, *persistence homology*, will be discussed and used extensively in this paper. By applying topological methods, an understanding of the topological structure of data can be used to formulate arguments and develop an understanding of system parameters on a manifold shape. This statement is especially true when working with higher-dimensional data sets, where an intuitive understanding of a manifold is not easy to visualize. By using topological methods, an understanding can be quantifiably analyzed. When considering engineering data, data sets are often embedded in higher-dimensional space, and topological information is often not well explored, leaving out potentially important and insightful information.

Topological data analysis (TDA) is a recently-developed and fast-growing field which has found its way into many areas of science and engineering. The general idea of TDA applies concepts from algebraic topology to data sets. The primary focus of TDA is to determine the shape of the manifold in which sampled data are embedded. This process is achieved by identifying 2D holes, 3D cavities, and higher-dimensional analogs within the data structures. From these sampled data, an approximation to the topological structure can be calculated by using *simplicial complexes*, which are higher-dimensional analogs of graphs. From the simplicial complexes, the persistent homology can be calculated, this is then used to understand the topological structure of the data. The persistent homology is invariant for each data set, and can be used to identify the data set; just like a fingerprint.

After the simplicial complexes have been constructed for some point data, these can be manipulated using ideas from algebraic

topology, to determine algebraic groups that will capture information about the shape and structure of the simplicial complex; these groups are topological invariants. For any simplicial complex, algebraic topology can deduce a property called *homology*, which encodes information about its number of k -dimensional holes. A generalization of homology will be discussed with respect to pioneering work by Edelsbrunner and coauthors [1,2], where the homology can be considered over a range of simplicial complexes; this is called as *persistent homology*, aptly named, as this uncovers how the homology persists over a range of scales.

The field of algebraic topology is well-studied, but the application of the ideas to discrete point clouds has contributed to a boom in computational topology. Many standard packages exist to analyze data with regard to topological methods: Gudhi [3] and Ripser [4] being the most influential throughout this piece of work.

The remainder of the paper is as follows: Sec. 2 will be devoted to persistent homology and its significance and will provide an intuitive understanding. Some use cases of persistent homology will also be discussed. Section 3 will introduce and analyze the topology of some common attractors from the field of nonlinear dynamics. Section 4 will look at the classic Z24 bridge structural health monitoring (SHM) case study, and explore the role of topology as a metric between partitions of the data set.

2 Topological Data Analysis

Only the strict mathematical formulations will be introduced. For more information regarding the mathematics of group theory and algebraic topology, the authors would advise referring to [5–8], and [1,2,9–15] for more on TDA. For some other interesting applications in economics and genomics, the readers can refer to [9,16,17].

2.1 Manifolds. Manifolds are continuous surfaces from which the data are assumed to be sampled. By understanding the topology of the sampled data points, TDA aims to extract topological

Manuscript received February 2, 2022; final manuscript received July 29, 2022; published online September 5, 2022. Assoc. Editor: Hameed Metghalchi.




Fig. 1 The first four simplices

information about the underlying manifold from the data. The manifold shape is unknown prior to analysis, and persistent homology will identify features within the manifold over a range of length scales. Thereby, understanding the shape of the sampled data, TDA conjectures that the shape of the manifold is also understood.

Formally, a manifold is a space that is locally homeomorphic to some n -dimensional Euclidean space, R^n . Manifolds are primarily thought of as the underlying space from which the generated or collected data are sampled. Throughout this work, there will be a reoccurring idea that a change in the parameters of a dynamic system may change the shape of the associated manifold; therefore the topology can potentially be used to identify a change in the system.

2.2 Simplicial Complex. Simplicial complexes will be used as a way of attributing quantifiable shape to the data; they can be thought of as higher-dimensional analogs of graphs, giving a way of encoding connections between vertices. In TDA, the vertices of the simplicial complexes are the observed data points. Simplicial complexes can be analyzed to output the homology of the data, and following this, the persistent homology. The homology and persistent homology are key topological invariants that can be used to describe the structure of the data.

A simplicial complex is a structure made up of fundamental building blocks called *simplices*, the first four types of which are shown in Fig. 1. Each vertex in the simplex is fully-connected to all the other vertices and the space enclosed by the vertices is part of that simplex. For instance, Δ^2 encloses a two-dimensional area, and Δ^3 encloses a three-dimensional volume. This sequence can be generalized for Δ^k enclosing a k -dimensional space between $(k + 1)$ fully-connected vertices.

A *filtration* of a simplicial complex, K , is a nested sequence of sub-complexes, such that, $\emptyset = K^0 \subset K^1 \subset \dots \subset K^m = K$, where $K^{i+1} = K^i \cup \Delta^i$ where Δ^i is a simplex of K [11], and \emptyset is the empty set, i.e., there are no points inside.

There are many ways to construct a simplicial complex from point data. For simplification, only one method will be discussed

in this paper, i.e., the *Vietoris–Rips* (VR) complex [18]. The VR complex can be constructed for point data present in a *metric space*.

A *metric* ∂_X is defined on a set X , which maps two elements from X into the positive real numbers. The mapping gives the associated distance between the two elements. For a set to be deemed a metric space, it must satisfy the following axioms:

- (1) for any $x, y \in X$, then $\partial_X(x, y) = \partial_X(y, x)$;
- (2) for any $x, y \in X$, then $\partial_X(x, y) \geq 0$, and $= 0$ if $x = y$;
- (3) for any $x, y, z \in X$, then $\partial_X(x, z) \leq \partial_X(x, y) + \partial_X(y, z)$.

If the aforementioned axioms are satisfied, then (X, ∂_X) forms a metric space, where ∂_X is the metric on X .

An *open ball* is defined on a metric space, (X, ∂_X) . For a point $a \in X$ and $\varepsilon > 0$, the subset of X consisting of all the points $x \in X$ such that $\partial_X(a, x) < \varepsilon$ is referred to as the open ball of radius ε at a .

For the VR complex $VR_\varepsilon(X, \partial_X)$, let (X, ∂_X) be a finite metric space and $\varepsilon \in R >_0$ be a fixed value that determines the scale of the VR complex [19], where X is the set of vertices and ∂_X is the metric on X . A VR complex is defined by the following condition

$$\{x_0, \dots, x_k\} \in VR_\varepsilon(X, \partial_X) \Leftrightarrow \|x_i - x_j\| \leq \varepsilon, \forall i, j \in \{0, \dots, k\} \quad (1)$$

As ε goes from 0 to $+\infty$, a nested sequence of complexes defines the *Vietoris–Rips filtration* [11].

This process is depicted in Fig. 2 for some randomly-sampled data and an arbitrary value of ε , which can be seen as the radius of the balls. The existence of a simplex is determined by how the balls intersect between the vertices. For a VR complex, a simplex between some set of vertices is formed if the Euclidean distance between all those vertices is less than ε .

2.3 Persistent Homology. Before delving into persistent homology, the homology of a topological space needs to be defined. Many great introductions and summaries of homology already exist; the more intrigued reader may want to consult Refs. [6,7,9–11,20]. In very quick terms, the homology is

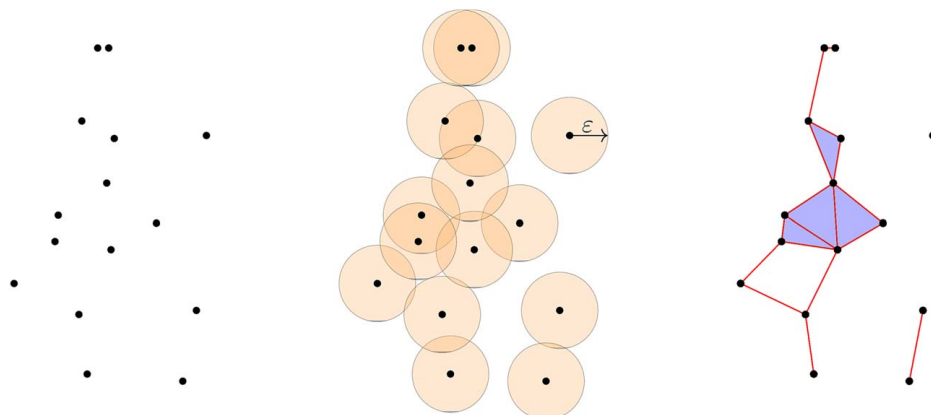


Fig. 2 The process of constructing a VR complex

calculated by computing successive boundary operations on a *chain complex* of a topological space. Homology is essentially the quantification of the number of voids present in a topological space, or the case of TDA, a simplicial complex. The homology groups, $H_k(X)$ are invariants for the data set X , where k refers to the dimension of the homology group. Generally, the k th homology group encodes information about the number of k -dimensional holes in the data. Under the rules of topology, discontinuities (voids) cannot be created or destroyed under continuous maps (homeomorphisms). Therefore, a simplicial complex can be categorized by the properties underpinned by homology. The homology can be used to categorize and compare simplicial complexes, and by extension, data sets.

From the homology, the *Betti numbers* are defined as the *rank* of the homology groups. If the Betti numbers for two topological spaces are different, these spaces are not topologically identical, meaning a continuous bijective map between the spaces does not exist. The zeroth Betti number, β_0 , is the rank of the zeroth homology group, $H_0(X)$, and refers to the number of connected sets in the topological space. The first Betti number, β_1 , is the rank of the first homology group, $H_1(X)$, and refers to the number of non-contractible holes present in the space. The second Betti number, β_2 , refers to the number of enclosed volumes in the topological space. This analogy carries on further for higher dimensions.

This discussion now raises the question: which length scale ε is representative of the topology of the data? When constructing the VR complexes, for the same data set, different values of ε will result in different values for the Betti numbers. The hyper-parameter ε determines the Betti numbers for that specific instance of some point cloud data. Additionally, when the feature present within the data is at a length scale less than ε this feature will not be expressed, as ε will span the feature. This reasoning means that only topological properties that are described at a length scale greater than ε can be captured. A problem arises here, as usually the feature scale is not known prior to analysis, and a manifold may have many multi-scale features.

The answer to this problem is to vary ε and see how the Betti numbers evolve and *persist*. When varying ε , a *filtered simplicial complex* is formed. This can be thought of as a chain or sequence of simplices; with n disconnected points at the beginning and a fully connected $(n-1)$ - simplex at the end. The next simplicial complex in the *filtration* is the earlier simplicial complex plus the next simplex to be formed as ε increases. Figure 3 shows some simplicial complexes in this process. A filtered simplicial complex is an ordering of simplices to show how they evolve as a distance scale is increased. From this, one can form an idea of how the topology evolves over the filtration.

Obtaining the homology for a single value of ε provides very limited information; this notion is almost redundant, because of potential varying feature length scales in the manifold. For this reason, it is vital to consider what homological features persist as ε is varied. The goal of persistent homology is to track the homology classes as ε is varied. The process of varying ε does not bias any disk size, as all are being considered. This process will give an initial value, ε_{\min} , where a specific homological feature comes to life and ε_{\max} , where the feature is no longer considered for that simplicial complex. This range of values $[\varepsilon_{\min}, \varepsilon_{\max}]$ is called the *persistence interval* for that homological feature. Each persistence interval is attributed to a Betti number. Following this, the set of all persistence intervals is descriptive for that manifold, giving information about in which dimension a hole exists in the data and over what range of values it persists.

The persistence intervals obtained can be represented in two ways: *barcodes* or *persistence diagrams*. In a barcode representation, the x -axis refers to the value of ε . As ε increases, the barcode shows which features persist. The set of intervals is plotted with each interval beginning at its ε_{\min} and ending at its ε_{\max} . The color of the interval on the barcode refers to the Betti number, β_k [15]. The value of the y -axis can simply be thought of as indexing of the intervals in the barcode.

An example of a barcode is shown in Fig. 3, with vertical dotted lines showing the intersections with the intervals, showing which features are present for the corresponding simplicial complex in the filtration. The bars starting at 0, refer to the persistence intervals for β_0 , similarly, the bars that do not start at 0 are for β_1 . A bar is formed when that feature begins and ends when that feature dies. The number of intersections of the persistence intervals with the slices (denoted by the dashed black line) at each ε refers to the Betti number for each feature. For instance, for $\varepsilon = 0.3$, the simplicial complex consists of one connected component (as there is only one intersection with the persistence intervals) and there are two holes present (denoted by two intersections with the β_1 persistence intervals). For $\varepsilon = 0.50, 1.00$, these are shown beside the barcode, as there is no change in the barcode for these values. For $\varepsilon > 0.45$, all of the holes have been spanned, and there is only one connected component. The vertices only become more and more connected as ε increases, and this has little topological interest. The length of the interval represents for how long the feature persists. The longer the feature persists, the higher the probability that this feature is characteristic of the manifold. Shorter intervals are generally regarded as topological noise.

The other method of visualizing the set of persistence intervals is the persistence diagram or *birth–death* diagram. In this representation, ε_{\min} is plotted on the x -axis and ε_{\max} is plotted on the y -axis, with each interval represented by the point $(\varepsilon_{\min}, \varepsilon_{\max})$. Intuitively, there is a line defined by $y=x$, below which points will not be plotted. This fact means that the gray region in Fig. 3 will never contain any points. The line $y=x$ has the interpretation that the feature must first exist before it can die. Reading these diagrams is as intuitive as reading the barcodes; the vertical height of the point from the line $y=x$ is analogous to the length of the interval, i.e., the further a point is from the line $y=x$, the more the feature is persistent. An example of a birth–death diagram is shown in Fig. 4.

On both the barcode and birth–death diagrams, it can be seen that one feature persists to ∞ ; this is because there will always be a fully-connected simplex that persists to infinity. There will be a value of ε_{fc} that results in a fully-connected simplex, where every vertex is connected to every other vertex. For values $\varepsilon > \varepsilon_{fc}$ the simplex will remain fully-connected, and therefore this will continue to infinity. The removal of this infinite interval is called *reduced homology*. The reduced homology is required for use in many calculations.

The space of barcodes actually forms a metric space [9]; the distance between the barcodes is a measure of the similarity between two barcodes. As the persistent intervals are invariant for a manifold, the data manifolds can be represented by their persistent homology. This notion of a metric space allows one to compare the similarity of manifolds. Metrics between barcodes are well established and the one used in this paper is the *p-Wasserstein distance* (WD) [21].

DEFINITION 2.1. Given two barcodes B_1 and B_2 . For $p > 0$, the *p-Wasserstein distance*, ∂_{W_p} , is given by

$$\partial_{W_p}(B_1, B_2) = \left(\inf \sum_{Z \in B_1} d_\infty(Z, \phi(Z))^p \right)^{\frac{1}{p}}$$

where ϕ is a matching between B_1 and B_2 , d_∞ is the supremum metric, and Z is a persistence interval in B_1 [9].

The Wasserstein distance is used to measure the similarity between two persistence barcodes, by computing the sum over their edge lengths. The q -Wasserstein distance is defined as the minimum value achieved by a perfect matching between the points of the two diagrams, where the value of matching is defined as the q th root of the sum of all edge lengths to the power q [21].

2.4 Calculating Fractal Dimension. The first application of persistent homology in this paper is on calculating *fractal dimension*, which is a measure of how much space a set occupies. To




Fig. 3 A persistence barcode with realizations showing which simplicial complexes are present at values of ϵ of 0:10, 0:15, 0:20, 0:25, 0:30, 0:35, 0:40, 0:50, and 1:00

gauge the idea of the fractal dimension intuitively, consider that a single point is zero-dimensional, a line spanning two points is one dimension, and the area enclosed by three points is two-dimensional, and so on. Now, suppose a curve, consisting of purely one-dimensional elements has infinite length and is bounded by a finite region. This curve may be believed to occupy more space than something that is only considered one-dimensional, but it also does not completely fill this arbitrarily-defined space. A logical conclusion is that the dimension of this curve is somewhere *between* one and two, depending on how efficiently it occupies this area.

The first notable work on quantifying fractal dimension was by Hausdorff [22], where a measure of roughness was conceived; this was later expanded on by Falconer, and generalized for point clustering in Ref. [23]. Mandelbrot popularized the idea of fractal dimension in the notable work [24].

The two most common methods of quantifying fractal dimensions are the *Hausdorff* and *box-counting* dimensions. The box-counting method exploits the dimensional relations to scalability and the difference in occupied space over different scales. The box-counting method is very popular because of its ease of computation.

The method of using the persistent intervals, which is described later, more closely represents the box-counting method.

To calculate the box-counting dimension, assume a space X . The smallest number of sets of diameter that can cover X , is referred to as $N(\epsilon)$. The scale determines $N(\epsilon)$, and the dimension of X determines the rate of change of $N(\epsilon)$. As the scale is altered, it is expected that

$$N(\epsilon) \simeq c \epsilon^{-d_B} \quad (2)$$

for positive constants d_B and c . It is said that X has a box dimension of d_B . To solve for the dimension of the set, this equation can be rearranged to give

$$d_B = \lim_{\epsilon \rightarrow 0} \frac{\log(N(\epsilon))}{-\log(\epsilon)} \quad (3)$$

with the constant term, c , disappearing in the limit [23].

Inherently, fractals are infinitely-complex objects that are self-similar, often over an infinite range of scales. This view gives a problem when working with finite point-data sets, as only finite information can be captured in a finite point-data set. When

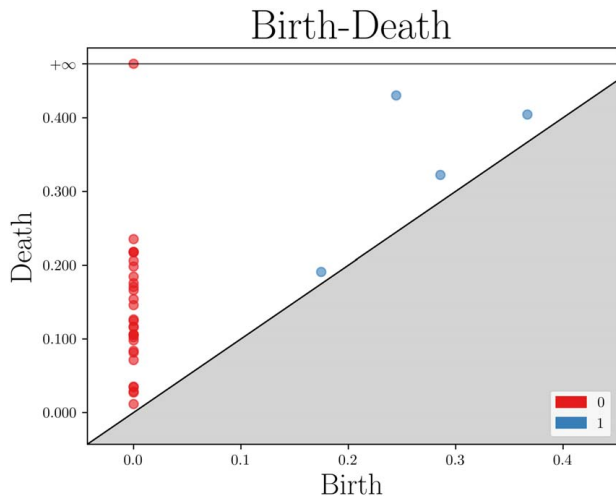


Fig. 4 Birth–death diagram for the same random data present in Fig. 3

zooming in excessively, points will become sparser and the approximations will inevitably become less exact. In reality for finite sets, self-similarity may only be visible over a small number of scales. This process can be seen later, in Fig. 9, where the self-similarity is clear over two length scales, and then thereafter points become sparse.

2.5 Fractal Dimension—Persistent Homology. Persistent homology can be used to calculate fractal dimension, given some point data sampled from a fractal. This idea first came from that of the minimal spanning tree (MST) [25,26]. The MST method was shown as a viable method to calculate the fractal dimension of a set [27]. Numerous works have shown that the MST method is equivalent to using the zeroth homology group for calculating the fractal dimension [27–30]. Unlike using the MST method, the persistent homology can be generalized for higher-order homology groups, therefore returning more information regarding the topology of the fractal shapes.

Often when analyzing the persistent data, the smaller intervals are discarded as noise, as this mean that this specific feature only persists for a short while. In the case of calculating fractal dimension, there is information contained in all of the persistent intervals, as the idea here is to deduce the fractal dimension of X_n points as $n \rightarrow \infty$

and see how the lengths of the persistence intervals in each homology group vary. The shorter intervals provide a good measure of how the local geometry is present in the finite random sample.

To calculate the fractal dimension from the persistent homology, the α -weighted sum for $\alpha > 0$, of the persistent intervals, in a given dimension should be computed, as follows:

$$E_\alpha^i(x_1, \dots, x_n) = \sum_{I \in PH_i(x_1, \dots, x_n)} |I|^\alpha \quad (4)$$

(x_1, \dots, x_n) are n points sampled from a manifold, most interestingly one that exhibits fractal behavior. $PH_i(x_1, \dots, x_n)$ is the i th dimension persistent homology group. $|I|$ is the length of the interval from PH_i . The α -weighted sum of the persistence intervals tracks the rate at which the topological noise decays, and similarly to the box-counting method, this can be linked to the fractal dimension [31]. The value of alpha will give a larger weighting to longer intervals.

DEFINITION 2.2. Let X be a bounded subset of a metric space and μ a measure on X . For all $i \in N$ and a value $\alpha > 0$, which is used to give a weighting to larger persistence intervals in the calculation [31], the persistent homology dimension can be defined as

$$\dim PH_i^\alpha(\mu) = \frac{\alpha}{1 - \beta} \quad (5)$$

where

$$\beta = \limsup_{n \rightarrow \infty} \frac{\log(\mathbb{E}(E_\alpha^i(x_1, \dots, x_n)))}{\log(n)} \quad (6)$$

Here, sup refers to the largest value in the set. The operator E is used as the expected value of a random variable. Finally, $E_\alpha^i(x_1, \dots, x_n)$ is the α -weighted cumulative sum over the n points. This result means that the embedding dimension $d = \dim_{PH_i} \alpha(\mu)$, of a manifold can be calculated if $E_\alpha^i(x_1, \dots, x_n)$ scales with $n^{d-\alpha d}$ [29,31].

For the case when $i = 0$, the persistence intervals are equivalent to the lengths of the edges in the MST. From the MST, there is a set of n vertices and a set of $n - 1$ edges, where each edge spans two vertices. This work is built on ideas from Kruskal [25] and Prim [26] in the 1950s. The edges in the MST are equivalent to the persistence intervals in the reduced zeroth homology group.

The MST approach formulates the problem in terms of graph theory, with V being the set of vertices and E being the edge-set, connecting the vertices. Two vertices are referred to as connected, if a path connects them. A path is the successive joining of adjacent edges from one vertex to another, i.e., one can walk from one vertex to the other without leaving the path. If these edges form a closed




Fig. 5 Hénon attractor constructed with the two methods of calculating fractal dimension: (a) VR complex and (b) MST

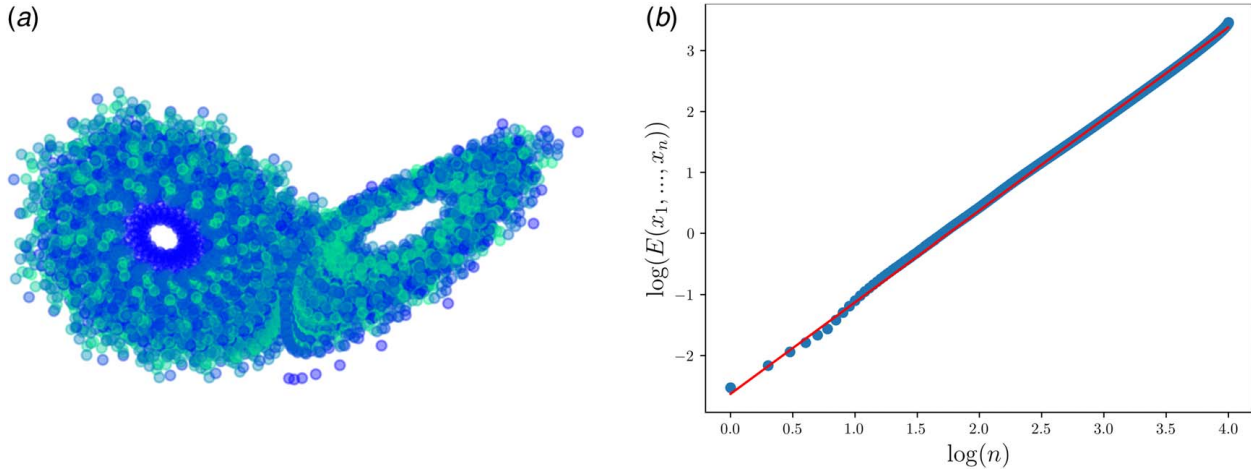


Fig. 6 Lorenz attractor: (a) dense Lorenz attractor and (b) log-log plot to determine fractal dimension, with gradient

loop, this is called a *circuit*. If a graph does not contain a closed circuit, it is called a *tree*.

A spanning tree has an attribute called its *length*. The length of a spanning tree is the sum of all the edges in the tree. The MST is the spanning tree that spans the points most efficiently, by minimizing the length of the edges. The algorithms used to calculate MSTs are very well optimized, meaning that competitive results for calculating fractal dimensions can be obtained from their usage [31]. One can use the lengths of the edges in an MST to approximate the dimension of the manifold from which the vertices are sampled. The asymptotic behavior of Eq. (4) is studied to calculate the fractal dimension.

This property is useful, as algorithms for calculating MSTs are much faster than calculating persistent intervals. However, the MST is only equivalent to the zeroth homology group. If more information is required for higher-dimensional homological features, the slower, but more informative persistence algorithms must be used. Figure 5 shows the VR complex and the MST for an attractor.

When dealing with higher-order persistent homology, $i \geq 1$, things become more tricky. With PH_0 , it is known that there are going to be $n - 1$ edges for n points. However, for higher-order persistent homology, this is not the case. There exist metric spaces where the number of persistent intervals grows faster than linearly in the number of points. To get around this problem, a limit for the upper bound can be proven. In the case of the VR complex, Schweinhart proved $|PH_1| = O(n)$ for the points (x_1, \dots, x_n) [29].

3 Attractors

Strange attractors arise from nonlinear dynamical systems. A dynamical system is modeled (and evolves) in a phase space embedded in R^n , where the geometric object is embedded in Euclidean space and parameterized by time. Attractors represent the state to which a dissipative dynamical system will eventually converge, seemingly regardless of initial conditions. Attractors are commonplace for study as they describe the asymptotic behavior of many dynamical systems [32–34].

3.1 Lorenz Attractor. The Lorenz attractor is defined by the differential equations

$$\dot{x} = -\sigma x \quad (7)$$

$$\dot{y} = \rho y - y - zw \quad (8)$$

$$\dot{z} = -\beta w + uv \quad (9)$$

The attractor was first discovered by Lorenz, when studying non-periodic turbulent flows [32]. It displays an interesting topology, with two holes being present in the manifold. It also shows a Cantor-set like behavior over its cross sections [35] for certain parameters.

For the system parameters $\rho = 28$, $\sigma = 10$, and $\beta = 8/3$, the fractal dimension has previously been calculated by Viswanath [35] to be 2.063. To obtain a structure with fine details, showing the Cantor-like fractal cross sections, a large number of points are required for the embedding. Ten thousand points were calculated here for a dense embedding of the Lorenz attractor. Using the MST method of calculating the fractal dimension [31,36,37], of the embedded attractor shown in Fig. 6, an approximation to the Hausdorff dimension of 2.0826 ± 0.03603 was calculated. This value is calculated by determining the linear gradient of the log-log plot of the cumulative sum of the MST edges, which are analogous to the persistent intervals in PH_0 .

Now, in order to capture the global topology of the system, fewer points are required to calculate the persistent homology. In this example, a more sparse Lorenz attractor is sampled with only 1000 points, as shown in Fig. 7. These samples were taken over the same range as the earlier case but now there is a greater time-step between each point. This procedure ensures a relatively dispersed distribution of points over the Lorenz attractor. By taking a dispersed distribution of points on the manifold, the global topology of the manifold should become clearer.

The birth–death diagram can be seen on the right-hand side of Fig. 7. There are no large differences in β_0 , represented by the points lying along the line $x=0$; this strongly indicates that the Lorenz attractor consists of only one connected component. For the case of the first homology group, there is a large amount of noise present close to the line $y=x$. Most interestingly from this plot, two features persist long enough to be deemed properties of the manifold. The points are (2.59, 7.29) and (1.97, 9.17). The holes represented by these points can be visibly seen on the left-hand side of Fig. 7 as the holes present within the manifold. The smaller of the two intervals represents the void on the left-hand side of the plot.

3.2 Hénon Attractor. The Hénon attractor is a two-dimensional quadratic map with a constant Jacobian; it was first conceived as a simplified discrete map of the Lorenz system. As the map is discrete, it has become a common object for study in dynamical systems, as computations for generating a large number of points are fast [34]. The Hénon attractor used in this work is defined by

$$x_{n+1} = 1 - ax_n^2 + y_n \quad (10)$$

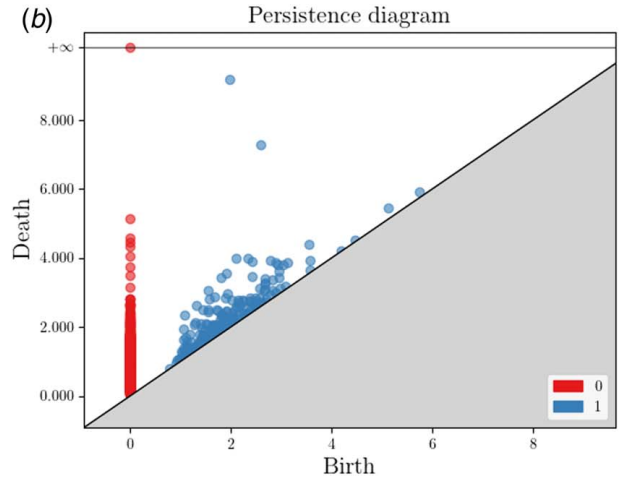



Fig. 7 Lorenz attractor: (a) Lorenz attractor and (b) persistence homology for the Lorenz attractor

$$y_{n+1} = bx_n \quad (11)$$

In this case, the well-studied parameters of $a = 1.4$, $b = 0.3$, and an initial point of $(0.1, 0.3)$ were used, as these give a convergent solution. Two thousand iterations were taken to give a finite approximation of the orbit of this specific Hénon attractor, with the initial conditions listed. Using the MST method of calculating the fractal dimension [31,36,37], a value of 1.2558 ± 0.04476 was obtained. A previously-calculated value from work by Grassberger and Procaccia [38], calculated the dimension to be 1.26 (Fig. 8).

The Hénon attractor is well known to exhibit a self-similar, Cantor-like, behavior over its cross sections [34]. The longitudinal structure of the Hénon attractor is simple, with each curve appearing to be a 1D manifold. The traversal structure is the interesting part; this can be seen to be self-similar in Fig. 9.

The dimensions of (a), (b), and (c) in Fig. 9 are 1.233, 1.245, and 1.526; this shows that, as the data become more sparse, this fractal dimension calculation method becomes less accurate as more data points are missing at the smaller length scales

3.3 Rössler Attractor. This attractor was first formulated in Refs. [33,39], a dense example of which is shown in Fig. 10. The Rössler attractor is obtained by solving the differential

equations

$$\dot{x} = -y - z \quad (12)$$

$$\dot{y} = x + ay \quad (13)$$

$$\dot{z} = b + z(x - c) \quad (14)$$

The values for the initial conditions and system parameters were $a = 0.2$, $b = 0.2$, and $c = 5.7$, with the initial point, $p_0 = (0, 0, 0)$. With 20,000 points sampled from the Rössler attractor, the dimension estimate gives a value of 2.025 ± 0.0246 [31,36,37]. Kuznetsov calculated the dimension of the Rössler attractor, for these parameters to be 2.0160 [40].

From looking at the persistence barcode in Fig. 11, it can be seen that all the dots, corresponding to β_0 are close and show no clear split between these values. This observation strongly implies that the Rössler attractor is a fully-connected manifold. The range in which these occur is most likely because of the data points along the “flick” of the Rössler attractor being more sparse; if this were a continuous embedding, these features would not be present.

For the case of the first-dimensional Betti numbers, represented by the points not lying along $x = 0$ in Fig. 11, there is a lot of topological noise close to the line $y = x$. Interestingly, there is a point far from the line $y = x$, which is representative of the topology of the Rössler attractor. This point has the coordinates $(4.00, 10.3)$ and





Fig. 8 Hénon attractor: (a) Hénon attractor generated with 2000 points and (b) log-log plot to determine fractal dimension, with gradient




Fig. 9 Hénon points sampled from three different scales, showing the Cantor cross section: (a) $x \in [0.5, 0.75]$, $y \in [0.15, 0.21]$, (b) $x \in [0.62, 0.64]$, $y \in [0.185, 0.191]$, and (c) $x \in [0.63, 0.6325]$, $y \in [0.1889, 0.1895]$

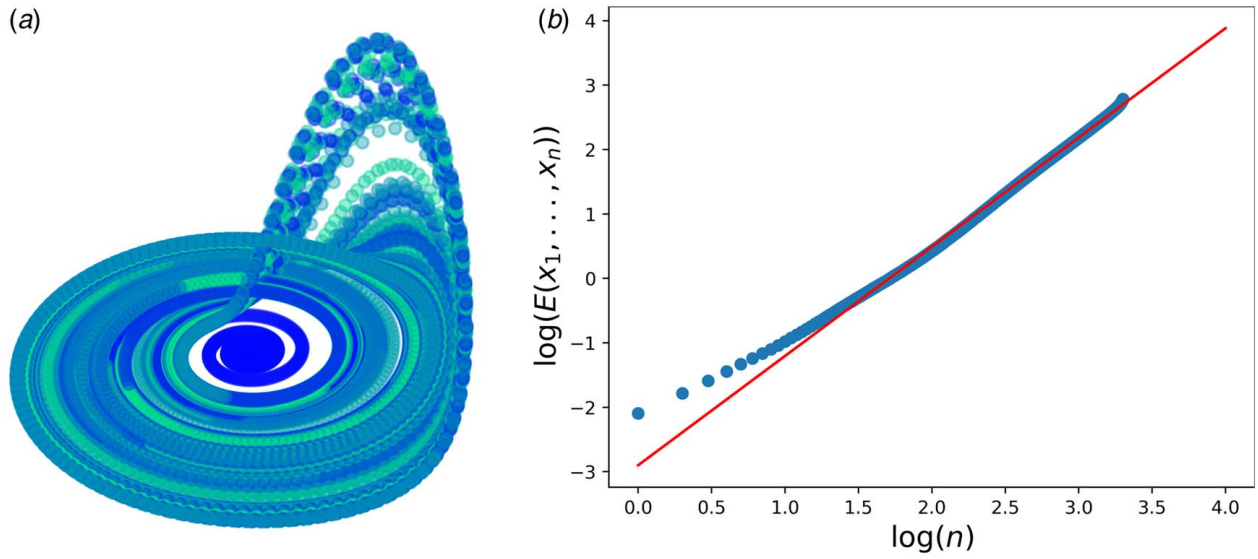


Fig. 10 Rössler attractor with 20,000 points, used to calculate the embedding dimension: (a) Rössler attractor and (b) log-log plot to determine fractal dimension, with gradient

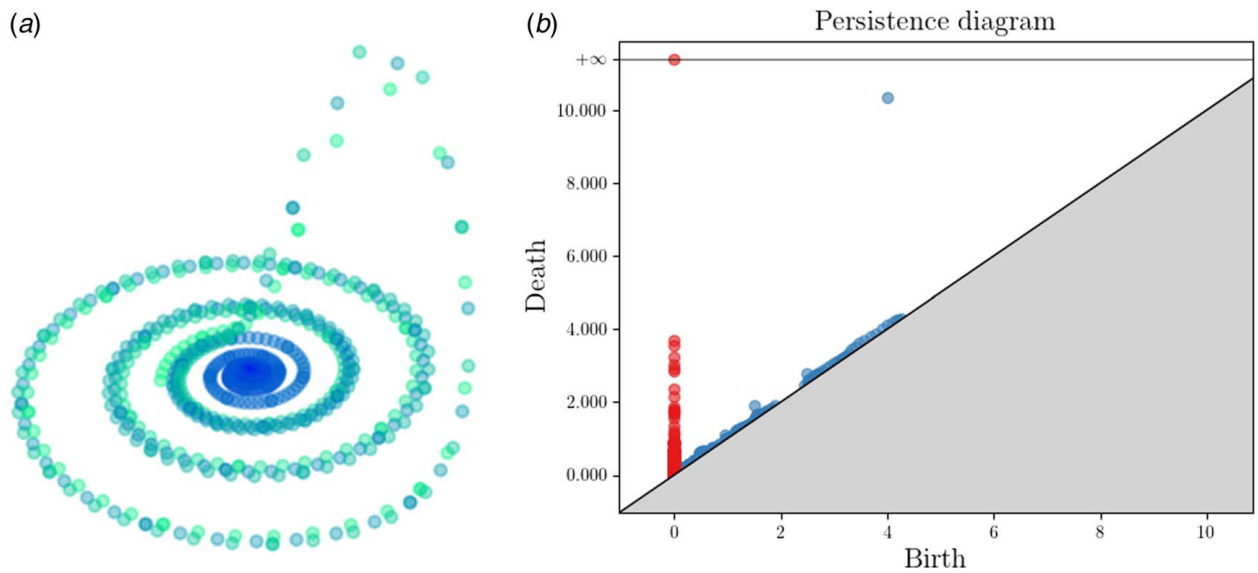


Fig. 11 Rössler attractor, with 700 points, used to analyze the topology: (a) Rössler attractor and (b) persistence homology of the Rössler attractor

is representative of the hole formed over the “flick” of the attractor. These persistence data will now act as the benchmark for the Rössler attractor, and multiple embeddings will be compared to this. The aim is to minimize the WD between the reconstructed object and the original attractor. The smaller the WD, the closer the topologies.

3.4 Reconstruction From Time-Delay Embedding. In this case, the topologies of a reconstructed phase space, formed from a 1D time series from the attractor, and the original point cloud are compared. Persistent homology will be used to determine the optimal embedding parameter from the 1D time series.

Time-delay embeddings were first introduced in Ref. [41] as a method of inducing geometry from a one-dimensional time series. This work was further expanded on by Takens [42], proving that the topology can be perfectly reconstructed for chaotic attractors. Skraba et al. [43] then showed that persistent homology could be used with the time-delay embedding to yield useful topological results.

Given a time-varying series, $f: t \rightarrow \mathbb{R}$, the time-delay embedding can be stacked d times, each with a delay α , to give a new embedding $\varphi: t \rightarrow \mathbb{R}^d$ where the new embedding is represented by

$$\varphi(t) = (f(t), f(t + \alpha), \dots, f(t + (d - 1)\alpha)) \quad (15)$$

From this embedding, the time series now has a topology induced. This embedding is a reconstruction of the topology from a one-dimensional time series into any desired dimension. Takens’ theorem shows that under certain circumstances, this reconstructed attractor is homeomorphic to the original attractor [42].

The delay embedding highlights periodic recurrent features in the time series. Recurrent behavior will be highlighted in the time-delay embedding as a loop. Persistent homology can then be used to quantify the size of these loops.

If the delay is too small, there will not be sufficient information to form meaningful topology, and the reconstruction will be too similar to a straight diagonal line. Therefore, by maximizing the size of the holes present in the attractor by using the persistent homology, an optimal time-delay embedding can be determined. Conversely, using a too-large delay will result in nonsense, as the gap between the readings will be too large and will not show a local change over the manifold, resulting in a deformed attractor.

By computing a range of delay embeddings, the persistent homology of these can be calculated, and then the WD can be used as a metric of topological similarity, to give the optimal delay embedding. Figure 12 shows the WD between the original Rössler attractor and attractors formed from a delay embedding over a range of delays. By taking the delay for the minimal WD,

this implies that the topology of the original attractor and the reconstructed delay embedding has the most similar topology. Therefore, the delay that minimizes the WD is the one that gives the optimal delay embedding.

There is a clear periodic trend emerging in Fig. 12; this trend occurs as the time-delay embedding comes in and out of phase with the original attractor. The general trend in Fig. 12 shows an increase in the WD as the delay is increased. This trend is because of artifacts induced in the reconstruction, as successive points are sampled over different periods from the attractor.

From Fig. 12, it can be seen that a delay equal to one gives the smallest WD; in this case, the reconstructed topology is too similar to the diagonal set, and the topology is relatively uninteresting. If one was to assume a delay embedding of $\alpha = 0$, this would result in the straight line $x = y = z$, which will clearly have an uninteresting topology. Therefore, in this analysis, only values after the first peak ($\alpha = 13$) will be considered, as these ensure an interesting topology is being formed from the embedding. On the other hand, using a too-large value for the delay gives a deformed manifold. Two bad examples of delay embeddings, $\alpha = 1$, and $\alpha = 57$, are shown in Fig. 13.

For $\alpha = 19$, not only the topology of the attractor is most accurately being captured, but from Fig. 14(a), there is also a strong likeness to the original geometry. This is not likely to be the case for all other reconstructions, as Takens’ only proved this for the topology [42].

4 A Case Study of Interest for Structural Health Monitoring: The Z24 Bridge

4.1 Introduction to the Problem. The Z24 bridge was a structure connecting the two regions, Koppigen and Utzenstorf, in Switzerland. Data were collected over a whole year, with a sensor network placed over the bridge to collect modal parameters. Sensors were also used to measure the air temperature, soil temperature, and humidity. Because of the extreme conditions of the Swiss weather, the air temperature was recorded as low as -9°C and as high as 36°C . As a result, the temperature effects are clearly visible on the calculated natural frequencies [44]. Shortly before the destruction of the Z24 bridge, there was controlled damage introduced to the system, which is also visible in the natural frequencies after a certain point in time.

The change in temperature is the biggest contributing factor to the change in the natural frequency ($\approx 30\%$). The change in temperature has a greater impact than introducing damage ($\approx 7\%$). For this reason, the change in the magnitude of the natural frequencies offers little insight into the presence of damage. The main problem here is to separate the damage case from the temperature effects.

The data set can be broken down into four categories, according to the air temperature at the time of the measurements, and whether the damage was present. Figure 15 shows the temperature readings and the first four calculated natural frequencies; the corresponding colors refer to:

- Light blue makes up the freezing data set. This is any value with a temperature reading of $T < 0^\circ\text{C}$.
- Dark blue makes up the cold data set. This is any value in the temperature range of $0^\circ\text{C} \leq T < 4^\circ\text{C}$.
- Red makes up the warm data set. This is any value with a temperature reading of $4^\circ\text{C} \leq T$.
- Black makes up the damage-state data set. This is any reading taken after an index of 3475, irrespective of the temperature.

At every measurement instance, the natural frequencies $\{\omega_i\}$, can be calculated; where ω_i is the resonance frequency for the i th mode. The first set of the n natural frequencies can be represented as a point in \mathbb{R}^n , where the i th axis is for the value of ω_i .

As the Z24 data set has been sampled over different environmental and operational variations (EOVs), it is expected that the

Wasserstein value vs Delay for Rossler Reconstruction




Fig. 12 WD versus time-delay embedding for Rössler

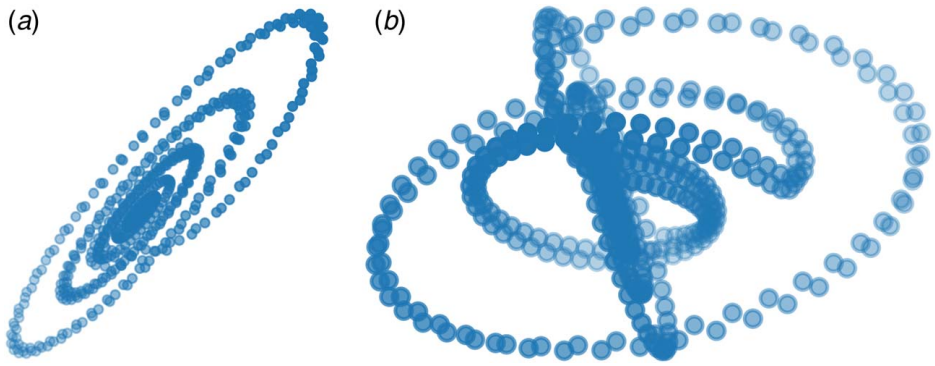


Fig. 13 Examples of bad Rössler reconstructions: (a) $\alpha = 1$ and (b) $\alpha = 57$

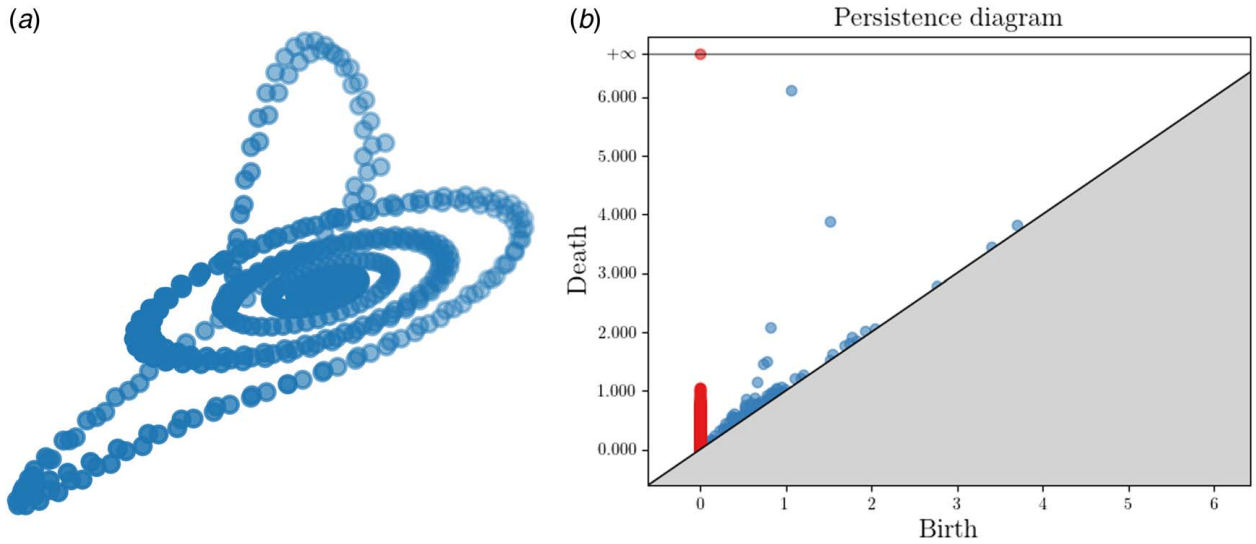


Fig. 14 Optimal reconstructed Rössler attractor: (a) $\alpha = 19$ and (b) persistence homology




Fig. 15 Natural frequency classes according to their temperature

Table 1 WD values over each partition

	Freezing	Cold	Warm	Damage	Warm1	Warm2
Freezing	0.00	9.39	22.92	10.62	12.62	12.50
Cold	9.39	0.00	21.47	5.35	8.78	8.50
Warm	22.92	21.46	0.00	23.44	14.26	14.51
Damage	10.62	5.35	23.44	0.00	10.09	9.59
Warm1	12.62	8.78	14.26	10.09	0.00	1.80
Warm2	12.50	8.50	14.51	9.59	1.80	0.00

Table 2 Summed and scaled WD data

	WD sum	Number of points	Scaled WD sum
Freezing	68.050	720	0.095
Cold	53.481	666	0.080
Warm	96.584	2089	0.046
Damage	59.095	457	0.129
Warm1	47.549	1044	0.046
Warm2	46.903	1045	0.045

resonance frequencies will vary, depending on the environmental variables on the day. The bigger the change in the EOVs, the greater the change in ω_i . As the temperature changes, it is expected that material properties will be affected, therefore resulting in a change in the natural frequency.

In the case of the Z24 Bridge, by sampling over the time frame of a year, there will be slight changes in humidity, air temperature, and soil temperature, which mean that each natural frequency will be slightly different from any other day. As the first four natural frequencies have previously been extracted for each reading [44], each point can be plotted in R^4 . Plotting these points will trace out a manifold shape that is parameterized by all of these EOVs. The data points are then assumed to lie on a manifold, representative of this specific bridge. TDA can then be used to form an understanding of what is expected for the shape of the natural frequency manifold for the Z24 data set.

Since the data have been partitioned into freezing, cold, warm, and damage, TDA can be used to compare the relative shapes of these manifolds. A significant change in the shape of the manifold could potentially be an indicator of the presence of damage.

4.2 Wasserstein Distance of Z24 Partitions. There will be three case studies presented, all on the Z24 data set. Each case

study will include a new aspect of analysis to make the analysis more concrete.

- (1) The first case study will compare the manifolds embedded in four-dimensional space, where each of the axes corresponds to a natural frequency $\omega_1, \omega_2, \omega_3,$ and ω_4 . This case will also show how normalizing by the number of points when using the WD makes a more robust metric for comparing data sets of different sizes.
- (2) The second case study will remove ω_2 , and plot the manifold embedded in R^3 . Earlier analysis of the Z24 bridge, has shown that a lot of the nonlinear features are present in the second natural frequency. This case study is useful, as now that the natural frequency data are embedded in three dimensions, the data can be plotted and visualized, whilst showing that topological arguments are still valid in lower-dimensional shadows.
- (3) The third case study shows the robustness of the manifolds to a linear dimension reduction algorithm. In this case, the four-dimensional embedding of the data will be compressed down to two and three dimensions [45]. The reduced-dimension manifolds are then analyzed with TDA.

4.2.1 Raw Natural Frequency Case. For the first case, a full walk-through of the calculation procedure will be displayed. For the succeeding cases, this will be omitted to limit repetition. This case will take the first four natural frequencies obtained from the Z24 data set [44]. The natural frequencies will be represented by a point in four-dimensional space. It is believed that the introduction of damage will change the shape of the manifold more substantially than the change in temperature. TDA can be used here to compare the different shapes of the partitioned manifolds.

The persistent homology of the manifolds is calculated accordingly. The partitions are not all of the same sizes. The warm data set is the largest data set, and this will be randomly split in half to form two data sets, the original warm data set will remain in the analysis. The second can be used to verify the results; the two new random subsets will have a very similar topological structure, as they are all sampled from the same manifold. There will be slight differences because of topological noise formed in the persistence intervals from missing points in the smaller samples. These effects should be negligible when over the true global structure of the manifold, there are enough points to adequately describe the topology. For smaller partitions, the absence of many missing points will affect the topology in a significant way.

Table 1 shows the WDs over the different partitions of the data set. These values are relatively uninteresting, as the WD values are a factor in the number of points present in the data. This effect can be seen clearly between the warm and the warm

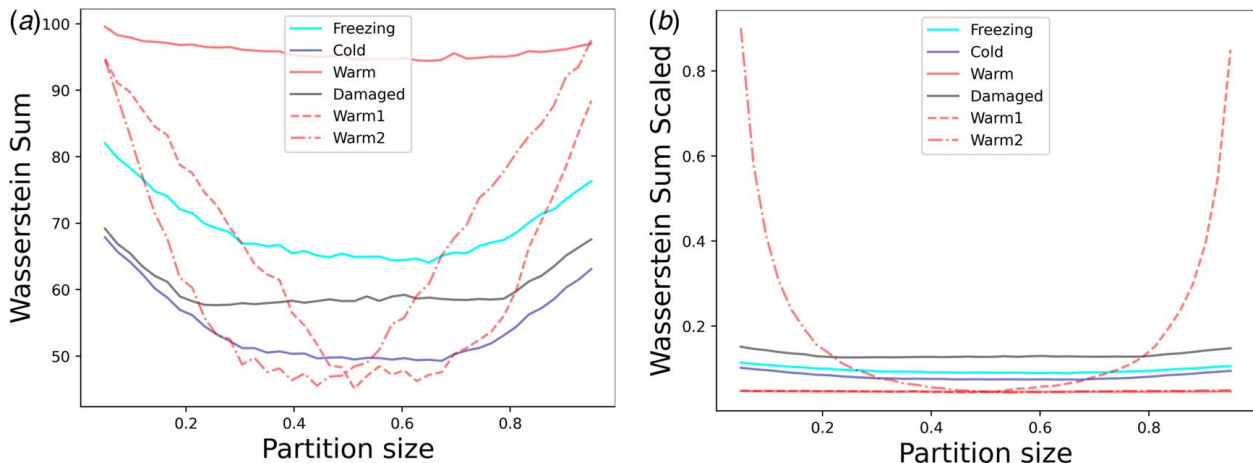


Fig. 16 Size of the WD values depending on the size of the warm partition size: (a) raw and (b) scaled

subsets, as all the WD values for the full warm data set are rough twice the size of the two random subsets. This effect shows that the number of points present in the point cloud is linked to the size of the WD value.

As a more informative measure, the WD values can be summed along the rows. Summing along the rows gives an understanding of how different a data set is from all the others. The larger the value, the more different it is. As well as this, one can normalize by the number of points present in the data set; this gives a normalized value, independent of the size of the data set. This measure then acts as a metric to discriminate between manifold shapes for manifolds with a varying number of points present (Table 2).

Here, the change in the partition size and its effect on the WD value will be discussed, (Fig. 16). The left-hand plot in Fig. 16 shows how the sum of the WD values for any data set varies as the partition size of the warm data set changes. It can be seen that the very cold, cold, warm, and damage sets are all changing proportionally, as the partition size is varied, as is expected. On the other hand, the two warm subsets vary a lot as the partition size changes. For small partition sizes, the topology of the manifold is subsampled to an extreme and a likeness between the manifolds cannot be established.

The right-hand plot of Fig. 16 shows the sum of all the WDs for a data set, and then scaled by the number of points present in the data set. This figure shows how dividing by the number of points almost perfectly maps the warm subsets to the full warm set, in the regions where there are enough points to adequately describe the topology.

As shown, the damage case is clearly the most different in terms of manifold structure, compared to the other data sets. The freezing data are the next most distinct.

For reference, if the partitions are not included in the analysis, results are obtained as in Tables 3 and 4. The damage manifold comes out as less different from the other values, as this now contains a greater weight in the analysis. When the warm subsets were included, this gave a much larger weight to this condition. Despite this, the damage manifold still comes out as the most topologically dissimilar manifold.

4.2.2 3D Shadow: ω_1 , ω_3 , and ω_4 . Earlier analysis of the Z24 Bridge data has determined that ω_2 contains the most nonlinear features. The analysis presented in the earlier section can be repeated, but after eliminating ω_2 and plotting the manifold in three dimensions (Fig. 17). This restriction acts as a visual aid as the data can now be plotted; it also presents a case where a lower-dimensional shadow is taken of the data, and the topology is preserved over this type of subsampling.

This example shows the resilience of the TDA procedure; how a lower-dimensional shadow still contains enough information to distinguish the damage data partition as the most different, shown in Table 5.

4.2.3 Principal Component Analysis. In this section, a linear dimension reduction algorithm, called principal component analysis (PCA) [45] will be applied to the data. Naturally, when reducing the dimension of the data, some information will be lost. The 4D space will be reduced to the principal components present in 3D and 2D; the results of both will be presented. This demonstration will show how the data's topological structure is preserved over linear transformations, and how the accuracy degrades as the dimension reduction becomes more distant from the true embedding dimension.

Table 3 WDs with subsets not included

	Freezing	Cold	Warm	Damage
Freezing	0.00	9.39	22.92	10.62
Cold	9.39	0.00	21.47	5.35
Warm	22.92	21.46	0.00	23.44
Damage	10.62	5.35	23.44	0.00

Table 4 Scaled summed WDs with subsets not included

	WD sum	Number of points	Scaled WD sum
Freezing	42.922	720	0.060
Cold	36.203	666	0.054
Warm	67.817	2089	0.032
Damage	39.411	457	0.086

The first principal component can be thought of as the direction that maximizes the variance of the data into a projected space. The successive principal components are the ones that maximize the variance of the data into the projected space that are also orthonormal to all the previous components.

To clarify how the topological structure is altered when taking the principal components, a case will be considered that uses some points randomly sampled from a torus. The example is shown in Fig. 18. This figure shows that as the PCA embedding dimension is reduced, topological information is lost. When going from a 3D embedding to a 2D embedding, the volume enclosed by the torus is lost. Following this progression, when going from 2D to 1D, the embedding is simply a straight line with a distribution of points weighted at the ends. These changes in the embedding dimension mean topological information in the dimensions higher than the current embedding are lost, i.e., in Fig. 18(b) the 3D-volume is lost, and in Fig. 18(c), the 2D-hole is lost. Despite this loss of information, the topology captured in the remaining principal components is still adequately represented, in Fig. 18(b),

Table 5 WD values with the highly nonlinear ω_2 removed

	WD sum	Number of points	Scaled WD
Freezing	43.373	720	0.060
Cold	32.461	666	0.049
Warm	54.486	2089	0.026
Damage	41.078	457	0.090
Warm1	29.190	1044	0.028
Warm2	29.013	1045	0.028

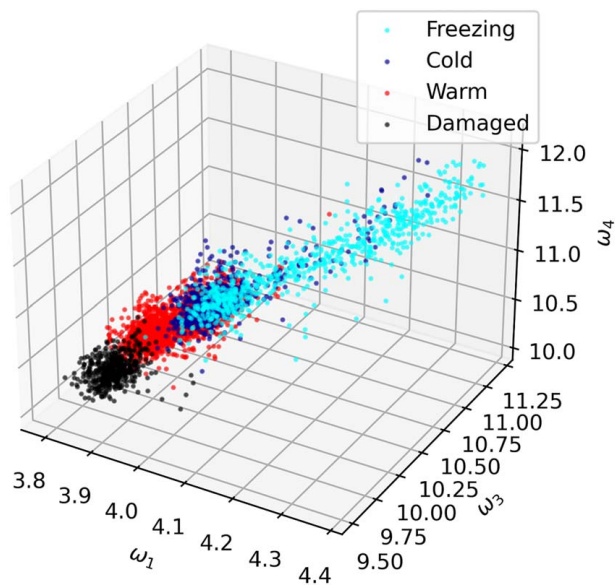


Fig. 17 ω_1 , ω_3 , ω_4 3D plot, with the colors specifying data partitions




Fig. 18 Torus principal components: (a) regular torus, (b) first two principal components, and (c) first principal component

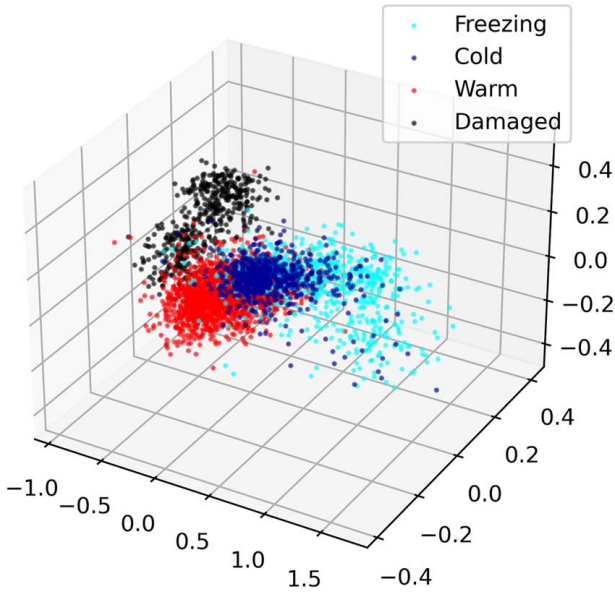


Fig. 19 A visualization of the first three principal components of Z24 natural frequency data

Table 6 WD values for data reduced from 4D to 3D

	WD sum	Number of points	Scaled WD
Freezing	56.763	720	0.079
Cold	42.665	666	0.064
Warm	69.743	2089	0.033
Damage	44.878	457	0.098
Warm1	36.561	1044	0.035
Warm2	37.414	1045	0.036

the 2D-hole is still visible, and in Fig. 18(c), the manifold still consists of one connected component.

The first PCA example presented here on the Z24 Bridge data set, is for the dimension reduction from 4D space into 3D space. As shown in Fig. 19, calculating the PCA of the data has still kept the clusters in somewhat separate regions of space. Under close inspection, the warm, cold, and freezing all seem to be single clusters. Whereas, the damage cluster appears to be comprised of two clusters positioned very close to one another. This property will result in the topology of the damage partition being different from the other clusters.

Performing the PCA of the data is interesting, as it shows that the topology and persistent homology is conserved over the first three




Fig. 20 A visualization of the first two principal components of Z24 natural frequency data

Table 7 WD values for data reduced from 4D to 2D

	WD sum	Number of points	Scaled WD
Freezing	28.074	720	0.039
Cold	19.816	666	0.030
Warm	25.156	2089	0.012
Damage	19.310	457	0.042
Warm1	16.637	1044	0.016
Warm2	16.043	1045	0.015




Fig. 21 ω_1 time-delay embedding $\alpha = 75$

principal components. As shown in Table 6, using the principal components of the data still results in the same orderings between the different partitions.

Now for a larger dimension reduction, into 2D space (Fig. 20). This result still preserves the topology, although the WD values are a lot less distinct between the different cases, shown in Table 7. This result is because more topological information is being lost, as the data are compressed down to only the first two principal components. Despite the loss of information, the ordering from the scaled WD calculations still remains the same.

4.3 Time-Delay Embedding. In this section, the topology of the natural frequency manifold will be reconstructed from each of the 1D time series. For visualization, the Z24 time series data will be embedded in 3D, so that an intuitive understanding of the topology can be formed.




Fig. 22 ω_2 time-delay embedding $\alpha = 75$

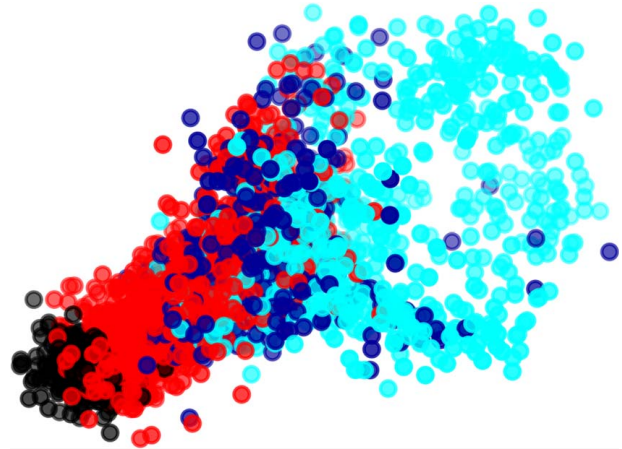


Fig. 23 ω_3 time-delay embedding $\alpha = 75$




Fig. 24 ω_4 time-delay embedding $\alpha = 75$

Figures 21–24 show the delay embeddings of the first four natural frequencies, projected into three dimensions. The embeddings for ω_1 , ω_3 , and ω_4 show a similar structure at the delay value $\alpha = 75$. However, ω_2 has a different topology defined around the colder temperatures. This can be inferred by the second mode having a more pronounced effect from the freezing temperature variations.

When delaying the time series, this means that the delayed dimensions will have different temperature readings from the one-dimensional time series data used to form the embedding. The value for the temperature used to represent the point is the one from the initial time series data. The other dimensions will have their own respective temperatures. The classes used for the temperatures in these plots are a little less certain than in earlier cases. This classification isn't too much of a problem, as the change in temperature is assumed to be a smooth change over each day, and as the value of $\alpha = 75$ is relatively small compared to the size of the data set, the time delay can be inferred as a local change.

5 Conclusion

The focus of this paper is to present an overview of TDA to the structural dynamics community, with a variety of use cases that highlight its depth and borderline limitless capabilities.

For the case of attractors, the calculation of the dimension of some common attractors was facilitated. A novel method for determining the optimal delay for constructing an attractor's topology from a 1D time series was also presented.

With respect to the Z24 data set, topological methods have been able to single out the damage data partition as the most topologically dissimilar. However, further analysis on topological methods for damage detection would need to be explored to understand the true limits and possibilities of TDA in SHM. An insight into the data structure provides powerful insight into the operating conditions of a machine or structure.

Future work on TDA, will look at different topological methods; aside from using the Wasserstein distance as a metric, other case studies will also be considered. A further journal paper, which extends on the ideas presented here, will be submitted.

Acknowledgment

The authors would like to thank the UK EPSRC for funding via the Established Career Fellowship EP/R003645/1 and the Programme Grant EP/R006768/1.

Conflict of Interest

There are no conflicts of interest. This article does not include research in which human participants were involved. Informed consent is not applicable. This article does not include any research in which animal participants were involved.

Data Availability Statement

The datasets generated and supporting the findings of this article are obtainable from the corresponding author upon reasonable request.

References

- Edelsbrunner, H., and Harer, J., 2010, *Computational Topology: An Introduction*, American Mathematical Society.
- Edelsbrunner, H., Letscher, D., and Zomorodian, A., 2000, "Topological Persistence and Simplification," Proceedings 41st Annual Symposium on Foundations of Computer Science, Redondo Beach, CA, Nov. 12–14, IEEE, pp. 511–533.
- Maria, C., Boissonnat, J. D., Glisse, M., and Yvinec, M., 2014, "The Gudhi Library: Simplicial Complexes and Persistent Homology," International Congress on Mathematical Software, Seoul, South Korea, Aug. 5–9, Springer, New York, pp. 167–174.
- Bauer, U., 2021, "Rips: Efficient Computation of Vietoris-Rips Persistence Barcodes," *J. Appl. Comput. Topol.*, **5**(3), pp. 391–423.
- Mendelson, B., 1990, *Introduction to Topology*, Dover, UK.
- Nash, C., and Sen, S., 1988, *Topology and Geometry for Physicists*, Elsevier, New York.
- Ghrist, R., 2018, "Homological Algebra and Data," *Math. Data*, **25**, pp. 273.
- Saunders, M., 1963, *Homology*, Springer, Berlin.
- Rabadan, R., and Blumberg, A. J., 2019, *Topological Data Analysis for Genomics and Evolution: Topology in Biology*, Cambridge University Press, Cambridge.
- Ghrist, R. W., 2014, *Elementary Applied Topology, Volume 1*, Createspace, Seattle, WA.
- Boissonnat, J. D., Chazal, F., and Yvinec, M., 2018, *Geometric and Topological Inference*, Cambridge University Press, Cambridge.
- Chazal, F., and Michel, B., 2017, An Introduction to Topological Data Analysis: Fundamental and Practical Aspects for Data Scientists. arXiv preprint arXiv:1710.04019.
- Zomorodian, A. J., 2005, *Topology for Computing*, Cambridge University Press, Cambridge.
- Carlsson, G., 2019, "Persistent Homology and Applied Homotopy Theory," *Handbook of Homotopy Theory*, pp. 297–331.
- Ghrist, R., 2008, "Barcodes: The Persistent Topology of Data," *Bull. Am. Math. Soc.*, **45**(1), pp. 61–75.
- Gidea, M., and Katz, Y., 2018, "Topological Data Analysis of Financial Time Series: Landscapes of Crashes," *Phys. A: Stat. Mech. Appl.*, **491**, pp. 820–834.
- Li, L., Cheng, W. Y., Glicksberg, B. S., Gottesman, O., Tamler, R., Chen, R., Bottinger, E. P., and Dudley, J. T., 2015, "Identification of Type 2 Diabetes Subgroups Through Topological Analysis of Patient Similarity," *Sci. Transl. Med.*, **7**(311), p. 174.
- Carlsson, E., Carlsson, G., and De Silva, V., 2006, "An Algebraic Topological Method for Feature Identification," *Int. J. Comput. Geom. Appl.*, **16**(4), pp. 291–314.
- Chambers, E. W., De Silva, V., Erickson, J., and Ghrist, R., 2010, "Vietoris-Rips Complexes of Planar Point Sets," *Discrete Comput. Geom.*, **44**(1), pp. 75–90.
- Schutz, B. F., 1980, *Geometrical Methods of Mathematical Physics*, Cambridge University Press, Cambridge.
- Lacombe, T., Cuturi, M., and Oudot, S., 2018, Large Scale Computation of Means and Clusters for Persistence Diagrams Using Optimal Transport, arXiv preprint arXiv:1805.08331.
- Hausdorff, F., 1918, "Dimension und Außeres Maß," *Math. Ann.*, **79**(1), pp. 157–179.
- Falconer, K., 1986, *The Geometry of Fractal Sets*, Cambridge University Press, Cambridge, pp. 7–17.
- Mandelbrot, B., 1982, *The Fractal Geometry of Nature*, WH Freeman, New York.
- Kruskal, J., 1956, "On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem," *Proc. Am. Math. Soc.*, **7**(1), pp. 48–50.
- Prim, R., 1957, "Shortest Connection Networks and Some Generalizations," *Bell Syst. Tech. J.*, **36**(6), pp. 1389–1401.
- Van de Weygaert, R., Jones, B., and Martinez, V., 1992, "The Minimal Spanning Tree as an Estimator for Generalized Dimensions," *Phys. Lett. A*, **169**(3), pp. 145–150.
- Schweinhart, B., 2020, "Fractal Dimension and the Persistent Homology of Random Geometric Complexes," *Adv. Math.*, **372**, p. 107291.
- Schweinhart, B., 2021, "Persistent Homology and the Upper Box Dimension," *Discrete Comput. Geom.*, **65**(2), pp. 331–364.
- Kozma, G., Lotker, Z., and Stupp, G., 2006, "The Minimal Spanning Tree and the Upper Box Dimension," *Proc. Am. Math. Soc.*, **134**(4), pp. 1183–1187.
- Jaquette, J., and Schweinhart, B., 2020, "Fractal Dimension Estimation With Persistent Homology: a Comparative Study," *Commun. Nonlinear Sci. Numer. Simul.*, **84**, p. 105163.
- Lorenz, E. N., 1963, "Deterministic Nonperiodic Flow," *J. Atmos. Sci.*, **20**(2), pp. 130–141.
- Rössler, O. E., 1976, "An Equation for Continuous Chaos," *Phys. Lett. A*, **57**(5), pp. 397–398.
- Hénon, M., 1976, "A Two-Dimensional Mapping With a Strange Attractor," *The Theory of Chaotic Attractors*, B. R. Hunt, J. A. Kennedy, T.-Y. Li, and H. E. Nusse, eds., Springer, New York, pp. 94–102.
- Viswanath, D., 2004, "The Fractal Property of the Lorenz Attractor," *Phys. D*, **190**(1–2), pp. 115–128.
- March, W. B., Ram, P., and Gray, A. G., 2010, "Fast Euclidean Minimum Spanning Tree: Algorithm, Analysis, and Applications," KDD '10 Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, July 25–28, pp. 603–612.
- Curtin, R. R., Edell, M., Lozhnikov, M., Mentekidis, Y., Ghaisas, S., and Zhang, S., 2018, "mlpack 3: A Fast, Flexible Machine Learning Library," *J. Open Source Soft.*, **3**(26), p. 726.
- Grassberger, P., and Procaccia, I., 1982, "Measuring the Strangeness of Strange Attractors," *The Theory of Chaotic Attractors*, B. R. Hunt, J. A. Kennedy, T.-Y. Li, and H. E. Nusse, eds., Springer, New York, pp. 170–189.
- Rössler, O. E., 1979, "An Equation for Hyperchaos," *Phys. Lett. A*, **71**(2–3), pp. 155–157.
- Kuznetsov, N. V., and Mokaev, T. N., 2018, A Note on Finite-Time Lyapunov Dimension of the Rössler Attractor. arXiv preprint arXiv:1807.00235.
- Packard, N. H., Crutchfield, J. P., Farmer, J. D., and Shaw, R. S., 1980, "Geometry From a Time Series," *Phys. Rev. Lett.*, **45**(9), p. 712.
- Takens, F., 1981, "Detecting Strange Attractors in Turbulence," *Dynamical Systems and Turbulence, Warwick 1980*, D. Rand, and L. Young, eds., Springer, New York, pp. 366–381.
- Skraba, P., Vin De, S., and Vejdemo-Johansson, M., 2012, "Topological Analysis of Recurrent Systems," NIPS 2012 Workshop on Algebraic Topology and Machine Learning, Lake Tahoe, NV, Dec. 8, pp. 1–5.
- Peeters, B., and De Roeck, G., 2001, "One-Year Monitoring of the Z24-Bridge: Environmental Effects Versus Damage Events," *Earthq. Eng. Struct. Dyn.*, **30**(2), pp. 149–171.
- Wold, S., Esbensen, K., and Geladi, P., 1987, "Principal Component Analysis," *Chemom. Intell. Lab. Syst.*, **2**(1–3), pp. 37–52.