

Stabilizing the unstable output of persistent homology computations

Paul Bendich, Peter Bubenik, and Alexander Wagner

ABSTRACT. We propose a general technique for extracting a larger set of stable information from persistent homology computations than is currently done. The persistent homology algorithm is usually viewed as a procedure which starts with a filtered complex and ends with a persistence diagram. This procedure is stable (at least to certain types of perturbations of the input). This justifies the use of the diagram as a signature of the input, and the use of features derived from it in statistics and machine learning. However, these computations also produce other information of great interest to practitioners that is unfortunately unstable. For example, each point in the diagram corresponds to a simplex whose addition in the filtration results in the birth of the corresponding persistent homology class, but this correspondence is unstable. In addition, the persistence diagram is not stable with respect to other procedures that are employed in practice, such as thresholding a point cloud by density. We recast these problems as real-valued functions which are discontinuous but measurable, and then observe that convolving such a function with a suitable function produces a Lipschitz function. The resulting stable function can be estimated by perturbing the input and averaging the output. We illustrate this approach with a number of examples, including a stable localization of a persistent homology generator from brain imaging data.

1. Introduction

Persistence diagrams, also called bar codes, are one of the main tools in topological data analysis [10, 35, 31, 37]. In combination with machine-learning and statistical techniques, they have been used in a wide variety of real-world applications, including the assessment of road network reconstruction [3], neuroscience [23], [6], vehicle tracking [5], object recognition [39], protein compressibility [36], and protein structure [38].

Put briefly, these persistence diagrams are multi-sets of points in the extended plane, and they compactly describe some of the multi-scale topological and geometric information present in a high-dimensional point cloud, or carried by a real-valued function on a domain. Several theorems [24, 14, 26] state that persistence diagrams are stable with respect to certain variations in the point-cloud or functional input, and so the conclusions drawn from them can be taken with some confidence.

On the other hand, there is additional potentially very useful but unstable information produced during the computation of persistence diagrams. For example, a point far from the diagonal in the degree-zero persistence diagram represents a connected component with high persistence. This component first appears somewhere and the computation that produces the persistence diagram can be used to find its location. However this location is not stable: as we will describe below, a small change in the input will cause only a small change in the persistence of this connected component, but it can radically alter the location of its birth. We summarize this as follows.

FUNDAMENTAL CONUNDRUM OF TOPOLOGICAL DATA ANALYSIS. *Users of topological data analysis would like to find the simplices or cycles corresponding to the birth of the most significant pairings of critical values. However, unlike the paired critical values, these simplices and cycles are unstable.*

In addition, persistent homology computations may rely on parameters such that the output persistence diagram is not stable with respect to changes of these parameters.

1.1. Our Contribution. This paper introduces a method for stabilizing desirable but unstable outputs of persistent homology computations. The main idea is the following. On the front end, we think of a persistent homology computation \mathcal{C} as being parametrized by a vector $\mathbf{a} = (a_1, \dots, a_n)$ of real numbers. These parameters could specify the input to the computation (e.g. the coordinates of the vertices of a simplicial complex) or they could specify other values used in the computation (e.g. threshold parameters used in de-noising or bandwidths for smoothing). For a given choice of \mathbf{a} , we get a persistence diagram. On the back end, we consider a function p that extracts a real-number summary from a persistence diagram. For example, p might extract the persistence of a homology class created by the addition of a specific edge in a filtered simplicial complex, or it might be an indicator function on whether or not the longest bar was born by the addition of a simplex contained in a fixed region of the input space, or it may indicate whether or not a chosen representative geometric cycle intersects a given region. The composite function h that maps the parameter vector to the real number need not be continuous, but it will in many cases be measurable. We convolve this function with a Gaussian (or indeed any Lipschitz function) to produce a new Lipschitz function that carries the persistence-based information we desire.

Our main theoretical results (Theorems 4.1, 4.2, and 4.3) give conditions on functions h and K (where K will usually be a kernel) that guarantee that the convolution $h * K$ is Lipschitz with specified Lipschitz constant. From these we obtain the following, where more precise statements are given as Corollaries 4.4, 4.5, and 4.7.

THEOREM 1.1. *If h is locally essentially bounded then for the triangular and Epanechnikov kernels, $h * K$ is locally Lipschitz. If h is essentially bounded then for the Gaussian kernel, $h * K$ is Lipschitz.*

In practice, this can be translated to a simple procedure for stabilizing unstable persistent homology computations: perturb the input by adding, for example, Gaussian noise, and redo the computation; repeat and average. See Algorithm 1. By the law of large numbers, the result converges to the desired stable value.

THEOREM 1.2. *Let $\epsilon_1, \dots, \epsilon_M$ be drawn independently from a kernel K . Then*

$$\frac{1}{M} \sum_{i=1}^M h(\mathbf{a} - \epsilon_i) \rightarrow (h * K)(\mathbf{a}).$$

Algorithm 1 Stabilizing unstable persistence computations

Input: $h : \mathbb{R}^n \rightarrow \mathbb{R}$, $\mathbf{a} \in \mathbb{R}^n$

Parameters: $M \in \mathbb{N}$, $\sigma > 0$

for $i \leftarrow 1, M$ **do**

for $j \leftarrow 1, n$ **do**

 Sample ϵ_j from $N(0, \sigma^2)$

end for

$y_i \leftarrow h(\mathbf{a} + \epsilon)$, $\epsilon = (\epsilon_1, \dots, \epsilon_n)$

end for

return the average value of y_1, \dots, y_M

We summarize our computational pipeline in the following algorithm. Say we have performed a persistence computation and obtained an unstable output. For example, we have determined that the longest interval in the degree one bar code of the Vietoris-Rips complex on points $X_1, \dots, X_N \in \mathbb{R}^d$ is born with the addition of the edge X_1X_2 . We encode this output as a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with input $\mathbf{a} \in \mathbb{R}^n$. For example, the coordinates of the above points give us $\mathbf{a} \in \mathbb{R}^n$ where $n = Nd$. We define a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ whose value is the length of the longest interval in the bar code if it is born with the addition of the edge X_1X_2 and is otherwise 0.

The choice of standard deviation σ (also called bandwidth) is discussed in Sections 4.7 and 6.7. In Section 4.6, we prove that Algorithm 1 is stable with respect to this choice.

1.2. Three examples. For the reader familiar with persistent homology who wants to see how this works in practice, we provide three examples, the first and third to synthetic data and the second to brain imaging data. Code for these examples is available at <https://github.com/peter-bubenik/stabilizing-paper-code>.

EXAMPLE 1.3. *A cycle generating a persistent homology class*

We sample 1000 points uniformly from two conjoined annuli of inner and outer radii $(20, 50)$ and $(40, 50)$. Using Dionysus [40], we compute the 1-dimensional persistent homology of the alpha complex of our sample and obtain a representative cycle for the longest and second-longest bars. See Figure 1, left panel. However, the embedded location of these cycles is unstable. We would like to quantify the uncertainty of this location. To do so, we consider a square grid with edge-length 1. Our function $h : \mathbb{R}^{2000} \rightarrow \mathbb{R}$ has input the coordinates of the sampled points and has output 1 if the geometric cycle produced by Dionysus intersects a given square in our grid and otherwise has output 0.

We perturb the sampled points 10,000 times by adding Gaussian noise with standard deviation 3. For each square, we find the proportion of trials in which the representative

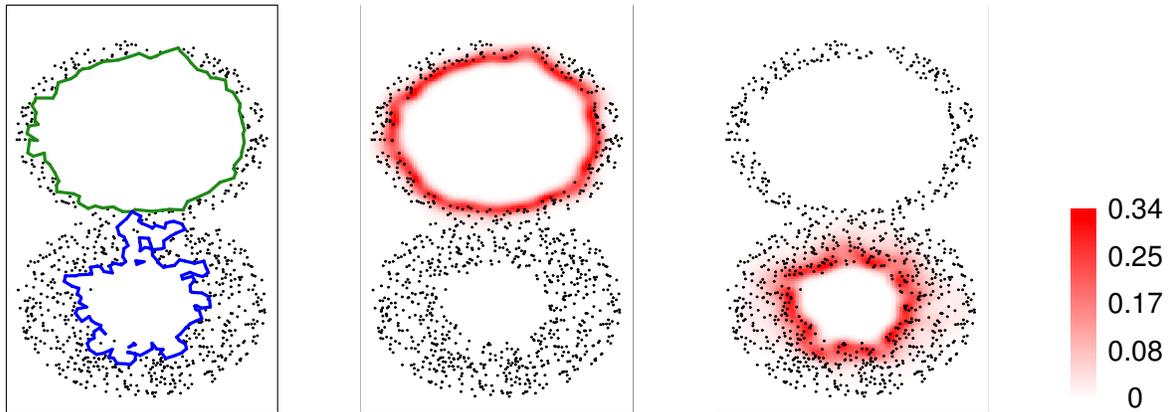


FIGURE 1. Finding the first and second longest bars. The first panel shows the original sample and the representative cycles produced by Dionysus [40] for the first and second longest bars. The latter two panels show the proportion of perturbations for which each square in a grid intersected the representative geometric cycle of the first or second longest bar, respectively.

geometric cycle for the longest or second-longest bar produced by Dionysus intersects the square. By performing this procedure simultaneously for every square in the grid, we obtain the second and third panels in Figure 1.

To see the effect of varying the choice of bandwidth, see Section 6.7.

EXAMPLE 1.4. *Location of a persistent homology generator in brain imaging data*

In [6], the authors apply topological data analysis to brain arteries extracted from magnetic resonance images. Mathematically, each of these brain arteries is a graph embedded in three-dimensional Euclidean space. Using the height (the z -coordinate) one obtains a filtration on this graph, which may be used to compute degree-zero persistent homology.

To facilitate statistical analysis of the resulting persistence diagrams (i.e. bar codes), they convert each persistence diagram to a vector consisting of the lengths of the 100 longest bars in decreasing order. In their analysis, the length of the 28th longest bar is a numerical feature that yields a correlation with age that is near-optimal among vector features consisting of the lengths of the i th through j th longest bars for any $1 \leq i \leq j \leq 100$.

If one wants to find a biological interpretation of this result, it is obvious to ask for the location of the generator of the 28th longest bar for each subject. It is easy to locate the generator responsible for the birth of the 28th longest bar. It will be a particular vertex of the graph, whose image is a point in space. However, the location of this point

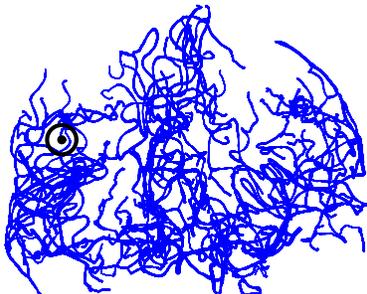


FIGURE 2. The brain arteries of the first subject in [6]. The black dot is the location of the generator of the 28th longest bar in degree-zero persistent homology. We consider the indicator function on the location of this generator with respect to the given sphere.

is unstable: as we will later explain, small perturbations of the spatial coordinates of the vertices of the graph can lead to large changes of this location.

We choose a ball centered at this point and consider the function whose value is 1 if the location of the generator of the 28th longest bar is located in this ball, and is otherwise 0. The resulting function $h : \mathbb{R}^{3V} \rightarrow \mathbb{R}$ (where V is the number of vertices in the graph) is unstable, but it may be stabilized using the method summarized in Section 1.1. Applying Algorithm 1 with $M = 1000$ and $\sigma = 0.1$ we obtain an estimate of the stable value of $h * K$ evaluated at the observed input, equal to 0.637. This shows that under small perturbations of the input, over half of the time the generator of the 28th longest bar is located in the chosen ball. This result holds for a large range of sizes of balls - see Section 6.8 for some further discussion.

We remark that this approach provides a resolution of the conflict between TDA theorists and TDA users expressed in the Fundamental Conundrum of Topological Data Analysis in the introduction. We can provide TDA users with a location of a generator of a persistent homology class together with an estimate of a stable real value of how often this location lies in a given region under certain perturbations.

EXAMPLE 1.5. *Persistence of a homology class born in a region*

Consider the function f on the square in Figure 3. This induces a function \bar{f} on the torus since $f(x, y) = 0$ on the boundary of the square. Suppose we are only given a finite sample of this induced function and we are interested in the presence of long-lived bars which are born in the region of the torus corresponding to the second quadrant of the square.

To be concrete, we start with a sample X of N points from the graph of \bar{f} , by sampling u_i, v_i independently from the uniform distribution on $[-\pi, \pi]$ and letting $z_i = f(u_i, v_i)$. Note that X is a random variable. We use X to construct a filtered simplicial complex approximating the unknown function \bar{f} as follows. From the points $\{(u_i, v_i)\}$ we construct a Delaunay triangulation of the torus. We filter this simplicial complex by assigning the

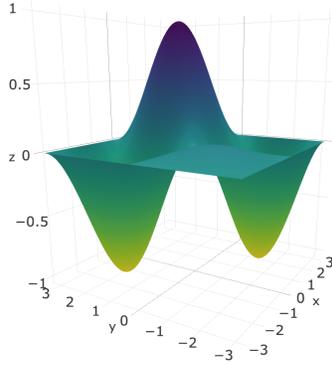


FIGURE 3. The graph of the function on the square $[-\pi, \pi]^2$ given by $f(u, v) = \sin(u) \sin(v)(1 - 0.9 \cdot 1_{\{u < 0, v < 0\}}) : [-\pi, \pi]^2 \rightarrow \mathbb{R}$. It induces a function on the torus, $\bar{f} : \mathbb{T}^2 \rightarrow \mathbb{R}$, with two global minima with value -1 , one global maximum with value 1 , one local maximum with value 0.1 , and four saddle points with value 0 . From

vertex (u_i, v_i) the value z_i and assigning edges and triangles the maximum value of their vertices.

We compute the 0-dimensional extended¹ persistence diagram of this filtered simplicial complex. Let $h(X)$ be the length of the longest bar if that bar was born in the region corresponding to the second quadrant (see Figure 3) and 0 otherwise.

This process defines a function $h : \mathbb{R}^{3N} \rightarrow \mathbb{R}$, but h is unstable. Consider the sample $X = x$ in Figure 4. We have $h(x) = 0$ since the global minimum, highlighted in red, is born outside the region corresponding to the second quadrant. Because of the symmetry of f , the random variable $h(X)$ is 0 approximately half the time and about 2 approximately half the time.

Let K denote the $3N$ -variate Gaussian with mean 0 and standard deviation 0.2. For $M \geq 1$, sample $\epsilon_1, \dots, \epsilon_M$ independently from K . Compute $\frac{1}{M} \sum_{i=1}^M h(x - \epsilon_i)$. See Figure 5. As M increases, this quantity converges to $g(x)$, where $g := h * K$ is the stabilized version of h .

1.3. Related Work. Partial inspiration for the main idea of our work (when faced with an instability caused by a near-interchange of values, perturb the values many times and take some sort of average) comes from the trembling-hand equilibrium solution [41] to the non-uniqueness problem for Fréchet means of persistence diagrams. Our approach should also be compared with the topological reconstruction results of Niyogi, Smale and Weinberger [44].

Several recent papers have advocated principled approaches for extracting features from persistence diagrams, including persistence landscapes [8], the stable multi-scale

¹Extended persistent homology follows the homology of increasing sublevel sets with the relative homology of the whole space relative to decreasing superlevel sets [25]. In the case considered here, it pairs the global minimum with the global maximum.

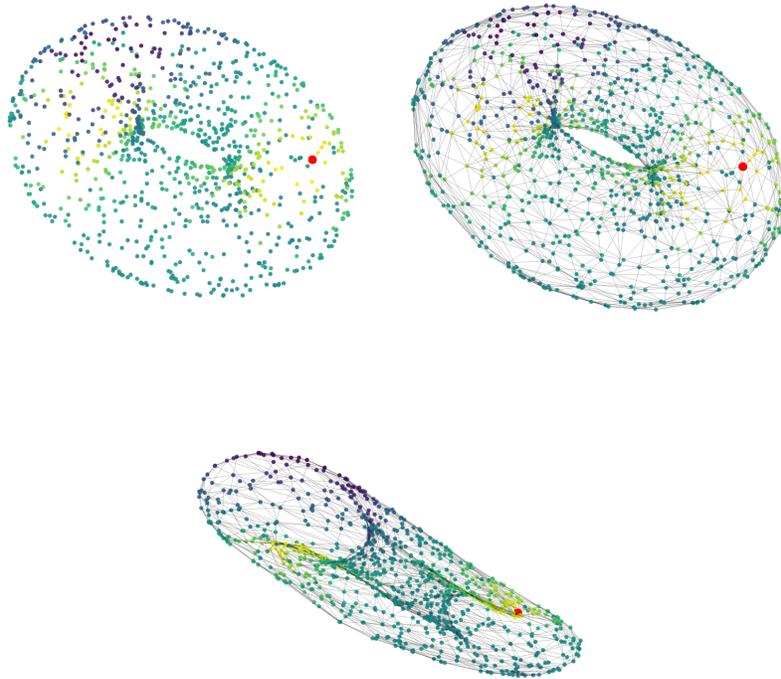


FIGURE 4. Sample of 1000 points from the graph $\{(x, f(x)) : x \in T^2\}$, where the function values are indicated using the same color scale as in Figure 3. The points on the torus are used to construct a Delaunay triangulation, which is filtered using the function values. On the right we indicate the filtration values by moving the points in the normal direction.

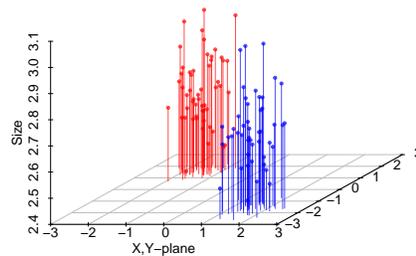


FIGURE 5. Locations and sizes of 100 longest bars from the trials. Averaging the lengths of the red bars over 1000 trials we get 1.291, which is consistent with the fact that the random variable $h(X)$ is 0 or about 2 with equal probability. We should not expect $\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M h(x + \epsilon_i)$ to converge to 1 because unlike f , a particular sample $X = x$ is not symmetric with respect to the second and fourth quadrants.

kernel [46], intensity functionals [21], persistence images [2], the stable topological signature [13], and the cover-tree entropy reduction [48]. Our result complements these ideas: once one identifies some specific parts of the persistence diagram as having good classification power, one can then attempt to locate, in a robust way, the portions of the domain responsible for these parts. Other papers (e.g. [15, 9, 1]) have developed sophisticated schemes for data-cleaning before persistent homology computation. These techniques are generally fragile to certain initial parameter choices, such as the m_0 parameter in [15]. Again, we provide a complementary role: any of these schemes can be run many times for several perturbations of an initial parameter choice, and the output can then be taken with confidence.

Dey and Wenger [29] have shown that the critical points of interval persistent homology are stable in the sense that they remain within some path-connected component.

Zomorodian and Carlsson [51] use Mayer-Vietoris as inspiration in their technique for localizing (relative to a cover) homology classes within a given simplicial complex. However, this works only for a fixed simplicial complex, not a simplicial complex endowed with a filtration, and the results are certainly fragile to changes in this fixed complex.

Weinberger [50] considers the sample complexity of some basic problems of topological inference. Specifically, he estimates the number of sample points necessary to determine the dimension, topological type, and to detect singularities for certain spaces.

Robust summaries of persistent homology are considered in the following papers; they do not consider the location of homology generators. In [7], Blumberg et al. show that persistent homology on a metric measure space induces a stable empirical measure in the space of persistence diagrams. Taking the distance to a reference distribution or a reference barcode, they obtain robust statistics. In [17], the authors derive limiting distributions and confidence sets for persistence diagrams based on the sub-level sets of the distance-to-a-measure.

Convolving with a kernel to obtain smoothness is a classical idea in statistics [47, 49]. It has been used to construct smooth estimators of discrete data as an initial step to computing persistent homology [9, 8, 33]. A related idea is the to perform subsampling (e.g. the bootstrap) to obtain convergence results and confidence intervals for persistence diagrams and persistence landscapes [33, 20, 18, 19]. These papers use ideas related to ones presented here, but to smooth initial data or to smooth stable outputs of persistence computations, not to stabilize unstable outputs of persistence computations.

1.4. Computational issues. Our methods (applying Theorem 4.8 in Algorithm 1) require repeated computation of persistence diagrams for similar filtrations. The computational cost may be considerable. In Examples 1.4 and 1.5 we repeat $M = 1000$ times. In Example 1.3 we repeat 10,000 times. Note that we do not have convergence results at this time. For repeated persistent homology calculations it is important to have efficient software. In Example 1.3 we use Dionysus [40], in Examples 1.4 and 5.1

we calculate persistent homology using a union-find data structure [32], in Example 1.5 we use Perseus [43], and in Example 5.2 we use Ripser [4].

However, our methods are trivially parallelizable. With access to many cores, our repeated computations can be computed in parallel without increasing the running time.

Note that for small perturbations, much of the persistent homology computation may be the same. In this case, there may be considerable computational savings by using vineyard updates [27].

Let us also remark that our methods combine nicely with subsampling, which is crucial for allowing persistent homology computations in the big data setting [18].

1.5. Replacing persistence diagrams with features. Our approach centers on converting a persistence diagram to a real number. This may seem simplistic and somewhat ad hoc. However, all effective methods of combining persistence diagrams with serious statistical analysis and machine learning techniques rely on replacing a persistence diagram with a vector in some Hilbert space or Banach space. For simplicity, we restrict ourselves to the vector space \mathbb{R} , but our approach can be extended to more general vector spaces.

1.6. Outline. Persistent homology computations and stability theorems are reviewed in Section 2, although we assume the reader is already somewhat familiar with them. Several examples of important but unstable persistence-based information are given in Section 3, and we then describe a general approach that stabilizes them in Section 4. In Section 5, we show how to apply these results to various persistent homology computations. Additional analyses and discussion are presented in Section 6. Potential future directions are discussed in Section 7.

2. Persistent Homology and Stability

The treatment of persistence diagrams here is adapted from [31]. For a more general discussion, see [45]. We assume the reader is familiar with the basics of homology groups: the textbook [42] is a good introduction. All homology groups are assumed to be computed over some fixed field. For concreteness, we restrict our attention to simplicial complexes, but our results also apply to more general complexes.

2.1. Persistent Homology. Persistent homology is computed for a finite *filtered abstract simplicial complex*. That is, we have a finite *abstract simplicial complex*, a collection, $K = \{\sigma\}$, of nonempty subsets of a fixed finite set that satisfy the condition that if $\emptyset \neq \tau \subseteq \sigma \in K$ then $\tau \in K$. In addition, we have a *filtration*, a function $f : K \rightarrow \mathbb{R}$ such that $\tau \subseteq \sigma$ then $f(\tau) \leq f(\sigma)$. That is, f is order preserving.

Fix a homological dimension p . Suppose the distinct values of f are $r_1 < \dots < r_m$. For each $1 \leq i \leq m$, define $K^i = \{\sigma \in K \mid f(\sigma) \leq r_i\}$. Since f is order preserving, each

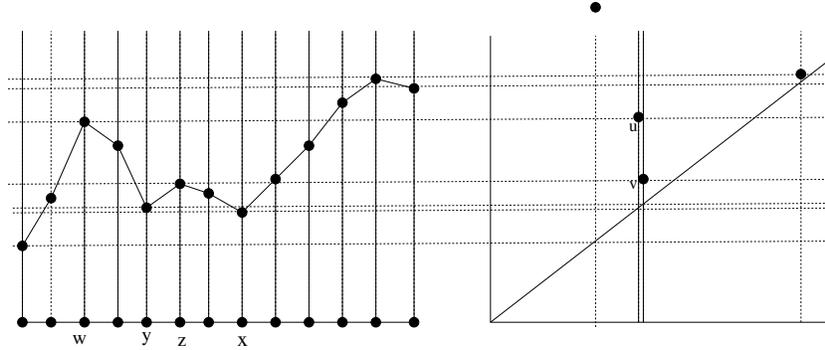


FIGURE 6. Left: The graph of a function F on a simplicial complex \mathcal{K} . Right: the degree-zero persistence diagram $\text{Dgm}_0(f)$ for the corresponding abstract simplicial complex K and filtration f . The labeled points have coordinates $u = (f(x), f(w))$ and $v = (f(y), f(z))$. The point on the very top has infinite y -coordinate.

K^i is a subcomplex. Whenever $i \leq j$, there is an inclusion $K^i \hookrightarrow K^j$, which induces a homomorphism:

$$f_p^{i,j} : H_p(K^i) \rightarrow H_p(K^j).$$

A homology class $\alpha \in H_p(K^i)$ is a *persistent homology class* that is *born* at level i if $\alpha \notin \text{im } f_p^{i-1,i}$, and that *dies* entering level j if $f_p^{i,j}(\alpha) = 0$ but $f_p^{i,j-1}(\alpha) \neq 0$. If α never dies, we say that it dies entering level $j = \infty$ and $r_\infty = \infty$. The *persistence* of α is defined to be $\text{pers}(\alpha) = r_j - r_i$. The set of classes which are born at i and die entering level j form a vector space, with rank denoted $\mu_p^{i,j}$. The degree- p *persistence diagram* of f , $\text{Dgm}_p(f)$, encodes these ranks. It is a multiset of points in the extended plane, with a point of multiplicity $\mu_p^{i,j}$ at each point (r_i, r_j) .

In practice, one constructs a filtered abstract simplicial complex from some other starting data. In addition, more information can be extracted from the persistent homology algorithm than just the persistence diagram.

We define a *persistent homology computation*, \mathcal{C} , to be a function whose input consists of real numbers a_1, a_2, \dots, a_n . These may include input values and also parameter values for the computation. Using this input, \mathcal{C} constructs an abstract simplicial complex K together with a filtration f . The output of \mathcal{C} consists of a degree- p persistence diagram together with for each (r_i, r_j) in the persistence diagram (counted with multiplicity), a p -simplex σ with $f(\sigma) = r_i$, a p -cycle α in K^i containing σ , a $(p+1)$ -simplex τ with $f(\tau) = r_j$, and a $(p+1)$ -chain β in K^j containing τ with $d\beta = \alpha$.

EXAMPLE 2.1. *Functions on simplicial complexes.* A filtered abstract simplicial complex, K , may be obtained from a real-valued function, F , on the vertices in a finite simplicial complex, \mathcal{K} . As a set $K \cong \mathcal{K}$. A filtration, f , on K is defined by $f(\sigma) = \sup_{x \in \sigma} F(x)$.

For example, let \mathcal{K} be the geometric line graph (i.e. an embedding of a graph - consisting of vertices and edges - in the plane), shown on the bottom of the left side of Figure 6. Above this, we have the graph of a function F on the points in \mathcal{K} . From this,

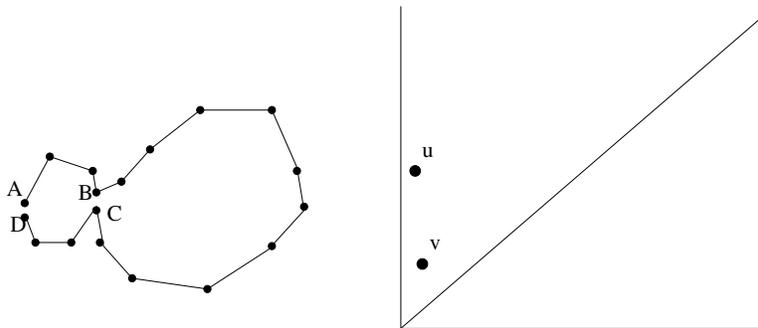


FIGURE 7. Left: a piecewise-linear curve in the plane. The distance between A and D is slightly smaller than the distance between B and C. Right: $\text{Dgm}_1(f)$, where f is as defined in the text. The points u and v correspond to one-cycles that are created by the additions of edges (A, D) and (B, C) , respectively.

we have a corresponding abstract simplicial complex K and filtration f . The persistence diagram $\text{Dgm}_0(f)$ is on the right. The input to \mathcal{C} consists of the function values (from left to right) a_1, a_2, \dots, a_n .

EXAMPLE 2.2. *Distance to a PL-Curve.* Consider the piecewise-linear curve C on the left side of Figure 7. Moving clockwise, we order its vertices $A = v_1, v_2, \dots, v_N = D$. Let K be the full simplex on these N vertices. For each vertex v , define $f(v) = 0$. For each edge of the form $e = (v_i, v_{i+1})$, define $f(e) = 0$, and for any other edge $e = (v_i, v_j)$, we set $f(e)$ to be the Euclidean distance between v_i and v_j . Finally, for any higher simplex σ , set $f(\sigma) = \max_{e \subseteq \sigma} f(e)$, where we take the maximum over the set of edges contained in σ . The degree-one persistence diagram $\text{Dgm}_1(f)$ appears on the right of Figure 7.

Here the input to \mathcal{C} consists of the $2n$ coordinates of the vertices. We note this paradigm can be extended to curves C in higher-dimensional ambient spaces, or even to higher-dimensional complexes.

EXAMPLE 2.3. *Point cloud – Vietoris-Rips.* Suppose that $X = \{x_1, \dots, x_N\}$ is a set of points in some metric space (Y, d) . We let K be the full simplex on these vertices. Define $f(v) = 0$ for each vertex and $f(e) = d(v, w)$ for each edge $e = (v, w)$. As above, we set $f(\sigma) = \max_{e \subseteq \sigma} f(e)$ for all higher-dimensional simplices. This is called the Vietoris-Rips filtration. We denote $\text{Dgm}_p(X) = \text{Dgm}_p(f)$. The input to \mathcal{C} consists of the coordinates of the points in X in some parametrization of Y . Alternatively, it consists of the entries of the distance matrix $D = (d(x_i, x_j))$.

For example, let X be the annular point cloud on the top-left of Figure 9. The corresponding $\text{Dgm}_1(X)$ appears on the top-right of the same figure.

EXAMPLE 2.4. *Point cloud – geometry.* Often, the simplicial complex in the previous example is too large to work with. Instead one applies some geometric ideas to construct a smaller filtered simplicial complex. Examples include witness complexes [28], the graph-induced complex [30], and the use of nudged elastic bands [1]. These constructions typically include one or more parameters, which we append to the input to \mathcal{C} .

EXAMPLE 2.5. *Point cloud – statistics.* Instead of using geometric ideas to construct a smaller point cloud we can use statistical ideas. For example, one can use a kernel to smooth the point cloud to obtain a density estimator on the underlying space Y and use this to filter a triangulation of Y [23, 9, 22]. Or one may use the local density to threshold the point cloud [11]; we consider this in more detail in Example 5.4. Again, these constructions include one or more parameters, which we append to the input for \mathcal{C} .

EXAMPLE 2.6. *Regression.* Here we present a variant of Example 2.1 in which we are not given the simplicial complex. Instead we sample points $X = (x_1, \dots, x_N)$, $x_i \in \mathbb{R}^d$ from some probability distribution on \mathbb{R}^d . We also sample corresponding perturbed function values $y_i \in \mathbb{R}$. For example, we may have $y_i = f(x_i) + \epsilon_i$, where ϵ_i is sampled from a univariate Gaussian. We use X to construct a Delaunay triangulation K . We then use $Y = (y_1, \dots, y_N)$ to filter K as follows: $f(\sigma) = \max_{x_i \in \sigma} y_i$. This is called the lower star filtration.

Instead of the sample points lying in \mathbb{R}^d , they may lie on some compact Riemannian manifold. See the torus example in Section 1.

Instead of filtering K directly using Y , one can instead use (X, Y) to construct an estimator \hat{f} of the unknown regression function f . We can then use \hat{f} to filter K [9].

2.2. Stability. The persistence diagram $\text{Dgm}_p(f)$ is a summary of the function f , and it turns out to be a stable one. The discussion here is adapted from [24]. For a broader description, see [14, 16].

For convenience, to each persistence diagram, we add every point (r, r) on the major diagonal, each with infinite multiplicity.

Now suppose that $\phi : D \rightarrow D'$ is some bijection between two persistence diagrams; bijections exist because of the infinite-multiplicity points along the diagonal. The cost of ϕ is defined to be $C(\phi) = \sup_{u \in D} \|u - \phi(u)\|_\infty$; that is, the largest box-norm distance between matched points. The *bottleneck distance* $W_\infty(D, D')$ is defined to be the minimum cost amongst all such bijections. For example, if D and D' are the black and red diagrams, respectively, on the right side of Figure 8, then the best bijection would pair u with u' , v with v' , the two infinite-persistence points with each other, and the other two points with the closest diagonal points. The bottleneck distance is the cost of this bijection. The Diagram Stability Theorem [24] guarantees that persistence diagrams of nearby functions are close to one another. More precisely, we have $W_\infty(D_p(f), D_p(g)) \leq \|f - g\|_\infty$. This is illustrated by Figure 8.

Note that the difference between f and g is measured in the L_∞ norm. In the point cloud context (Ex. 2.3), this translates into requiring that the two point cloud inputs be Hausdorff-close. However, the persistence diagram is not stable with respect to the addition of outliers. We discuss this problem in more detail in Section 3.2 and propose a solution in Section 4.

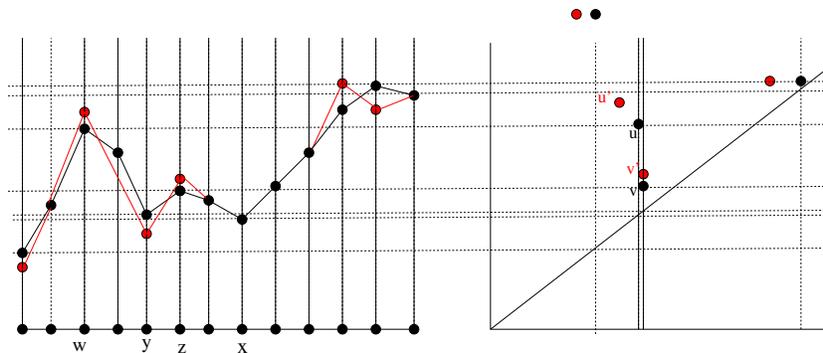


FIGURE 8. Left: The graphs of functions f (black) and g (red), both on the same domain \mathcal{K} . Right: the persistence diagrams $\text{Dgm}_0(f)$ and $\text{Dgm}_0(g)$, using the same color scheme.

3. Instability

The Diagram Stability Theorem tells us that the persistence diagram obtained in the output of a persistent homology computation is stable with respect to certain perturbations of the input used to construct a filtered abstract simplicial complex. However, other outputs of persistent homology computations are not stable. This includes the simplices and cycles that generate persistent homology classes. These are of great interest to practitioners hoping to interpret persistence calculations more directly. In addition, many persistence computations rely on choices of parameters and the resulting persistence diagrams may be unstable with respect to these choices.

3.1. Instability of Generating Cycles/Simplices. Persistence diagrams are useful and robust measures of the *size* of topological features. What they are less good at, on the other hand, is robustly pinpointing the *location* of important topological features. We use Figure 8 to illustrate this problem. Suppose that we have the fixed domain \mathcal{K} and we observe the function f . One of the most prominent points in $\text{Dgm}_0(f)$ is u , which corresponds to the pair of values $f(x)$ and $f(w)$. We might thus be tempted to say that f has an important feature, a component of high-persistence, *at* x . But consider the nearby function g instead. Its diagram $\text{Dgm}_0(g)$ has a point u' that is very close to u , but this point corresponds to the pair of values $f(y)$ and $f(w)$. There is still a component born at $g(x)$, but it corresponds to the much smaller persistence point v' . And so while the persistence of the point u is a stable summary of the function f , the actual location x of the topological feature it corresponds to is not.

This is unfortunate. Several recent works ([5], [6], among others) have shown that the presence of points in certain regions of the persistence diagram has strong correlation with covariates under study. For example, each diagram in the second cited work came from a filtration of the brain artery tree in a specific patient's brain, and it was found that the density of points in a certain middle-persistence range gave strong correlations with patient age. It would of course be tempting to hold specific locations in the brain responsible for these points with high distinguishing power.

Unsurprisingly, this problem remains for persistent homology in higher degrees. Consider Figure 7 again. It is easy to see that edge (A, D) creates the large loop which corresponds to point $u \in \text{Dgm}_1(f)$. However, a slight perturbation of the vertex configuration could render (B, C) responsible for this loop instead, and so we cannot robustly locate the persistence of this loop at (A, D) .

In Section 4, we both rigorously define this non-robustness and give a method for addressing it.

3.2. Outliers and Instability of Parameter Choices. The Diagram Stability Theorem guarantees the persistence diagrams associated to two Hausdorff-close point clouds will themselves be close. However, it says nothing about the outlier problem. For example, consider again the point cloud X (Figure 9, top-left) from Example 2.3 to which we apply the Vietoris-Rips construction. Its persistence diagram $\text{Dgm}_1(X)$ (top-right of same figure) has one high-persistence point, which corresponds to the “circle” that we qualitatively see when looking at the points. On the other hand, consider the point cloud X' on the bottom-left, which consists of X and three “outlier” points spread across the interior of the circle. The diagram $\text{Dgm}_1(X')$ (bottom-right) is not close to $\text{Dgm}_1(X)$: there is still one point of fairly high persistence, but it’s much closer to the diagonal than before.

In practice, this problem is often addressed by first de-noising the point cloud in some way. For example, Carlsson et. al. [12] first thresholded by density before computing Vietoris-Rips filtrations when they discovered a Klein bottle in the space of natural images. There are no guarantees that a different, nearby choice of density threshold parameter would not give a qualitatively different persistence diagram. Section 4 addresses this by introducing a general method for handling parameter choice in persistence computations.

4. Theory: Stability from convolutions

In this section we show how functions may be stabilized by convolving them with a kernel. In Section 5, we will apply these results to the function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ discussed in the introduction. First, we give three general results with various assumptions on the function and the kernel. Next, we apply them in three particular cases: the simple triangular kernel and the commonly used Epanechnikov and Gaussian kernels. We then outline a few specific examples, some of which will be explored via experiment in the next section.

4.1. Lipschitz functions and convolution. Let us start by recalling a few definitions. For $C \geq 0$, a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be C -Lipschitz if for all $u, v \in \mathbb{R}^n$, $|f(u) - f(v)| \leq C|u - v|$, where $|x|$ denotes the Euclidean norm. We will call a function *Lipschitz* if it is C -Lipschitz for some $C \geq 0$. The support of f , denoted $\text{supp}(f)$, is the closure of the subset of \mathbb{R}^n where f is non-zero.

Let $h, g : \mathbb{R}^n \rightarrow \mathbb{R}$ be (Lebesgue) measurable functions that are defined almost everywhere. The 1 -norm of h , is given by $\|h\|_1 = \int_{\mathbb{R}^n} |h(t)| dt$, if it exists. The *essential supremum*

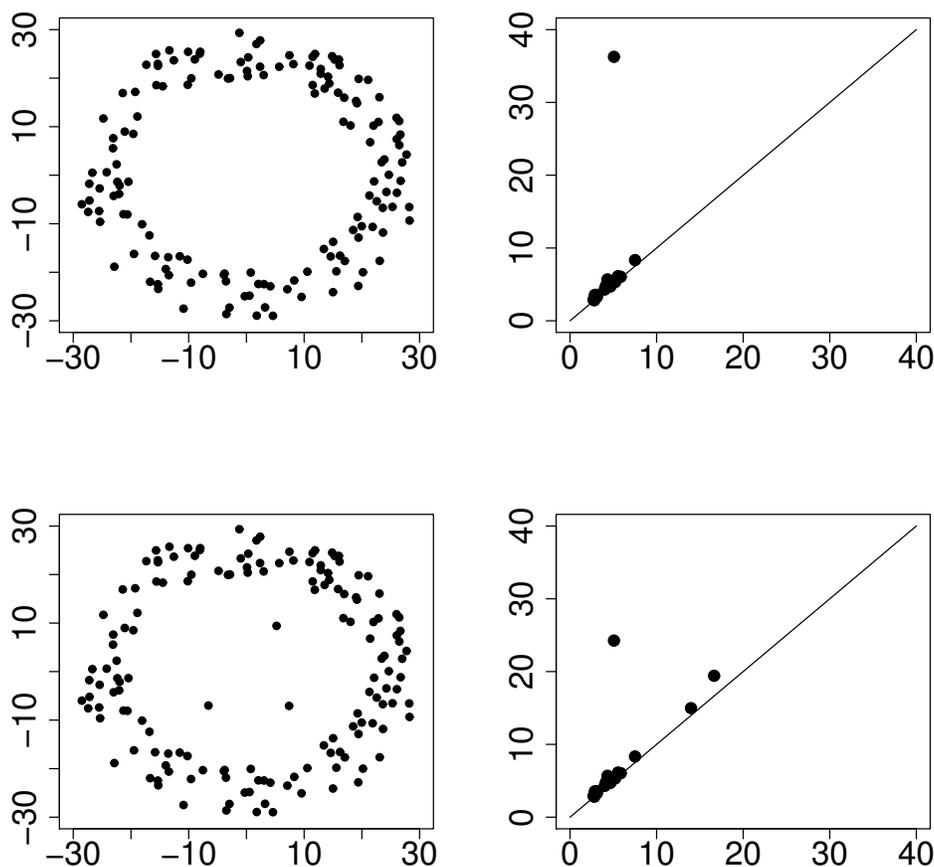


FIGURE 9. Illustration of the outlier problem for the persistent homology of the Vietoris-Rips complex of a point cloud. Top left: 150 points X , sampled from an annulus. Top right: $Dgm_1(X)$. Bottom left: 153 points X' , which is X plus three outlier points. Bottom right: $Dgm_1(X')$.

of h , denoted by $\|h\|_\infty$, is the smallest number a such that the set $\{x \mid |f(x)| > a\}$ has measure 0. If it exists, the *convolution product* of h and g , is given by

$$(h * g)(t) = \int_{\mathbb{R}^n} h(s)g(t-s)ds = \int_{\mathbb{R}^n} h(t-s)g(s)ds.$$

It exists everywhere, for example, if one function is essentially bounded and the other is integrable; or if one function is bounded and compactly supported and the other is locally integrable [34, Section 473D].

Throughout this section we assume that $h : \mathbb{R}^d \rightarrow \mathbb{R}$ is defined almost everywhere, $K : \mathbb{R}^d \rightarrow \mathbb{R}$ and that that the convolution product $h * K$ exists almost everywhere.

4.2. Stability theorems. We now give several conditions on a pair of functions which imply that their convolution product is (locally) Lipschitz.

The first result appears in [34, 473D(d)], but the proof is included here for completeness.

THEOREM 4.1. *If $\|h\|_1 = a$ and K is b -Lipschitz, then $h * K$ is ab -Lipschitz.*

PROOF. Let $g = h * K$. First we have, $g(u) - g(v) = \int_{\mathbb{R}^n} h(s) (K(u-s) - K(v-s)) ds$. Then, $|g(u) - g(v)| \leq \int_{\mathbb{R}^n} |h(s)| |K(u-s) - K(v-s)| ds \leq \int_{\mathbb{R}^n} |h(s)| b|u-v| ds \leq ab|u-v|$. \square

Let $B_\alpha(x)$ denote the closed ball of radius α centered at $x \in \mathbb{R}^d$, and let V_d denote the volume of the d -dimensional ball of radius 1.

THEOREM 4.2. *Let $x \in \mathbb{R}^d$ and let $\alpha > 0$. If $\|h\|_\infty \leq M$ on $B_{2\alpha}(x)$, K is b -Lipschitz and $\text{supp}(K) \subseteq B_\alpha(0)$, then $h * K$ is $2Mb\alpha^d V_d$ -Lipschitz in $B_\alpha(x)$.*

PROOF. Let $g = h * K$. Let $u, v \in B_\alpha(x)$. As in the previous proof, $|g(u) - g(v)| \leq \int_{\mathbb{R}^n} |h(s)| |K(u-s) - K(v-s)| ds \leq \int_{B_\alpha(u) \cup B_\alpha(v)} |h(s)| b|u-v| dx \leq 2Mb\alpha^d V_d |u-v|$. \square

THEOREM 4.3. *If $\|h\|_\infty \leq M$ and $\int |K(s+t) - K(s)| ds \leq b|t|$ for all $t \in \mathbb{R}^d$, then $h * K$ is Mb -Lipschitz.*

PROOF. Let $g = h * K$. Again, $|g(u) - g(v)| \leq \int_{\mathbb{R}^n} |h(s)| |K(u-s) - K(v-s)| ds \leq \int M |K(u-v+s) - K(s)| ds \leq Mb|u-v|$. \square

4.3. Application to kernels. We now apply the above theorems to smooth a function h , obtaining a Lipschitz function. That is, we will take K to be a *kernel*, a non-negative integrable real-valued function on \mathbb{R}^n satisfying $\int K(x) dx = 1$, $\int xK(x) dx = 0$ and $\int x^2 K(x) dx < \infty$. For example, we can choose K to be the *triangular kernel*, $K(x) = c \max(1 - \|x\|, 0)$, for appropriate normalization constant c (see Figure 10). The most common choices are the Gaussian kernel and the Epanechnikov kernel, which are described below (see Figure 10). Notice that if K is a kernel, then so is $K_\alpha(x) = \frac{1}{\alpha^n} K(\frac{x}{\alpha})$.² The parameter α is called the *bandwidth* and allows one to control the amount of smoothing.

The triangular kernel. Let $\alpha > 0$. Let V_d denote the volume of the n -dimensional ball of radius 1. For $A \subseteq \mathbb{R}^d$, let I_A denote the indicator function on A . That is, $I_A(x) = 1$ if $x \in A$ and 0 otherwise. The *triangular kernel* is given by

$$K_\alpha(x) = \frac{d+1}{\alpha^d V_d} \left(1 - \frac{|x|}{\alpha}\right) I_{B_\alpha(0)}.$$

Note that $\text{supp}(K_\alpha) = B_\alpha(0)$ and K_α is $\frac{d+1}{\alpha^{d+1} V_d}$ -Lipschitz. Applying Theorem 4.2, we have the following.

²More generally, we can choose the bandwidth to be a symmetric positive definite matrix H and let $K_H(x) = \frac{1}{\sqrt{\det H}} K(H^{-1/2}x)$.

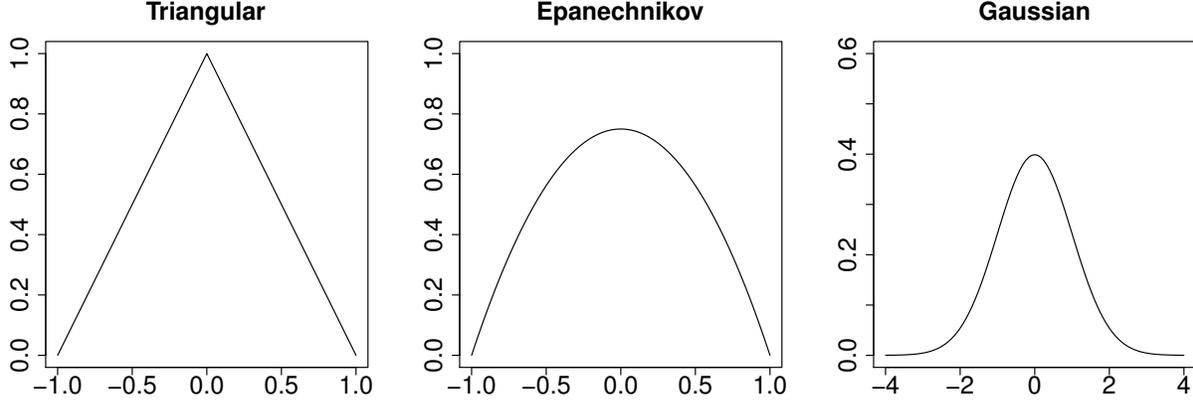


FIGURE 10. Graphs of three common kernels.

COROLLARY 4.4. *Let $x \in \mathbb{R}^d$. If $\|h\|_\infty \leq M$ on $B_{2\alpha}(x)$ then $h * K_\alpha$ is $\frac{2M(d+1)}{\alpha}$ -Lipschitz in $B_\alpha(x)$.*

Note that it follows that if the bound on h is global then so is the Lipschitz bound.

The Epanechnikov kernel. Let $\alpha > 0$. The Epanechnikov kernel is given by

$$K_\alpha(x) = \frac{d+2}{2\alpha^d V_d} \left(1 - \frac{|x|^2}{\alpha^2}\right) I_{B_\alpha(0)}.$$

Now $\text{supp}(K_\alpha) = B_\alpha(0)$ and K_α is $\frac{d+2}{\alpha^{d+1} V_d}$ -Lipschitz. Applying Theorem 4.2, we have the following.

COROLLARY 4.5. *Let $x \in \mathbb{R}^d$. If $\|h\|_\infty \leq M$ on $B_{2\alpha}(x)$ then $h * K_\alpha$ is $\frac{2M(d+2)}{\alpha}$ -Lipschitz in $B_\alpha(x)$.*

The Gaussian kernel. Let $\alpha > 0$. The Gaussian kernel is given by

$$K_\alpha(x) = \frac{1}{\alpha^d (2\pi)^{d/2}} e^{-|x|^2/2\alpha^2}.$$

LEMMA 4.6. *For the Gaussian kernel K_α , let $f(t) = \int |K_\alpha(s+t) - K_\alpha(s)| ds$. Then $f(t) \leq \frac{2}{\alpha\sqrt{2\pi}}|t|$ for all $t \in \mathbb{R}^d$.*

PROOF. Change coordinates so that $s = -\frac{|t|}{2}e_1$ and $s + t = \frac{|t|}{2}e_1$. Then by symmetry

$$\begin{aligned} f(t) &= 2 \left[\int_{x_1 \geq -\frac{|t|}{2}} K_\alpha(x) \, dx - \int_{x_1 \geq \frac{|t|}{2}} K_\alpha(x) \, dx \right] \\ &= 4 \int_{0 \leq x_1 \leq \frac{|t|}{2}} K_\alpha(x) \, dx \\ &= \frac{4}{\alpha^d (2\pi)^{d/2}} \int_0^{\frac{|t|}{2}} e^{-x_1^2/2\alpha^2} \, dx_1 \int_{-\infty}^{\infty} e^{-x_2^2/2\alpha^2} \, dx_2 \cdots \int_{-\infty}^{\infty} e^{-x_d^2/2\alpha^2} \, dx_d \\ &= \frac{4}{\alpha\sqrt{2\pi}} \int_0^{\frac{|t|}{2}} e^{-x_1^2/2\alpha^2} \, dx_1 \end{aligned}$$

It follows that $f(t) \leq \frac{4}{\alpha\sqrt{2\pi}} \int_0^{|t|/2} dx_1 = \frac{2}{\alpha\sqrt{2\pi}}|t|$. \square

Thus by Theorem 4.3 we have the following

COROLLARY 4.7. *If $\|h\|_\infty \leq M$ then $h * K_\alpha$ is $\frac{2M}{\alpha\sqrt{2\pi}}$ -Lipschitz.*

In practice, the function h will rarely be essentially bounded, but this can be arranged by setting it to be 0 outside a closed ball centered at a specified configuration.

4.4. Sharpness of the Lipschitz constants. Assume that $K_\alpha : \mathbb{R}^d \rightarrow \mathbb{R}$ is symmetric in the first variable. Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be defined by $h(x) = 1$ if $x_1 \geq 0$ and -1 otherwise. Let $g(t) = h * K_\alpha(te_1) - h * K_\alpha(-te_1)$. Then we calculate

$$\begin{aligned} h * K_\alpha(te_1) &= \int_{\mathbb{R}^d} h(te_1 - x) K_\alpha(x) \, dx \\ &= \int_{x_1 \leq t} K_\alpha(x) \, dx - \int_{x_1 > t} K_\alpha(x) \, dx \\ &= \text{sign}(t) \int_{-|t| \leq x_1 \leq |t|} K_\alpha(x) \, dx. \end{aligned}$$

So

$$g(t) = 2 \text{sign}(t) \int_{-|t| \leq x_1 \leq |t|} K_\alpha(x) \, dx = 4 \text{sign}(t) \int_{0 \leq x_1 \leq |t|} K_\alpha(x) \, dx.$$

Let K_α be the Gaussian kernel. Then

$$\begin{aligned} g(t) &= \frac{4}{\alpha^d (2\pi)^{d/2}} \int_0^t e^{-x_1^2/2\alpha^2} \, dx_1 \int_{-\infty}^{\infty} e^{-x_2^2/2\alpha^2} \, dx_2 \cdots \int_{-\infty}^{\infty} e^{-x_d^2/2\alpha^2} \, dx_d \\ &= \frac{4}{\alpha\sqrt{2\pi}} \int_0^t e^{-x_1^2/2\alpha^2} \, dx_1 \end{aligned}$$

It follows that $g(t)$ converges to $\frac{4t}{\alpha\sqrt{2\pi}}$ as t approaches 0 by the first fundamental theorem of calculus. Hence, the Lipschitz constant given in Corollary 4.7 is optimal.

When $d = 1$ and K_α is the triangular kernel, $g(t) = \frac{4.2}{\alpha V_1} \int_0^t (1 - \frac{|x|}{\alpha}) dx \rightarrow \frac{4t}{\alpha}$ as $t \rightarrow 0$. So the Lipschitz constant of $h * K_\alpha$ is at least $\frac{2}{\alpha}$. Hence, the Lipschitz constant given in Corollary 4.4 is optimal up to at most a factor of 2.

When $d = 1$ and K_α is the Epanechnikov kernel, $g(t) = \frac{4.3}{2\alpha V_1} \int_0^t (1 - \frac{x^2}{\alpha^2}) dx \rightarrow \frac{3t}{\alpha}$ as $t \rightarrow 0$. So the Lipschitz constant of $h * K_\alpha$ is at least $\frac{3}{2\alpha}$. Hence, the Lipschitz constant given in Corollary 4.5 is optimal up to at most a factor of 4.

4.5. Stable Computations in Practice. Suppose that we can compute $h(x)$ for values of x for which it is defined, we can sample from K , and that for a fixed $a \in \mathbb{R}^d$ we want to compute $g(a) = (h * K)(a) = \int_{\mathbb{R}^d} h(a - x)K(x)dx$. In practice, we will not be able to evaluate this integral analytically. We approximate $g(a)$ as follows. Let V be a random variable with probability distribution given by the kernel K (one writes $V \sim K$). Let W be the random variable given by $h(a - V)$. Then the expected value of W is given by $E[W] = \int_{\mathbb{R}^d} h(a - x)K(x)dx = g(a)$. We will approximate $E[W]$ by drawing a sample $\epsilon_1, \dots, \epsilon_M$ where $\epsilon_i \sim K$ are independent. Then $E[W]$ can be approximated by $\bar{W}_M = \frac{1}{M} \sum_{i=1}^M h(a - \epsilon_i)$. By the law of large numbers, $\bar{W}_M \rightarrow E[W]$, where the convergence may be taken to be in probability (the weak law) or almost surely (the strong law). This is the justification for the computations in Section 6. Let us record this result.

THEOREM 4.8. *Let $a \in \mathbb{R}^d$ and $\epsilon_1, \dots, \epsilon_M$ be drawn independently from K . Then*

$$\frac{1}{M} \sum_{i=1}^M h(a - \epsilon_i) \rightarrow g(a).$$

4.6. Stability of the choice of kernel. As should be clear, and as borne out by the experiments in Section 6, the value of $(h * K)(a)$, for fixed h and a , will certainly depend on K . However, there is no fragility of output with respect to this choice, as shown by the following fact.

THEOREM 4.9. *Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be an essentially bounded function. Then the map $K \rightarrow h * K$ is Lipschitz, from $L^1(\mathbb{R}^d)$ to $L^\infty(\mathbb{R}^d)$.*

PROOF. Let $\phi : L^1(\mathbb{R}^d) \rightarrow L^\infty(\mathbb{R}^d)$ be given by $\phi(K) = h * K$. For $x \in \mathbb{R}^d$,

$$|[\phi(K) - \phi(K')](x)| \leq \int |h(x - t)| |K(t) - K'(t)| dt \leq \|h\|_\infty \|K - K'\|_1.$$

□

4.7. Bandwidth selection – theoretical considerations. After choosing a family of kernels, such as the Gaussian kernels K_α described in Section 4.3, the most important choice in implementing the method described here is the choice of bandwidth α .

Choosing the amount of smoothing is a well-studied problem in nonparametric regression, where increasing the bandwidth decreases the estimation variance, but increases the squared bias. Both of these terms contribute to the error. A bandwidth which optimizes this trade-off may be estimated using cross-validation. A proper understanding of this problem in our situation requires analysis that goes beyond the scope of the present paper.

However, we offer some heuristics for the choice of bandwidth. First, it may be chosen to obtain a desired amount of smoothness of $h * K_\alpha$. For example, we may want $h * K_\alpha$ to be 1-Lipschitz. Second, it seems reasonable to choose the bandwidth to (at least) equal the level of estimated noise of the input data. One may combine these two to find the minimum bandwidth that satisfies both requirements.

5. Application to persistent homology computations

Now let us apply the results of the previous section to persistent homology. Assume we have a persistent homology computation, \mathcal{C} , with input the real numbers a_1, \dots, a_n . If our computation is defined for all $a = (a_1, \dots, a_n) \in \mathbb{R}^d$ and results in real values then we may proceed.

If not, we may reduce the more general situation to the one above as follows. We need h to be defined on all of \mathbb{R}^n so that convolutions with a Gaussian are well defined. Let \mathcal{O} be the set of outputs of this computation. Let $D \subseteq \mathbb{R}^n$ be the set of all inputs for which \mathcal{C} is defined. If $D \neq \mathbb{R}^n$ then add a state \emptyset to \mathcal{O} and say that the computation sends all points in $\mathbb{R}^n - D$ to \emptyset . Thus we encode this computation as a function $H : \mathbb{R}^n \rightarrow \mathcal{O}$. Let p be a real-valued function on \mathcal{O} with $p(\emptyset) = 0$. Let $h = p \circ H : \mathbb{R}^n \rightarrow \mathbb{R}$. We will need h to be (Lebesgue) measurable.

To make this less abstract, we show how the instabilities described in Sections 3.1 and 3.2 can be addressed by this method.

EXAMPLE 5.1. *Stable persistence located at a point.* We return to Example 2.1, where we have a geometric line graph \mathcal{K} with n vertices v_1, \dots, v_n , and edges $e_i = (v_i, v_{i+1})$ for $i = 1, \dots, n - 1$. To produce a filtration of the type used in this example, we just need to know n function values. More precisely, our persistence computation takes as input a vector $a = (a_1, \dots, a_n) \in \mathbb{R}^n$, from which we obtain a piecewise linear function, F_a , on \mathcal{K} determined by $F_a(v_i) = a_i$. Next we consider the corresponding abstract simplicial complex K and filtration f_a . Then we compute the persistence diagram $\text{Dgm}_0(f_a)$. This defines a function $H : \mathbb{R}^n \rightarrow \mathcal{O}$, where $H(a) = \text{Dgm}_0(f_a)$.

Now fix a specific vertex x in K . A given diagram $\text{Dgm}_0(f_a)$ either contains a point $u(x) = (b(x), d(x))$ that represents a persistent connected component born at x in the filtration, or it does not. In the former case, we define $p_x(\text{Dgm}_0(f_a)) = d(x) - b(x)$, and in the latter we define $p(x)(\text{Dgm}_0(f_a)) = 0$; that is, we map the diagram to the persistence of the connected component created by the addition of this specific vertex. Note that whether or not p_x is non-zero depends on whether or not x is a local minimum.

The discontinuity of the function $h_x = p_x \circ H : \mathbb{R}^n \rightarrow \mathbb{R}$ expresses the instability of localizing the persistence of a connected component. Referring to Figure 8, suppose that the vectors a and e produce the functions f and g , respectively, and that the vertex x is as marked in the figure. Then $h_x(a)$ is the persistence of u , while $h_x(e)$ is the persistence of v' . Corollaries 4.4 and 4.5 guarantee that smoothing h_x by a Triangular kernel and Epanechnikov kernel will result in a locally Lipschitz function.

To be able to convolve with the Gaussian kernel and apply Corollary 4.7 we need h_σ to be essentially bounded. We can arrange this by specifying that the domain D of \mathcal{C} be compact and that p_σ be bounded. This requires that all of the persistence pairs in the output of \mathcal{C} be finite. This can be arranged by truncating at some value M or by applying extended persistence [25]. The resulting h_x is Lipschitz.

The experiments in Section 6 show how this works in practice.

EXAMPLE 5.2. *Stable persistence located at an edge.* We return to Example 2.2. In this case, K is the full complex on n vertices, and we start with n ordered points in the plane which lead to a piecewise-linear curve C . That is, \mathcal{C} takes as input a vector $a \in \mathbb{R}^{2n}$ and places a vertex v_i at (a_{2i-1}, a_{2i}) , thus creating a curve C_a . This leads to a filtration f_a of K and finally we produce $\text{Dgm}_1(f_a) \in \mathcal{O}$. As before, $H(a) = \text{Dgm}_1(f_a)$ defines a function $H : \mathbb{R}^{2n} \rightarrow \mathcal{O}$.

If we fix a specific edge σ , we can proceed as in Example 5.1 by defining the function p_σ and thus $h_\sigma = p_\sigma \circ H$. For example, taking $\sigma = (A, D)$ in Figure 7 and letting a be the vector which led to that specific point configuration, we have $h_\sigma(a)$ equal to the persistence of u . As above, $g_\sigma = h_\sigma * K_\alpha$ is (locally) Lipschitz.

EXAMPLE 5.3. *Stable persistence of generating cycles.* Instead of tracking which j -simplex creates a persistent homology class, a persistent homology algorithm may record a j -cycle, γ , that represents the persistence class. In this case, we can define $p_\gamma : \mathcal{O} \rightarrow \mathbb{R}$ to be $d - b$ if γ represents a persistence pair $[b, d)$ or otherwise 0. Let $h_\gamma = p_\gamma H$ and then $g_\gamma = h_\gamma * K_\alpha$ is (locally) Lipschitz.

EXAMPLE 5.4. *Stability in density-thresholding choice.* Let Y be the point cloud on the bottom-left of Figure 9, which we recall was created from the point cloud on the top-left by adding three outlier points. Consider any de-noising process parametrized by some real numbers. For a specific example, let $k = (\delta, \epsilon)$. For each $y \in Y$, let $C_\delta(y) = \{x \in Y \mid \|x - y\| \leq \delta\}$. Then define

$$Y_\epsilon^\delta = \{y \in Y \mid \frac{|C_\delta(y)|}{|Y|} \geq \epsilon\}.$$

One then applies the Vietoris-Rips construction to obtain a filtered abstract simplicial complex from Y_ϵ^δ , and then computes $\text{Dgm}_1(Y_\epsilon^\delta)$. We may consider the input of our persistent homology computation \mathcal{C} to be $a_1, \dots, a_{2n}, \delta, \epsilon$: that is, the coordinates of the vertices and the parameter values. However, we may also take the coordinates to be fixed and only consider the parameters to be our input. Doing this, we obtain $H : \mathbb{R}^2 \rightarrow \mathcal{O}$. In this case, define $p(D) = \max_{u \in D} \text{pers}(u)$ for any degree-one diagram D . Then the discontinuity of the function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$k = (\delta, \epsilon) \mapsto \text{Dgm}_1(Y_\epsilon^\delta) \mapsto p(\text{Dgm}_1(Y_\epsilon^\delta))$$

expresses the instability of the threshold-parameter choice referred to in Section 3.2. If k is chosen so that the three outlier points are removed, then $h(k)$ will be the persistence of the most prominent point on the top-right of Figure 9. On the other hand, a very nearby choice of k might fail to remove these points, and we would get the persistence of the most prominent point on the bottom-right of Figure 9. As above, $g_\sigma = h_\sigma * K_\alpha$ is (locally) Lipschitz.

6. Further analysis and discussion

This section more deeply investigates some of the examples above and some related practical issues.

6.1. First line-graph experiment. First we explore Example 5.1, where the input to a persistent homology computation is a choice of function-values on the vertices of a simplicial complex. Specifically, we consider a line graph \mathcal{K} with vertices v_1, \dots, v_7 , and the initial input choice $\mathbf{a} = (10, 11, 12.5, 13, 9.9, 20, 1)$. The left side of Figure 11 shows the graph of the PL-function $F_{\mathbf{a}}$, and the persistence diagram $H(\mathbf{a})$ is in the middle. Ignoring the highest-persistence point $(1, 20)$, the high-persistence point $(9.9, 20)$ and the medium-persistence one $(10, 13)$ are created by the additions of v_5 and v_1 , respectively; that is, $h_5(\mathbf{a}) = 10.01$ and $h_1(\mathbf{a}) = 3$. These values are of course unstable to perturbations of \mathbf{a} : for instance, if we switch the first and fifth entries of \mathbf{a} , the reader can check that $h_5((9.9, 11, 12.5, 13, 10, 20, 1)) = 3$. Let K_α be a seven-dimensional Gaussian kernel with

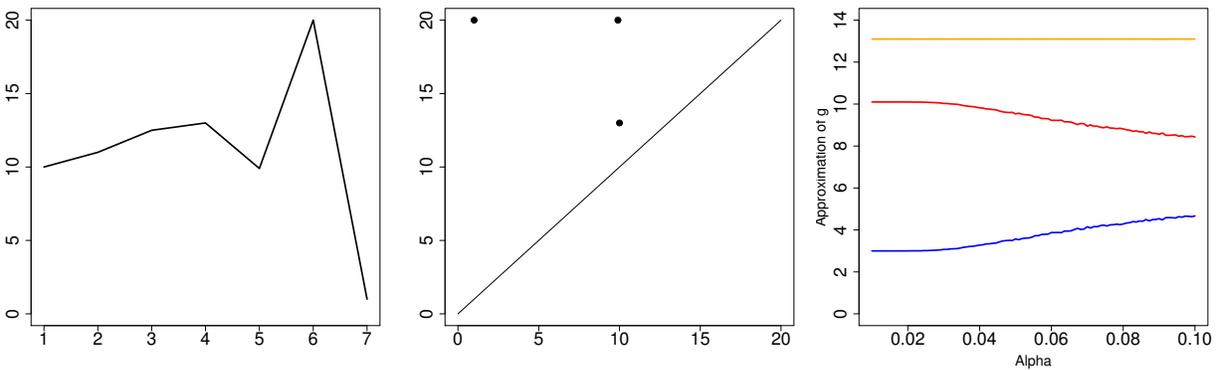


FIGURE 11. Input and results of first line graph experiment Left: a graph of the function $F_{\mathbf{a}}$ defined on a line graph with seven vertices. Middle: the persistence diagram $H(\mathbf{a}) = \text{Dgm}_0(f_{\mathbf{a}})$. We follow the extended persistence convention and pair the global min with the global max. Right: results of the experiment. Middle graph shows the value of $g_{5,\alpha}(\mathbf{a})$ versus α ; bottom graph shows $g_{1,\alpha}(\mathbf{a})$ versus α ; top graph shows their sum.

mean at the origin and bandwidth α . For each $i = 1, \dots, 7$, put $g_{i,\alpha} = h_i * K_\alpha$. The right side of Figure 11 shows graphs of the approximate values of $g_{5,\alpha}(\mathbf{a})$ and $g_{1,\alpha}(\mathbf{a})$, plotted against α , as well as a graph of their sum. There were 100 evenly spaced values of α used, ranging from $\alpha = 0.01$ to $\alpha = 0.1$. To make these graphs, we followed the

approximation procedure based on Theorem 4.8. For each fixed α , we took $N = 10000$ independent draws $\epsilon_1, \dots, \epsilon_{10000}$ from K_α , and computed

$$g_5(a) \approx \frac{1}{10000} \sum_{i=1}^{10000} h_5(a + \epsilon_i),$$

with an identical procedure for $g_1(a)$. Figure 11 also plots $g_1(a) + g_5(a)$ to demonstrate that, for the bandwidths considered, vertices v_1 and v_5 are together responsible for a consistent total amount of persistence but are competing for which generates the high persistence point.

6.2. Second line-graph experiment. Again we explore Example 5.1, this time with the input $a = (5, 1.1, 1, 1.05, 15)$ to a persistent homology computation that builds a filtration on a line graph with five vertices v_1, \dots, v_5 . The function F_a , whose graph is on the left of Figure 12, has a global min at v_3 . From the diagram in the middle, we see $h_3(a) = 15 - 1 = 14$. Note that $h_i(a) = 0$ for $i \neq 3$, since only one component is created during the entire filtration. On the right, we see convolved values of these functions for 100 evenly spaced choices of bandwidth between 0.01 and 0.5, with notation and computation procedure exactly as in 6.1 above.

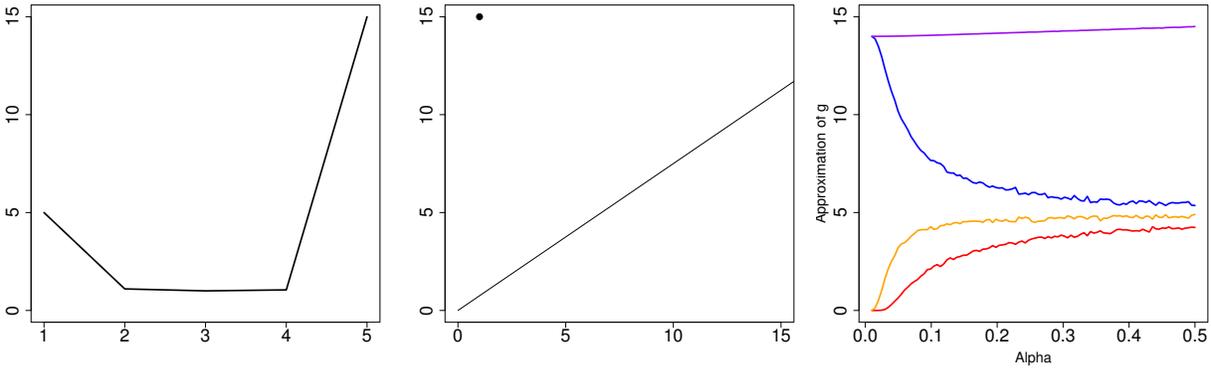


FIGURE 12. Input and results of second line graph experiment Left: a graph of the function F_a defined on a line graph with seven vertices. Middle: the persistence diagram $H(a) = \text{Dgm}_0(f_a)$. We follow the extended persistence convention and pair the global min with the global max. Right: results of experiment. Moving from bottom to top, the values of $g_{2,\alpha}(a)$, $g_{4,\alpha}(a)$, $g_{3,\alpha}(a)$, and their sum, all plotted against α .

6.3. Distance-to-a-curve experiment. Next we reconsider Example 5.2. Let C be the PL-curve with nine vertices on the left side of Figure 13. In our language, $C = C_a$, where the input vector a specifies the coordinates of the nine vertices: $v_1 = (0, 0.1)$, $v_2 = (1, 1)$, $v_3 = (2, 0.12)$, $v_4 = (7, 5)$, $v_5 = (12, 0)$, $v_6 = (7, -5)$, $v_7 = (2, -0.12)$, $v_8 = (1, -1)$, $v_9 = (0, -0.1)$. Following the vocabulary of Example 2.2, this curve placement leads to an order-preserving function f_a on the abstract full complex K on nine vertices.

Its degree-one persistence diagram $H(a) = \text{Dgm}_1(f_a)$, in the middle of the same figure, has only two off-diagonal points. The first, at $(0.2, 10)$, is created by the positive edge between v_1 and v_9 , while the second, at $(0.23, 2)$, comes from the edge between v_3 and v_7 . Thus we have $h_{1,9}(a) = 9.8$ and $h_{3,7}(a) = 1.77$. As usual, these values are highly unstable to small perturbations in the vertex positions. In very similar fashion to the

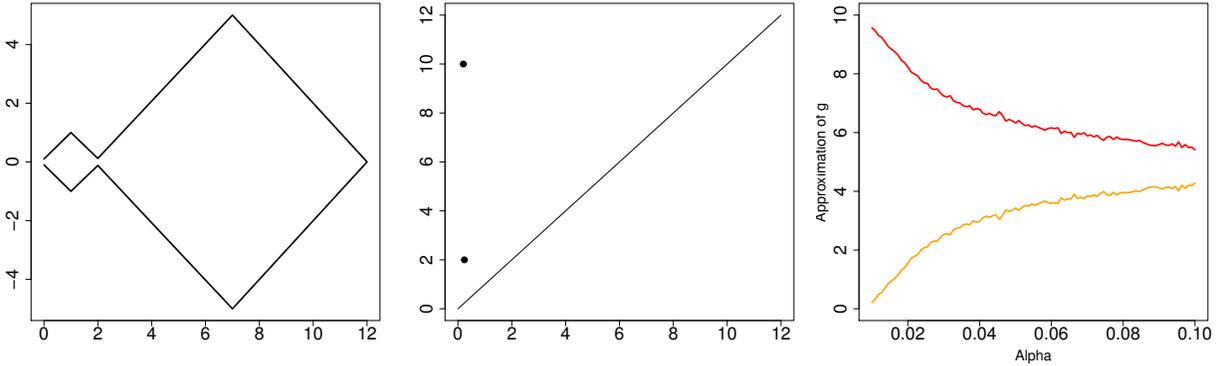


FIGURE 13. Input and results of distance-to-curve experiment. Left: the PL plane curve C_a whose nine vertices are defined in the text. Middle: the persistence diagram $H(a) = \text{Dgm}_1(f_a)$. Right: Results of experiment. Top graph shows the value of $g_{1,9,\alpha}(a)$ versus α , bottom graph shows $g_{3,7,\alpha}(a)$ versus α .

last experiment, we then computed approximate values at a for the convolved functions $g_{1,9,\alpha} = h_{1,9} * K_\alpha$ and $g_{3,7,\alpha} = h_{3,7} * K_\alpha$ where K_α was an eighteen-dimensional Gaussian kernel with bandwidth α . The results appear on the right side of Figure 13.

6.4. Locating a point in the domain. Let $u = (9.9, 20)$ be one of the high-persistence points in the diagram for our first line-graph experiment. It is accurate to say that u was created, for this specific persistent homology computation, by the addition of v_5 to the filtration. However, it is also a potentially misleading thing to say.

We propose that the difference between the persistence of u and the values of the convolutions $g_{5,\alpha}(a)$ might be seen as an indicator for how confidently one should locate u at v_5 . The graphs on the right side of Figure 11 tell us that this confidence should be low. On the other hand, the other high-persistence point $w = (1, 20)$ is created by the addition of v_7 . It turns out that $g_{7,\alpha}(a)$ remains very close to 19 for all α within a reasonable range.

6.5. Spreading out a point in the domain. Alternatively, one might choose to give u a more fuzzy location. A reasonable idea would be to spread out its location between vertices v_5 and v_1 , since v_1 is responsible for creating the same component in a nearby filtration. The graphs in Figure 11 bear this out: note that the sum of the two convolution values $g_{5,\alpha}(a) + g_{1,\alpha}(a)$ is always very close to the sum of the persistences of the components created by v_1 and v_5 . Similarly, in the second line-graph experiment, it

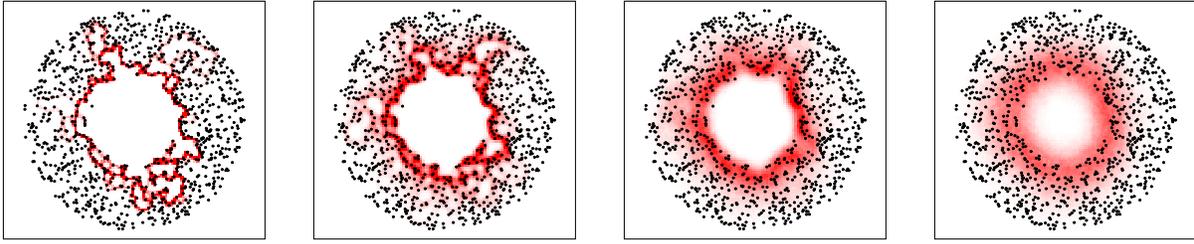


FIGURE 14. The probability of intersecting a representative cycle for each square in a grid. The bandwidths of the Gaussian kernel are 0.2, 1, 3, and 10. The color scale is given in Figure 1.

would be reasonable to smear the location of the only point throughout the immediate neighborhood of v_3 .

6.6. Convolved values as features. One could also use the values of g_i or $g_{i,j}$ as features in a machine-learning scheme. That is, the vector $(g_{1,\alpha}(a), \dots, g_{7,\alpha}(a))$ could be used as a summary feature of both the filtration created by a and the noise model K_α . The stabilities offered by Corollary 4.7 and Theorem 4.9 make this an appealing option.

6.7. Bandwidth selection – in practice. Our procedures depend on a free parameter, the bandwidth. For example, we chose a bandwidth of 3 in Example 1.3. Here we consider a slightly simpler example and consider the effect of varying the bandwidth.

We sample 1000 points uniformly from an annulus of inner and outer radius 20 and 40. Using Dionysus [40], we compute the 1-dimensional persistent homology of the alpha complex of our sample and obtain a representative cycle for the longest bar. However, the embedded location of this cycle is unstable. We would like to quantify and visualize the uncertainty of this location. To do so, we consider a grid of squares with edge-length 1. We perturb the sampled points 1000 times by adding Gaussian noise and find the proportion of trials in which the representative cycle produced by Dionysus intersects each square. By performing this procedure simultaneously for every square in the grid, we obtain Figure 14. When the standard deviation of the Gaussian noise is very small, the representative cycle barely changes between perturbations, resulting in a small number of squares with high probability of intersecting the cycle. As the standard deviation increases, the picture becomes more diffuse.

6.8. Possible choices for the location of the generator in the brain imaging data. In this section and the following section we expand on Example 1.4, which applies our method to real data.

We can obtain the estimate of $h * K$ for the observed data in Example 1.4 for a ball of any radius by considering the distance from the location of the generator of the 28th longest bar in one of the iterations of Algorithm 1 to the location of the generator of

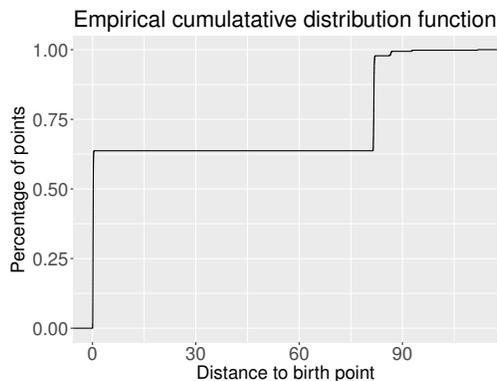


FIGURE 15. The empirical cumulative distribution function of the distance between the the location of the generator of the 28th longest bar in one of the iterations of Algorithm 1 to the location of the generator of the 28th longest bar in the observed data.

the 28th longest bar in the observed data, and computing the empirical cumulative distribution function of this function. See Figure 15. In this figure we see that there is a competitor for the location of the generator of the 28th longest bar. So in fact, the situation with this real data is quite similar to that in the elementary synthetic examples in Sections 6.1 and 6.3.

6.9. A concrete use of our brain imaging analysis for the practitioner. In Example 1.4, we show that under certain small perturbations, nearly two-thirds of the generators for the 28th longest bar in degree-zero persistent homology are born in a particular ball. Using the results of Section 6.8 we see that there is another ball containing a substantial proportion of the likely generators for this persistent homology class.

By repeating this computation for all subjects, we allow the clinician wanting to compare subjects with respect to some clinically important variable (in this case age) to focus their attention on a couple of small regions of each image. In particular, the clinician can compare the morphology of the brain arteries in these neighborhoods.

7. Future work

The work presented here opens many interesting questions and possible directions for future research.

Machine learning. Persistence diagrams have been used to produce features for machine-learning and statistical methods. This paper takes a first step towards the extraction of stable features that describe much of the other information produced during a persistent homology computation. For example, one could apply the ideas presented here to construct topological features for machine learning that are not only based on critical values but also on the locations of critical points.

Visualization. The ideas of this paper could be used to build visualization tools. For example, one might want to compute a persistence diagram, click on a point, and have the possible location candidates shown on the domain, perhaps with some sort of heat map of likelihood.

Convergence results. What is the rate of convergence of Algorithm 1 as the number of repetitions M increases?

Samples with increasing numbers of points. In all of the examples we consider, the number of sampled points is fixed. Expand the framework presented here to allow for increasing sample sizes, allowing asymptotic results to be considered.

A continuous theory. Finally, we also hope to enrich the theory whose development has started here. Work in the category of metric spaces, to define some versions of the functions h_x and g_x , and to prove Lipschitz-continuity of the latter. We believe that Example 2.6 points us in the right direction.

Features in other vector spaces. The present paper has only considered the stabilization of real-valued functions. However, one could consider functions in \mathbb{R}^n or more generally into Banach spaces.

Acknowledgements. The authors would like to thank Justin Curry, Francis Motta, Chris Tralie, and Ulrich Bauer for helpful conversations and to thank the referees for suggestions that improved the paper. The first author would like to thank the University of Florida for hosting him during the initial phase of this research. The second author would like to acknowledge that this research was partially supported by the Southeast Center for Mathematics and Biology, an NSF-Simons Research Center for Mathematics of Complex Biological Systems, under National Science Foundation Grant No. DMS-1764406 and Simons Foundation Grant No. 594594, and that this material is based upon work supported by, or in part by, the Army Research Laboratory and the Army Research Office under contract/grant number W911NF-18-1-0307.

References

- [1] Henry Adams, Atanas Atanasov, and Gunnar Carlsson. Nudged elastic band in topological data analysis. *Topol. Methods Nonlinear Anal.*, 45(1):247–272, 2015.
- [2] Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8):1–35, 2017.
- [3] Mahmuda Ahmed, Brittany Terese Fasy, and Carola Wenk. Local persistent homology based distance between maps. In *SIGSPATIAL*. ACM, Nov. 2014.
- [4] Ulrich Bauer. Ripser: a lean C++ code for the computation of vietoris–rips persistence barcodes. Software available at <https://github.com/Ripser/ripser>, 2017.
- [5] P. Bendich, S. P. Chin, J. Clark, J. Desena, J. Harer, E. Munch, A. Newman, D. Porter, D. Rouse, N. Strawn, and A. Watkins. Topological and statistical behavior classifiers for tracking applications. *IEEE Transactions on Aerospace and Electronic Systems*, 52(6):2644–2661, December 2016.
- [6] Paul Bendich, J. S. Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. *Ann. Appl. Stat.*, 10(1):198–218, 03 2016.

- [7] Andrew J. Blumberg, Itamar Gal, Michael A. Mandell, and Matthew Pancia. Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Found. Comput. Math.*, 14(4):745–789, 2014.
- [8] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research*, 16:77–102, 2015.
- [9] Peter Bubenik, Gunnar Carlsson, Peter T. Kim, and Zhi-Ming Luo. Statistical topology via Morse theory persistence and nonparametric estimation. In *Algebraic Methods in Statistics and Probability II*, volume 516 of *Contemp. Math.*, pages 75–92. Amer. Math. Soc., Providence, RI, 2010.
- [10] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, January 2009.
- [11] Gunnar Carlsson, Tigran Ishkhanov, Vin de Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *Int. J. Comput. Vision*, 76(1):1–12, 2008.
- [12] Gunnar Carlsson, Tigran Ishkhanov, Vin de Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1):1–12, 2008.
- [13] M. Carriere, S. Y. Oudot, and M. Ovsjanikov. Stable topological signatures for points on 3d shapes. In *Proc. Sympos. on Geometry Processing*, 2015.
- [14] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th annual symposium on Computational geometry, SCG '09*, pages 237–246, New York, NY, USA, 2009. ACM.
- [15] Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric Inference for Measures based on Distance Functions. *Foundations of Computational Mathematics*, 11(6):733–751, 2011. RR-6930 RR-6930.
- [16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. SpringerBriefs in Mathematics. Springer, [Cham], 2016.
- [17] Frédéric Chazal, Brittany Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Robust topological inference: distance to a measure and kernel distance. *J. Mach. Learn. Res.*, 18:Paper No. 159, 40, 2017.
- [18] Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Subsampling methods for persistent homology. In *Proceedings of the 32nd International Conference on Machine Learning, Lille, France*, volume 37. JMLR: W&CP, 2015.
- [19] Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Aarti Singh, and Larry Wasserman. On the bootstrap for persistence diagrams and landscapes. *Modeling and Analysis of Information Systems*, 20(6):96–105, 2014.
- [20] Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, and Larry Wasserman. Stochastic convergence of persistence landscapes and silhouettes. *J. Comput. Geom.*, 6(2):140–161, 2015.
- [21] Y.-C. Chen, D. Wang, A. Rinaldo, and L. Wasserman. Statistical analysis of persistence intensity functions. *ArXiv e-prints*, 2015, 1510.02502.
- [22] Yen-Chi Chen, Christopher R. Genovese, and Larry Wasserman. Statistical inference using the Morse-Smale complex. *Electron. J. Stat.*, 11(1):1390–1433, 2017.
- [23] Moo K. Chung, Peter Bubenik, and Peter T. Kim. Persistence diagrams in cortical surface data. In *Information Processing in Medical Imaging (IPMI) 2009*, volume 5636 of *Lecture Notes in Computer Science*, pages 386–397, 2009.
- [24] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37(1):103–120, January 2007.
- [25] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.*, 9(1):79–103, 2009.
- [26] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have l_p -stable persistence. *Found. Comput. Math.*, 10(2):127–139, February 2010.
- [27] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Computational geometry (SCG'06)*, pages 119–126. ACM, New York, 2006.
- [28] Vin de Silva and Gunnar Carlsson. Topological estimation using witness complexes. *Eurographics Symposium on Point-Based Graphics*, 2004.
- [29] Tamal K. Dey and Rephael Wenger. Stability of critical points with interval persistence. *Discrete Comput. Geom.*, 38(3):479–512, 2007.

- [30] Tamal Krishna Dey, Fengtao Fan, and Yusu Wang. Graph induced complex on point data. In *Proceedings of the Twenty-ninth Annual Symposium on Computational Geometry, SoCG '13*, pages 107–116, New York, NY, USA, 2013. ACM.
- [31] Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.
- [32] Herbert Edelsbrunner and John L. Harer. *Computational topology*. American Mathematical Society, Providence, RI, 2010.
- [33] Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Larry Wasserman, Sivaraman Balakrishnan, and Aarti Singh. Confidence sets for persistence diagrams. *Ann. Statist.*, 42(6):2301–2339, 2014.
- [34] D.H. Fremlin. *Measure Theory, Volume 4*. Torres Fremlin, 2000.
- [35] P. Frozini and B. Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9:596–603, 1999.
- [36] Marcio Gameiro, Yasuaki Hiraoka, Shunsuke Izumi, Miroslav Kramar, Konstantin Mischaikow, and Vidit Nanda. A topological measurement of protein compressibility. *Jpn. J. Ind. Appl. Math.*, 32(1):1–17, 2015.
- [37] Robert Ghrist. Barcodes: the persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)*, 45(1):61–75, 2008.
- [38] Violeta Kovacev-Nikolic, Peter Bubenik, Dragan Nikolić, and Giseon Heo. Using persistent homology and dynamical distances to analyze protein binding. *Stat. Appl. Genet. Mol. Biol.*, 15(1):19–38, 2016.
- [39] Chunyuan Li, M. Ovsjanikov, and F. Chazal. Persistence-based structural recognition. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 2003–2010, June 2014.
- [40] Dimitriy Morozov. Dionysus: a C++ library with various algorithms for computing persistent homology. Software available at <http://www.mrzv.org/software/dionysus/>, 2012.
- [41] Elizabeth Munch, Katharine Turner, Paul Bendich, Sayan Mukherjee, Jonathan Mattingly, John Harer, et al. Probabilistic fréchet means for time varying persistence diagrams. *Electronic Journal of Statistics*, 9:1173–1204, 2015.
- [42] James R. Munkres. *Elements of Algebraic Topology*. Addison Wesley, 1993.
- [43] Vidit Nanda. Perseus: the persistent homology software. Software available at <http://www.math.rutgers.edu/~vidit/perseus/index.html>, 2013.
- [44] P. Niyogi, S. Smale, and S. Weinberger. A topological view of unsupervised learning from noisy data. *SIAM J. Comput.*, 40(3):646–663, 2011.
- [45] Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.
- [46] Jan Reininghaus, Stefan Huber, Ulrich Bauer, and Roland Kwitt. A stable multi-scale kernel for topological machine learning. *Proc. CVPR*, pages 4741–4748, 2015.
- [47] B. W. Silverman. *Density estimation for statistics and data analysis*. Monographs on Statistics and Applied Probability. Chapman & Hall, London, 1986.
- [48] Abraham Smith, Paul Bendich, John Harer, and Jay Hineman. Supervised learning of labeled point-cloud differences via cover-tree entropy reduction. 02 2017.
- [49] M. P. Wand and M. C. Jones. *Kernel smoothing*, volume 60 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, Ltd., London, 1995.
- [50] Shmuel Weinberger. The complexity of some topological inference problems. *Found. Comput. Math.*, 14(6):1277–1285, 2014.
- [51] Afra Zomorodian and Gunnar Carlsson. Localized homology. *Comput. Geom.*, 41(3):126–148, 2008.

DEPARTMENT OF MATHEMATICS, DUKE UNIVERSITY, AND GEOMETRIC DATA ANALYTICS, INC.

Email address: bendich@math.duke.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF FLORIDA

Email address: peter.bubenik@ufl.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF FLORIDA

Email address: wagnera@ufl.edu