Stability of Critical Points with Interval Persistence *

Tamal K. Dey Rephael Wenger †

Abstract

Scalar functions defined on a topological space Ω are at the core of many applications such as shape matching, visualization and physical simulations. Topological persistence is an approach to characterizing these functions. It measures how long topological structures in the sub-level sets $\{x \in \Omega : f(x) \leq c\}$ persist as c changes. Recently it was shown that the *critical values* defining a topological structure with relatively large persistence remain almost unaffected by small perturbations. This result suggests that topological persistence is a good measure for matching and comparing scalar functions. We extend these results to *critical points* in the domain by redefining persistence and critical points and replacing *sub-level sets* $\{x \in \Omega : f(x) \leq c\}$ with *interval sets* $\{x \in \Omega : a \leq f(x) < b\}$. With these modifications we establish a stability result for critical points. This result is strengthened for maxima that can be used for matching two scalar functions.

1 Introduction

A *scalar field* is a scalar function $f : \Omega \to \mathbb{R}$ defined on some topological space Ω. Examples of scalar fields are fluid pressure in computational fluid dynamics simulations, temperature in oceanographic or atmospheric studies, and density in medical CT or MRI scans. A *level set* of a scalar field is a set of points with the same scalar value, i.e., $\{x \in \Omega : f(x) = c\}$. One way of deriving quantitative information about scalar fields is by studying the topological structures of its level sets or the sub-level sets, such as $\{x \in \Omega : f(x) \leq c\}$. The mathematical field of Morse Theory is the study of these topological structures.

Among the most basic problems on scalar fields is simplifying a scalar field for compact representation, identifying important features in a scalar field, and characterizing the essential structure of a scalar field. Extracting and representing the topological structure of the level sets is one way of approaching all these problems. However, this topological structure may contain "small" topological features which are insignificant or caused by noise. Small topological features should be removed in simplification and ignored in characterizing essential structure or identifying important features. How does one determine which topological features are small?

*Research partly supported by NSF CARGO grant DMS-0310642

†Department of Computer Science and Engineering, The Ohio State U., Columbus, OH 43210, USA.

Edelsbrunner, Letscher, and Zomorodian [8] introduced the notion of *topological persistence*. As $c \in \mathbb{R}$ increases, topological features appear and disappear in the sub-level set $\{x \in \Omega : f(x) \leq c\}$. If a topological feature appears at "time" a and disappears at "time" b, then its persistence is the difference, $b - a$, between these two times. Edelsbrunner et al. [8] use homology groups over $\mathbb{Z}/2\mathbb{Z}$ to define topological features. Zomorodian and Carlsson [13] generalized the theory of topological persistence to homology groups over any fields.

Topological persistence gives an approach to comparing scalar fields. Two fields are similar if they have matching topological features with approximately the same persistence. This approach to comparing fields makes sense only if persistence remains stable under relatively small perturbations of the scalar fields. Cohen-Steiner, Edelsbrunner, and Harer [4] proved that "large" persistence values remain almost unaffected. More precisely, let scalar field $\hat{f} : \Omega \to \mathbb{R}$ be a small perturbation of field $f : \Omega \to \mathbb{R}$, (i.e., $|\hat{f}(x) - f(x)| \leq \delta$ for all $x \in \Omega$.) If f has a topological structure with relatively large persistence which appears at a and disappears at b, then $\hat{f}$ has a corresponding topological structure which appears around a and disappears around b.

Critical point stability. The result of Cohen-Steiner et al. [4] showed that the critical values for structures with large persistence remain stable under small perturbations of the scalar field. Scalar fields also have *critical points*, points in the domain which change the topological structure of the level sets. It is natural to ask if critical points for structures with large persistence remain stable under perturbations of the field. If two scalar fields are close, then are their significant critical points "close"?

At the very onset, the problem of stability for critical points appears almost hopeless. Consider functions f and $\hat{f}$ in Figure 1 where $|f(x) - \hat{f}(x)| < \delta$ for all $x \in \Omega$. The maxima, p and p', of f and $\hat{f}$, respectively, can be made arbitrarily far apart even as δ is made arbitrarily small. So, there is no result for stability of critical points in terms of Hausdorff distances as considered by Cohen Steiner et al. [4] for critical values.

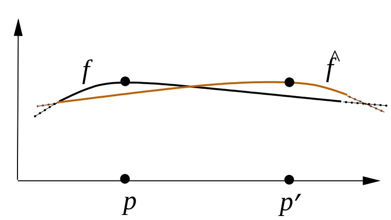

Figure 1: The maximum p for a real valued function f has moved by large distance even for an arbitrarily close approximant $\hat{f}$.

One main contribution of this paper is to overcome this difficulty by defining a new notion of stability for the critical points. Instead of using a metric in the domain, we use the range to determine neighborhoods of points. These neighborhoods help defining the stability. A (η_1, η_2)-*neighborhood* of a point p is the *pathwise connected component* of $\{x \in \Omega : f(p) - \eta_1 \leq f(x) \leq f(p) + \eta_2\}$ containing p. A point which is in this neighborhood for small values of η_1 and η_2 is "close" to p. Note that points p and p' in Figure 1 are close in this sense. We show that if p destroys a 'persistent homology element' in f, then a (η_1, η_2)-neighborhood of p contains a point p' which destroys a persistent homology element in $\hat{f}$. The values of η_1 and η_2 depend upon the persistence of the homology element and the difference δ between f and $\hat{f}$. Theorem 3 states this result formally. Although this result relates critical points for functions under perturbations, it does not match them. To construct a matching of critical points, we need each destroying critical point of $\hat{f}$ to be in the neighborhood of only one critical point of f. We establish this stronger result for local *maxima* of functions on manifolds.

Interval persistence. As with the critical values, not all critical points can be stable under perturbations. One may expect that only those critical points responsible for persistent homology structures remain stable. Therefore, we need a notion of persistence for critical points. Instead of using topological persistence as in Cohen Steiner et al. [4], we use *interval persistence* for this purpose. There are two reasons to introduce this new notion of persistence. First, it can certify more critical points as persistent than the original topological persistence does. Second, interval persistence gives a natural setting for our critical point stability result in Theorem 3. We illustrate the first point with an example below. The second point is discussed in Section 4 after Theorem 3.

Let $f(x)$ be the z-coordinate of any point x on the surface in $\mathbb{R}^3$ shown in Figure 2. There are eight critical points $\{p_i, i = 0, .., 7\}$. For sub-level sets, point p_1 creates the homology element $[c_2]$ generated by the cycle c_2. Homology element $[c_2]$ is destroyed by p_2. Thus, p_1 is paired with p_2. Similarly, point p_3 creates the homology element $[c_1]$ which is destroyed by p_4, so p_3 is paired with p_4. The critical points p_0, p_5, p_6, and p_7 create homology elements that are never destroyed and so these critical points remain unpaired by topological persistence.

On the other hand, if we consider the interval set $\{x \in \Omega : f(p_0) < f(x) < f(p_7)\}$, the homology element $[c_0]$ starts just after p_0 and ends at p_7 (note that the interval set does not include p_0 or p_7). Interval persistence pairs p_0 with p_7. The interval set $\{x \in \Omega : f(p_5) < f(x) < f(p_6)\}$ has two connected components which get joined by p_5 and p_6. Again, interval persistence pairs p_5 with p_6. We give formal definitions and explanations of interval persistence and this pairing in Section 2.6.

Agarwal et al. in [1] proposed a different extension of the topological persistence pairings using Reeb graphs. This pairing is deduced only for 2-manifolds as opposed to the broader class of topological spaces covered by interval persistence.

2 Definitions and assumptions

2.1 Homology groups

For a topological space X, the kth homology group $H_k(X)$ is an algebraic encoding of the connectivity of X in the kth dimension. For a good exposition on homology groups we refer to Hatcher [9]. We will use *singular homology* which is more general than simplicial or cellular homology. We will also use *reduced homology groups*, usually represented as $\tilde{H}_k(X)$. Groups $H_k(X)$ and $\tilde{H}_k(X)$ are exactly the same for all $k > 0$,

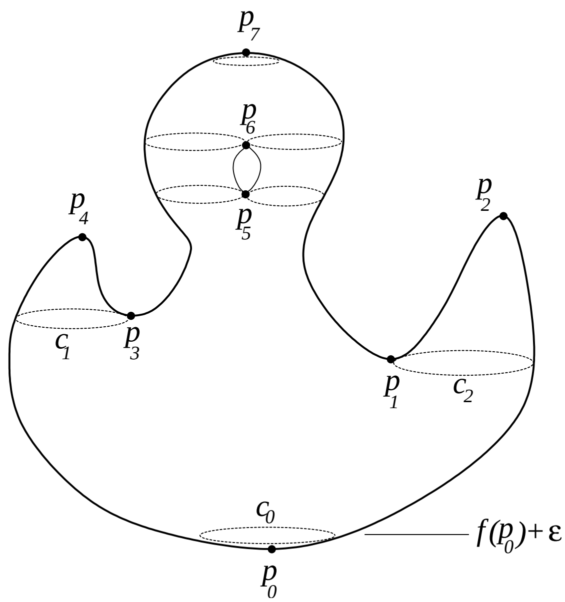

Figure 2: Critical points $p_0, p_1, \ldots, p_7$. Pairs under topological persistence are (p_1, p_2) and (p_3, p_4). Pairs under interval persistence are (p_1, p_2), (p_3, p_4), (p_0, p_7) and (p_5, p_6).

and their rank differs by one when k equals zero. Because we use reduced homology almost everywhere in this paper, we will drop the '~' mark and represent $\tilde{H}_k(X)$ as $H_k(X)$, except where otherwise noted.

Although homology groups are defined for coefficients drawn from any ring, we will consider only fields such as $\mathbb{R}, \mathbb{Q}, \mathbb{Z}_p = \mathbb{Z}/p\mathbb{Z}$ for a prime p as in the previous works [4, 13]. Over fields the homology groups are vector spaces and the rank of $H_k(X)$, denoted $\beta_k(X)$, is called the kth Betti number of X.

A continuous map $f\colon X \to Y$ between two topological spaces X and Y induces a homomorphism, say f_*, between their homology groups, $H_k(X) \xrightarrow{f_*} H_k(Y)$. This property is carried over the composition of maps, that is, $(f \circ g)_* = f_* \circ g_*$. In our case, the maps between spaces will be *inclusion* maps. This means, if $X \subseteq Y$, we will consider the map $H_k(X) \xrightarrow{i_*} H_k(Y)$ where i_* is induced by the inclusion map $i\colon X \to Y$.

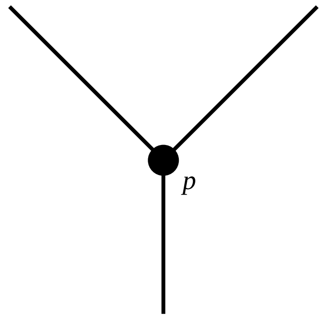

Figure 3: Function f is the height function in the y-direction. The mapping $H_k(F_{-\infty}^{f(p)}) \to H_k(F_{-\infty}^{f(p)} \cup \{p\})$ is an isomorphism for any $k \geq 0$ but $H_0(F_{f(p)}^{b}) \to H_0(F_{f(p)}^{b} \cup \{p\})$ is not for any $b > f(p)$.

2.2 Interval sets

We use the following notation to define the interval sets bounded by the level sets of a function. For $a, b \in \mathbb{R}$ and functions f and g, let

$$\begin{aligned} F_a^b &= \{x \in \Omega : a < f(x) < b\} \text{ and} \\ G_a^b &= \{x \in \Omega : a < g(x) < b\}. \end{aligned}$$

In our results and proofs we need the space F_a^b and G_a^b closed at the bottom. So, we define

$$\begin{aligned} \underline{F}_a^b &= \{x \in \Omega : a \leq f(x) < b\} \text{ and} \\ \underline{G}_a^b &= \{x \in \Omega : a \leq g(x) < b\}. \end{aligned}$$

Notice that a could be $-\infty$ and b could be ∞. With these notations, the (η_1, η_2)-neighborhood of a point p is the connected component of $\mathrm{cl}(F_{f(p)-\eta_1}^{f(p)+\eta_2})$ containing p. Here cl denotes the closure.

2.3 Critical values and points

Intuitively, a critical value is a value at which the homology of an interval set changes. A critical point is a point on the *boundary* of an interval set whose addition to that interval set changes its homology. We give the following formal definitions:

Definition 1. *Value $b \in \mathbb{R}$ is H_k-critical for the interval sets of $f\colon \Omega \to \mathbb{R}$ if $H_k(F_a^b) \to H_k(F_a^b \cup f^{-1}(b))$ is not an isomorphism for some $a < b$ or if $H_k(F_b^c) \to H_k(F_b^c \cup f^{-1}(b))$ is not an isomorphism for some $c > b$.*

Definition 2. *Point $p \in \Omega$ is H_k-critical for the interval sets of $f\colon \Omega \to \mathbb{R}$ if $H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ is not an isomorphism for some $a < f(p)$ or if $H_k(F_{f(p)}^{c}) \to H_k(F_{f(p)}^{c} \cup \{p\})$ is not an isomorphism for some $c > f(p)$.*

Note that in both definitions, a could be $-\infty$ or c could be ∞. A value or point can be H_k-critical for different values of k.

Cohen-Steiner et al. [4] give a similar definition for critical values. However, the definition of Cohen-Steiner et al. uses only sub-level sets instead of interval sets. Point $p \in \Omega$ is a H_k-critical point for the <u>sub-level</u> sets of f if $H_k(F_{-\infty}^{f(p)}) \to H_k(F_{-\infty}^{f(p)} \cup \{p\})$ is not an isomorphism for some integer $k \geq 0$.

As shown in Figure 3, H_k-critical points for interval sets are not quite equivalent to H_k-critical points for sub-level sets. In Figure 3, the mapping $H_k(F_{-\infty}^{f(p)}) \to H_k(F_{-\infty}^{f(p)} \cup \{p\})$ is an isomorphism for any integer k. On the other hand, $H_0(F_{f(p)}^{b}) \to H_0(F_{f(p)}^{b} \cup \{p\})$ is not an isomorphism and so p is a critical point for interval sets.

Of course, if we replace function f in Figure 3 by function $-f$ then the mapping $H_0(F_{-\infty}^{-f(p)}) \to H_0(F_{-\infty}^{-f(p)} \cup \{p\})$ is no longer an isomorphism. Thus, p is an H_0-critical point for the sub-level sets of the function $-f$.

Figure 3 illustrates the following relationship between critical points for interval sets and critical points for sub-level sets. A point which is critical for the interval sets of f is either critical for the sub-level sets of f or critical for the sub-level sets of $-f$ or both (Theorem 1.)

To prove Theorem 1, we first give the following relationship between $H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ and $H_{-\infty}(F_a^{f(p)}) \to H_k(F_{-\infty}^{f(p)} \cup \{p\})$. The proof is based on homological algebra and is left to the appendix.

Lemma 1. *Let a be some value less than $f(p)$. Mapping $H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ is an isomorphism for all integers $k \geq 0$ if and only if $H_k(F_{-\infty}^{f(p)}) \to H_k(F_{-\infty}^{f(p)} \cup \{p\})$ is an isomorphism for all integers $k \geq 0$.*

Note that Lemma 1 does not hold if k is fixed. For instance, in Figure 2 the mapping $H_1(F_{-\infty}^{f(p_7)}) \to H_1(F_{-\infty}^{f(p_7)} \cup \{p_7\})$ is an isomorphism while the mapping $H_1(F_{f(p_6)}^{f(p_7)}) \to H_1(F_{f(p_6)}^{f(p_7)} \cup \{p_7\})$ is not. There is no contradiction to Lemma 1, since $H_2(F_{-\infty}^{f(p_7)}) \to H_2(F_{-\infty}^{f(p_7)} \cup \{p_7\})$ is not an isomorphism.

Theorem 1 follows directly from Lemma 1.

Theorem 1. *For some $k \geq 0$, a point p is H_k-critical for the interval sets of $f : \Omega \to \mathbb{R}$ if and only if, for some $j \geq 0$, p is H_j-critical for the sub-level sets of f or $-f$ (or both.)*

Proof. A point p is an H_k-critical point for the interval sets of f if $H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ is not an isomorphism for some $a < f(p)$ or if $H_k(F_{f(p)}^{c}) \to H_k(F_{f(p)}^{c} \cup \{p\})$ is not an isomorphism for some $c > f(p)$. Since a can be $-\infty$, a point which is H_k-critical for the sub-level sets of f is H_k-critical for the interval sets of f. Since c can be ∞, a point which is H_k-critical for the sub-level sets of $-f$ is also H_k-critical for the interval sets of f. Thus, if p is H_k-critical for the sub-level sets of f or of $-f$, then p is H_k-critical for the interval sets of f.

Assume p is H_k-critical for the interval sets of f. By definition, either mapping $H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ is not an isomorphism, $a < f(p)$, or $H_k(F_{f(p)}^{c}) \to H_k(F_{f(p)}^{c}) \cup \{p\})$ is not an isomorphism, $c > f(p)$, (or both.) By Lemma 1, if

$H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ is not an isomorphism, then $H_j(F_{-\infty}^{f(p)}) \to H_j(F_{-\infty}^{f(p)} \cup \{p\})$ is not an isomorphism for some integer j. Similarly, if $H_k(F_{f(p)}^c) \to H_k(F_{f(p)}^c) \cup \{p\})$ is not an isomorphism, then $H_j(F_{f(p)}^{\infty}) \to H_j(F_{f(p)}^{\infty} \cup \{p\})$ is not an isomorphism for some integer j. If $H_j(F_{-\infty}^{f(p)}) \to H_j(F_{-\infty}^{f(p)} \cup \{p\})$ is not an isomorphism, then p is H_j-critical for the sub-level sets of f. If $H_j(F_{f(p)}^{\infty}) \to H_k(F_{f(p)}^{\infty} \cup \{p\})$ is not an isomorphism, then p is H_j-critical for the sub-level sets of $-f$. □

For the rest of this paper, all critical points are H_k-critical for interval sets, unless otherwise noted.

2.4 Destruction

Let Ω be a topological space. We define destruction of homology elements in sets by critical points and other sets in general.

Definition 3. *For $X \subseteq \Omega$ and $Y \subseteq \Omega$, set Y* destroys *non-zero $h \in H_k(X)$ if the image of h under the mapping $H_k(X) \to H_k(X \cup Y)$ is zero. In particular, if q is a point in Ω, point q* destroys *non-zero $h \in H_k(X)$ if the image of h under the mapping $H_k(X) \to H_k(X \cup \{q\})$ is zero.*

The above definition does not apply to points that are not in the closure of X though we encounter this situation repeatedly. We would like to say that a cycle generated in a level set is destroyed by a point disjoint from the level set. So, we extend the definition of destruction slightly.

Definition 4. *If $X \subseteq Z \subseteq \Omega$ and $Y \subseteq \Omega$, then we say that Y* destroys *the image of $h_x \in H_k(X)$ in $H_k(Z)$, if $h_z \in H_k(Z)$ is the image of h_x under the mapping $H_k(X) \to H_k(Z)$ and h_z is non-zero and Y destroys h_z.*

We apply this definition repeatedly where X is some level set $f^{-1}(a)$ and Y is a point. For brevity, we say that point q destroys $h \in H_k(f^{-1}(a))$ if a is less than $f(q)$ and point q destroys the image of h in $H_k(\underline{F}_a^{f(q)})$. In Figure 2, $[c_0]$ is a non-zero element in $H_1(f^{-1}(f(p_0) + \varepsilon))$. It is also a non-zero element in $H_1(\underline{F}_{f(p_0)+\varepsilon}^{f(p_7)})$ which is destroyed by p_7.

If point q destroys $h \in H_k(f^{-1}(a))$, then $H_k(F_{a'}^{f(q)}) \to H_k(F_{a'}^{f(q)} \cup \{q\})$ is not an isomorphism for any a' where $a < a' < f(q)$. (See Appendix, Lemma 14.) Thus, if q destroys $h \in H_k(f^{-1}(a))$, then q is an H_k-critical point for the interval sets of f.

A function $f : \Omega \to \mathbb{R}$ is *point destructible* if whenever $h \in H_k(\underline{F}_a^b)$ is destroyed by $f^{-1}(b)$, then h is destroyed by some point $q \in f^{-1}(b)$. Morse functions on smooth manifolds are point destructible. (See the next section for the definition of Morse functions.) Piecewise linear functions which have a different scalar value at each vertex of the underlying simplicial complex are also point destructible.

2.5 Morse functions

Let Ω be a smooth, compact d-manifold and let $f : \Omega \to \mathbb{R}$ be a smooth map on Ω. The *critical points* in Morse theory are the points p such that the gradient of f at p is the zero vector. A critical point is *non-degenerate* if its Hessian has full rank. The *index* of a non-degenerate critical point is the number of negative eigenvalues of the Hessian. (See [11] or [10] for definition of the Hessian and further explanation of Morse theory.) Function f is Morse if all its critical points are non-degenerate. If function f is Morse, then its index k critical points in Morse theory correspond to its H_k-critical points as defined in Section 2.3.

2.6 Interval persistence

Similar to the topological persistence, the interval persistence of a point $p \in \Omega$ measures the "age" of the "oldest" homology element destroyed by p. However, the homology elements are considered over interval sets as opposed to the sub-level sets. Formally, for each $k \geq 0$, the *interval persistence* of point $p \in \Omega$ is

$$\begin{aligned} \Pi_k^f(p) &= \sup\{f(p) - a : p \text{ destroys some} \\ &\quad \text{non-zero } h \in H_k(f^{-1}(a))\}. \end{aligned}$$

We use $\sup$ in place of $\max$ because it is possible that p destroys non-zero elements of $H_k(f^{-1}(a+\epsilon))$ for any $\epsilon > 0$ but not elements of $H_k(f^{-1}(a))$. For example, in Figure 2, p_7 destroys $[c_0]$ in $H_1(f^{-1}(f(p_0)+\epsilon))$ but no element of $H_1(f^{-1}(f(p_0)))$.

The persistent Betti numbers relate the homology groups of one space into the other. For $X \subseteq Y$ and $k \geq 0$, let $H_k^{X,Y}$ be the image of the map $H_k(X) \to H_k(Y)$ induced by inclusion $X \to Y$. Define

$$\xi_k(X,Y) = \dim H_k^{X,Y}.$$

In words, $\xi_k(X,Y)$ counts the number of non-zero generators of $H_k(X)$ that remain so in the larger space Y. The persistent Betti numbers are defined as $\beta_a^b = \xi_k(F_{-\infty}^a, F_{-\infty}^b)$ for $a \leq b$ [4, 13].

Assume that f has only a finite number of critical values for interval sets. Let δ be the minimum difference between any two critical values of f and let ϵ equal $\delta/2$. Define

$$\mu_a^b = (\beta_{a+\epsilon}^{b-\epsilon} - \beta_{a+\epsilon}^{b+\epsilon}) - (\beta_{a-\epsilon}^{b-\epsilon} - \beta_{a-\epsilon}^{b+\epsilon}).$$

As noted in Cohen-Steiner et al. [4], μ_a^b counts the number of homology generators born at a which die at b.

Edelsbrunner et. al [8] introduced a critical point pairing algorithm based on when a homology generator is born and killed. This algorithm is used in the context of Morse functions in [7]. It can be shown that a pair (p,q) can be characterized by $\mu_{f(p)}^{f(q)}$. Specifically, two critical points p and q with $f(p) < f(q)$ are paired if and only if $\mu_{f(p)}^{f(q)}$ is positive. Edelsbrunner et al. [7, 8] define the *topological persistence* of p and q as $f(q) - f(p)$.

Interval persistence generalizes the notion of topological persistence in the following sense. One can show that a pair (p,q) has a topological persistence $f(q) - f(p)$ only if it has an interval persistence $f(q) - f(p)$ though the converse is not necessarily true.

Theorem 2. *Let f be a Morse function where no two critical points share the same critical value. If a pair of critical points (p, q) with $f(p) < f(q)$ has topological persistence $f(q) - f(p)$, then point q has interval persistence $f(q) - f(p)$.*

Proof of Theorem 2 is left to the appendix. The converse of Theorem 2 is not true. Critical point p_7 in Figure 2 destroys $[c_0]$ in $H_1(F_{f(p_0)}^{f(p_7)})$. Thus p_7 has interval persistence $\Pi_1^f(p_7) = f(p_7) - f(p_0)$. Critical point p_6 in Figure 2 joins the two connected components in $F_{f(p_5)}^{f(p_6)}$, destroying a homology element in $H_0(F_{f(p_5)}^{f(p_6)})$. Thus p_6 has interval persistence $\Pi_0^f(p_6) = f(p_6) - f(p_5)$. On the other hand, for sub-level sets, points p_6 and p_7 create homology elements in $H_1(F_{-\infty}^{f(p_6)} \cup \{p_6\})$ and $H_2(F_{-\infty}^{f(p_7)} \cup \{p_7\})$, respectively, and these homology elements are never destroyed. Thus, points p_6 and p_7 have infinite (or undefined) topological persistence.

3 Maps and spaces

We will be dealing with continuous functions on a compact, connected topological space, Ω. We need some conditions that these functions will be well-behaved, i.e. have properties similar to Morse functions. However, we do not want to restrict ourselves to differentiable functions or to Morse functions.

For a function $f : \Omega \to \mathbb{R}$ and $a \in \mathbb{R}$, define the *open ϵ-neighborhood* of $f^{-1}(a)$ as:

$$N_\epsilon(f^{-1}(a)) = \{x \in \Omega : a - \epsilon < f(x) < a + \epsilon\}.$$

The first property we require is that the topology of $f^{-1}(a)$ is similar to the topology of a ϵ-neighborhood of $f^{-1}(a)$ for suitably small ϵ. The second property is that f is point destructible. These properties are similar to the Morse condition that critical points are isolated. We define the first property more formally below.

Represent the unit interval $[0, 1]$ by I. Subspace $X \subseteq Y$ is a *strong deformation retract* of Y if there is a continuous $\phi : Y \times I \to Y$ such that $\phi(y, 0) = y$ and $\phi(y, 1) \in X$ for all $y \in Y$ and $\phi(x, t) = x$ for all $x \in X$ and $t \in [0, 1]$. In other words, ϕ continuously deforms Y into X without moving any points in X. If X is a strong deformation retract of Y, then $H_k(X)$ is isomorphic to $H_k(Y)$.

Definition 5. *The continuous function $f : \Omega \to \mathbb{R}$ is* LR *(locally retractible) if for all $a \in \mathbb{R}$, there exists some $\epsilon_0 > 0$ such that for all $\epsilon \leq \epsilon_0$, the set $f^{-1}(a)$ is a strong deformation retract of $N_\epsilon(f^{-1}(a))$.*

Piecewise linear functions on compact spaces are LR. (See Appendix, Lemma 18.) If a continuous function is Morse, then the function is LR. (See Milnor [11, pp. 12–20] for a proof. Milnor actually proves that $\{x : f(x) \leq a\}$ is a deformation retract of $\{x : f(x) \leq a + \epsilon\}$ but his proof also shows that $f^{-1}(a)$ is a strong deformation retract of $\{x : a - \epsilon < f(x) < a + \epsilon\}$.) We use the following property of LR functions. (See appendix for proof.)

Lemma 2. *Let $f : \Omega \to \mathbb{R}$ be a continuous, point destructible, LR function. For any non-zero $h \in H_k(f^{-1}(a))$, if $\underline{F}_a^\infty$ destroys h, then some point $q \in \Omega$ destroys h.*

If point q destroys some element of $H_k(f^{-1}(a_1))$ and $a_1 < a_2 < f(q)$, does q destroy some element of $H_k(f^{-1}(a_2))$? The answer is yes.

Lemma 3. *Let $f : \Omega \to \mathbb{R}$ be a continuous, LR function. If $q \in \Omega$ destroys some non-zero $h_1 \in H_k(f^{-1}(a_1))$, then, for every a_2 where $a_1 < a_2 < f(q)$, there exists some non-zero $h_2 \in H_k(f^{-1}(a_2))$ such that point q destroys h_2 and elements h_1 and h_2 have the same image h in $H_k(\underline{F}_{a_1}^{f(q)})$.*

The essence of the lemma is that if h is the image of $h_1 \in H_k(f^{-1}(a_1))$ in $H_k(\underline{F}_{a_1}^{f(q)})$ and q destroys h, then some cycle in $f^{-1}(a_2)$ also generates h in $H_k(\underline{F}_{a_1}^{f(q)})$. The homology element in $H_k(f^{-1}(a_2))$ generated by this cycle is also destroyed by q.

We also need a version of Lemma 3 which combines two functions (See Figure 4).

Lemma 4. *Let $f : \Omega \to \mathbb{R}$ and $g : \Omega \to \mathbb{R}$ be continuous, LR functions and let $X = \{x \in \Omega : a_2 \leq g(x)$ and $f(x) < f(q)\}$ for some $a_2 \in \mathbb{R}$ and point $q \in \Omega$. If $q \in \Omega$ destroys some non-zero $h_1 \in H_k(f^{-1}(a_1))$ and $g^{-1}(a_2) \subseteq F_{a_1}^{f(q)}$, then there exists some non-zero $h_2 \in H_k(g^{-1}(a_2))$ such that point q destroys the image of h_2 in $H_k(X)$ and elements h_1 and h_2 have the same image h in $H_k(\underline{F}_{a_1}^{f(q)})$.*

Proofs of Lemmas 3 and 4 are left to the appendix.

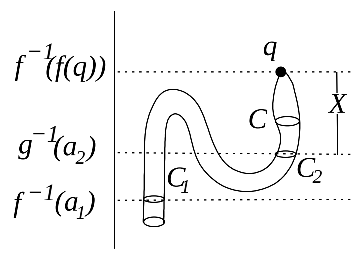

Figure 4: Cycle C generates $h \in H_1(\underline{F}_{a_1}^{f(q)})$. Cycle C_1 and C_2 generate $h_1 \in H_1(f^{-1}(a_1))$ and $h_2 \in H_1(g^{-1}(a_2))$ respectively which are destroyed by q. Element h is the image of both h_1 and h_2 in $H_1(F_{a_1}^{f(q)})$.

4 Stability

In this section we prove one of our main results, Theorem 3. Let f and g be two functions defined on Ω. We say $|f - g| < \delta$ if $|f(x) - g(x)| < \delta$ for all $x \in \Omega$. Let γ be some value greater than 2δ. We show that if q is H_k-critical for f with interval persistence greater than or equal to γ, then there is a H_k-critical point q' for g in the $(\gamma, 2\delta)$-neighborhood of q with interval persistence greater than $\gamma - 2\delta$. Moreover, the values $f(q)$ and $g(q')$ are close. This theorem not only relates q and q' in the range as in Cohen-Steiner et al. [4] but also in the domain.

Consider Figure 5 where f is a function defined on the surface in $\mathbb{R}^3$ shown in the figure and $f(x)$ is the z-coordinate of the point x. The set $\{x \in \Omega : f(x) \leq f(r)\}$ is homeomorphic to a pinched cylinder with circles c_1 and c_2 bounding each end. The two maxima p and q destroy the elements $[c_1]$ and $[c_2]$ respectively in $H_1(f^{-1}(f(r)))$. We have $\Pi_1^f(p) = f(p) - f(r)$ and $\Pi_1^f(q) = f(q) - f(r)$. Now consider a slightly perturbed f denoted as g. Set g equal to f everywhere except in the vicinity of p and q so that $|f - g|$ is much smaller than $\Pi_1^f(p)$ and $\Pi_1^f(q)$. Let a be any value in between $f(p)$ and $f(r)$ where $f(p) - a > 2|f - g|$. Our stability result asserts that there will be a maximum p' for g where $p' \in \underline{F}_a^{f(p)}$. Similar result holds for q as well.

Theorem 3. *Let $f, g : \Omega \to \mathbb{R}$ be continuous, point destructible, LR functions on Ω where $|f - g| < \delta$. Further, let q be a H_k-critical point that destroys a non-zero $h_f \in H_k(f^{-1}(a))$. If $\gamma = f(q) - a > 2\delta$, then $\Pi_k^f(q) > \gamma$ and there is a H_k-critical point q' for g in the $(\gamma, 2\delta)$-neighborhood of q such that*

i. $f(q) - \delta \le g(q') \le f(q) + \delta$, *and*

ii. $\Pi_k^g(q') > \gamma - 2\delta$.

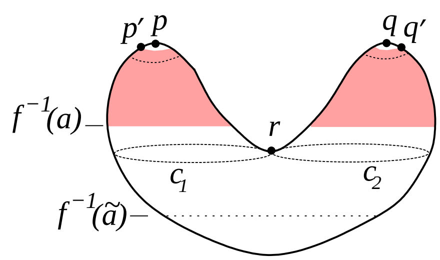

Figure 5: Illustration that interval persistence cannot be replaced by topological persistence in Theorem 3.

Topological persistence as in [4, 7, 8] is based on sub-level sets. Trying to use topological persistence instead of interval persistence for matching critical points causes difficulties. Returning to Figure 5, assume that $f(p) < f(q)$ whereas $g(p') > g(q')$. Points p and q' are H_1-critical for interval sets. While p' is a critical point for sub-level sets of g, it is H_2-critical for sub-level sets, not H_1-critical, and it creates an H_2 homology element, instead of destroying one. The H_2 homology element created by p' is never destroyed and so p' is never matched and has infinite (or undefined) persistence.

Theorem 3 states that for every critical point q for f with persistence greater than γ, there is a critical point q' for g in the $(\gamma, 2\delta)$-neighborhood of q with persistence greater than $\gamma - 2\delta$. The point q' depends upon the choice of γ. For instance, in Figure 5, if γ is set to $|f(p) - a|$, then point p' is in the $(\gamma, 2\delta)$-neighborhood of p and has persistence greater than γ. Similarly, point q' is in the $(\gamma, 2\delta)$-neighborhood of q and also has persistence greater than γ. If, instead of γ, we use $\tilde{\gamma} = |f(p) - \tilde{a}|$, then p' does not have persistence greater than $\tilde{\gamma}$ (nor does q.) Thus, point p' does not qualify for the role of q' in Theorem 3. Instead, q' is the point in the $(\tilde{\gamma}, 2\delta)$-neighborhood of p which has persistence greater than $\tilde{\gamma}$.

For the proof of Theorem 3, we need to restrict the homology element destroyed by some point q to a connected component containing q. If τ is the connected component

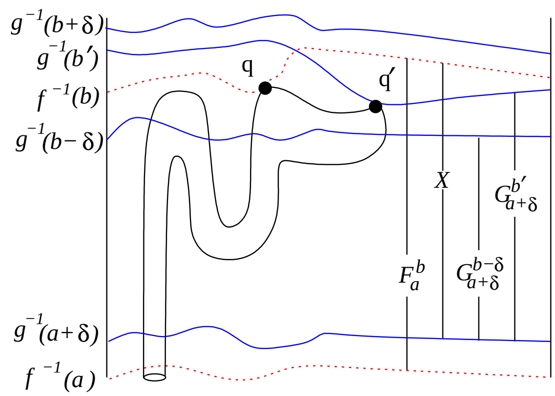

Figure 6: Sets $\underline{F}_a^b$, $X = \{x \in \Omega : a + \delta \leq g(x)$ and $f(x) < b\}$, $\underline{G}_{a+\delta}^{b-\delta}$ and $\underline{G}_{a+\delta}^{b'}$.

of $\underline{F}_a^b \cup \{q\}$ containing q and q destroys some non-zero element of $H_k(f^{-1}(a))$, then q destroys some non-zero element of $H_k(f^{-1}(a) \cap \tau)$. More generally, if q destroys the non-zero image of some element of $H_k(\Omega)$ in $H_k(\underline{F}_a^b \cup \{q\})$ for some $\Omega \subseteq \underline{F}_a^b$, then q destroys the non-zero image of some element of $H_k(\Omega \cap \tau)$ in $H_k(\underline{F}_a^b \cup \{q\})$. The statement of this lemma, Lemma 11, and its proof, are left to the Appendix.

Proof of Theorem 3. Let b equal $f(q)$. Let

$$X = \{x \in \Omega : a + \delta \leq g(x) \text{and} f(x) < b\}.$$

Since $|f(x) - g(x)| < \delta$ and $b - a > 2\delta$, set $g^{-1}(a + \delta)$ is a subset of F_a^b. Note that $\underline{G}_{a+\delta}^{b-\delta} \subseteq X \subseteq \underline{F}_a^b$ and so there exist homomorphisms $H_k(\underline{G}_{a+\delta}^{b-\delta}) \rightarrow H_k(X) \rightarrow H_k(\underline{F}_a^b)$ induced by inclusions. (See Figure 6.)

Let $h'_f \in H_k(\underline{F}_a^b)$ be the non-zero image of h_f under the mapping $H_k(f^{-1}(a)) \rightarrow H_k(\underline{F}_a^b)$. By Lemma 4, point q destroys the image of some non-zero $\hat{h}_g \in H_k(g^{-1}(a+\delta))$ where h'_f is the image of $\hat{h}_g$ under the mapping $H_k(g^{-1}(a + \delta)) \rightarrow H_k(\underline{F}_a^b)$.

Let σ^f be the pathwise connected component of $\underline{F}_a^b \cup \{q\}$ containing q. Let σ equal $\sigma^f \cap g^{-1}(a + \delta)$. The mapping

$$H_k(g^{-1}(a + \delta)) \rightarrow H_k(\underline{F}_a^b) \rightarrow H_k(\underline{F}_a^b \cup \{q\})$$

sends $\hat{h}_g$ to non-zero $h'_f \in H_k(\underline{F}_a^b)$ to zero. By Lemma 11 in the Appendix, there is some non-zero $h_\sigma \in H_k(\sigma)$ such that the mapping

$$H_k(\sigma) \rightarrow H_k(\underline{F}_a^b) \rightarrow H_k(\underline{F}_a^b \cup \{q\})$$

sends h_σ to h'_f to zero.

Let $h_g \in H_k(g^{-1}(a + \delta))$ be the image of h_σ under the mapping $H_k(\sigma) \rightarrow H_k(g^{-1}(a + \delta))$. Let h_x be the image of h_g under the mapping $H_k(g^{-1}(a + \delta)) \rightarrow H_k(X)$. Since the mapping

$$H_k(\sigma) \rightarrow H_k(g^{-1}(a + \delta)) \rightarrow H_k(X) \rightarrow H_k(\underline{F}_a^b)$$

takes h_σ to h_g to h_x to non-zero h'_f, elements h_g and h_x must be non-zero. Since h'_f is destroyed by q, so is h_x.

The following commutative diagram gives the relevant mappings between homology groups:

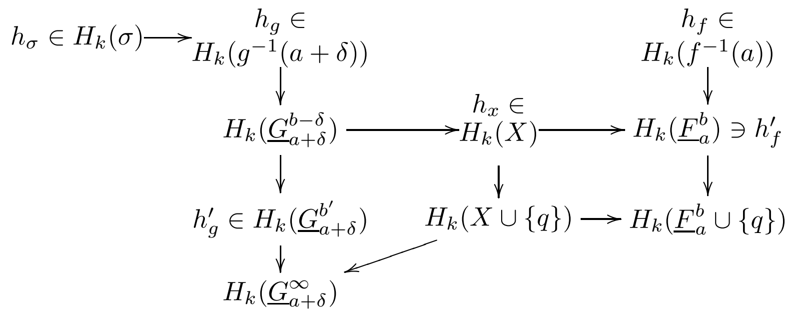

The value b' will be defined below.

Since h_x is destroyed by q, the mapping $H_k(X) \to H_k(X \cup \{q\}) \to H_k(\underline{G}_{a+\delta}^{\infty})$ sends h_x to zero. Thus the composition of mappings $H_k(g^{-1}(a+\delta)) \to H_k(X) \to H_k(\underline{G}_{a+\delta}^{\infty})$ sends h_g to h_x to zero.

By Lemma 2, there exists a point $q' \in \Omega$ such that h_g is destroyed by q' (i.e., the image of h_g under the mapping $H_k(g^{-1}(a+\delta)) \to H_k(\underline{G}_{a+\delta}^{g(q')})$ is destroyed by q'.) Let b' equal $f(q')$. Since $\underline{G}_{a+\delta}^{b-\delta}$ is a subset of X, the image of h_g under the mapping $H_k(g^{-1}(a+\delta)) \to H_k(\underline{G}_{a+\delta}^{b-\delta})$ is non-zero and so

$$b' \geq b - \delta. \tag{4.1}$$

Since $|f-g| \leq \delta$, set $X \cup f^{-1}(b)$ is a subset of $\underline{G}_{a+\delta}^{b+\delta} \cup g^{-1}(b+\delta)$ and so h_g is destroyed by $\underline{G}_{a+\delta}^{b+\delta} \cup g^{-1}(b+\delta)$. Thus,

$$b' \leq b + \delta. \tag{4.2}$$

Let σ^g be the connected component of $\underline{G}_{a+\delta}^{b'} \cup \{q'\}$ containing q'. We claim that σ^f intersects σ^g.

The mapping

$$H_k(\sigma) \to H_k(g^{-1}(a+\delta)) \to H_k(\underline{G}_{a+\delta}^{b'}) \to H_k(\underline{G}_{a+\delta}^{b'} \cup \{q'\})$$

takes h_σ to h_g to some $h'_g \in H_k(\underline{G}_{a+\delta}^{b'})$ to zero. Since point q' destroys h_g, element h'_g must be non-zero. By Appendix Lemma 11, there is some non-zero $h'' \in H_k(\sigma \cap \sigma^g)$ such that $H_k(\sigma \cap \sigma^g) \to H_k(\underline{G}_{a+\delta}^{b'})$ takes h'' to h'_g. Since h'' is non-zero, $H_k(\sigma \cap \sigma^g)$ is not the zero group and so σ intersects σ^g. Since σ is a subset of σ^f, set σ^f also intersects σ^g.

Since $b' \leq b + \delta$ and $|f-g| < \delta$, set $\underline{G}_{a+\delta}^{b'} \cup \{q'\}$ is contained in $\underline{F}_a^{b+2\delta}$. Thus σ^g is contained in $\underline{F}_a^{b+2\delta}$.

Let $\hat{\sigma}$ be the connected component of $\underline{F}_a^{b+2\delta}$ containing q. Its closure, $\mathrm{cl}(\hat{\sigma})$, is the $(\gamma, 2\delta)$-neighborhood of q. Set $\hat{\sigma}$ contains σ^f. Since σ^g intersects σ^f and is connected, set $\hat{\sigma}$ also contains σ^g and thus contains point q'.

The claim $\Pi_k^f(q) > \gamma$ follows from definition since $f(q) - a = \gamma$. The point q' is H_k-critical since it destroys an element of k'th homology. It lies in $\mathrm{cl}(\hat{\sigma})$, the $(\gamma, 2\delta)$-neighborhood of q. Proof of (i) follows from inequalities 4.1 and 4.2. Proof of (ii) follows from (i) and the fact that q' destroys $h_g \in H_k(g^{-1}(a+\delta))$. □

5 Computations

5.1 Computing interval persistence

Theorem 3 can be used to compare two real valued functions f and g defined on a topological space Ω. The key computation to apply Theorem 3 is:

(i) determine if a point p which destroys some $h \in H_k(f^{-1}(a))$ has an interval persistence greater than a given value γ.

We use Betti numbers and their persistent counterparts to compute (i). We discuss the computations for the function f. It is clear that similar computations are needed for g as well. In general, for a point p and a value $a < b = f(p)$ we want to compute if an element of $H_k(f^{-1}(a))$ gets destroyed by p. Let

$$\begin{aligned} \xi_a^b &= \xi_k(f^{-1}(a), \underline{F}_a^b) \text{ and} \\ \lambda_a^p &= \xi_k(f^{-1}(a), \underline{F}_a^{f(p)} \cup \{p\}). \end{aligned}$$

Note that p is a point whereas a and b are real values. The number ξ_a^b counts the generators of $H_k(f^{-1}(a))$ surviving in F_a^b and λ_a^p counts the generators of $H_k(f^{-1}(a))$ surviving in $F_a^b \cup \{p\}$. Therefore,

$$\pi_a^p = \xi_a^b - \lambda_a^p$$

counts the generators of $H_k(f^{-1}(a))$ destroyed by p. So, if $\pi_a^p > 0$, we have an element of $H_k(f^{-1}(a))$ that is destroyed by p where $f(p) = b$.

Let $p_0, p_1, ..., p_m$ be the H_k-critical points of f ordered according to the increasing values, that is, $f(p_i) > f(p_{i-1})$ for all $i \geq 0$. We compute the interval persistence $\Pi_k^f(p_j)$ for these critical points p_i as follows. For $1 \leq i \leq k-1$, let a_i be a value with $f(p_{i-1}) < a_i < f(p_i)$. Compute $\pi_{a_i}^{p_j}$ for any pair i, j where $j \geq i > 0$. Since $\pi_a^{p_j}$ is constant for all a where $f(p_{i-1}) < a < f(p_i)$, if $\pi_{a_i}^{p_j}$ is greater than 0, then the interval persistence $\Pi_k^f(p_j)$ is at least $f(p_j) - f(p_{i-1})$. Thus we compute $\Pi_k^f(p_j)$ as

$$\max_i |f(p_j) - f(p_{i-1})| \text{ so that } \pi_{a_i}^{p_j} > 0.$$

Similarly, we can compute the critical points $q_0, q_1, ..., q_n$ and a set of intermittent values $b_1, b_2, ..., b_{n-1}$ for the function g. The interval persistence of a H_k-critical point q of g is measured similarly by $\Pi_k^g(q)$.

To compare f and g, one can check for each $k \geq 0$, if any critical point p of f has $\Pi_k^f(p)$ greater than a user supplied parameter τ. If so, search for a critical point q of g in the (τ, τ)-neighborhood of q so that $\Pi_k^g(q) > \tau$ and $|f(p) - g(q)| \leq \frac{\tau}{2}$. If $\tau > 2\delta$, such a q exists by Theorem 3.

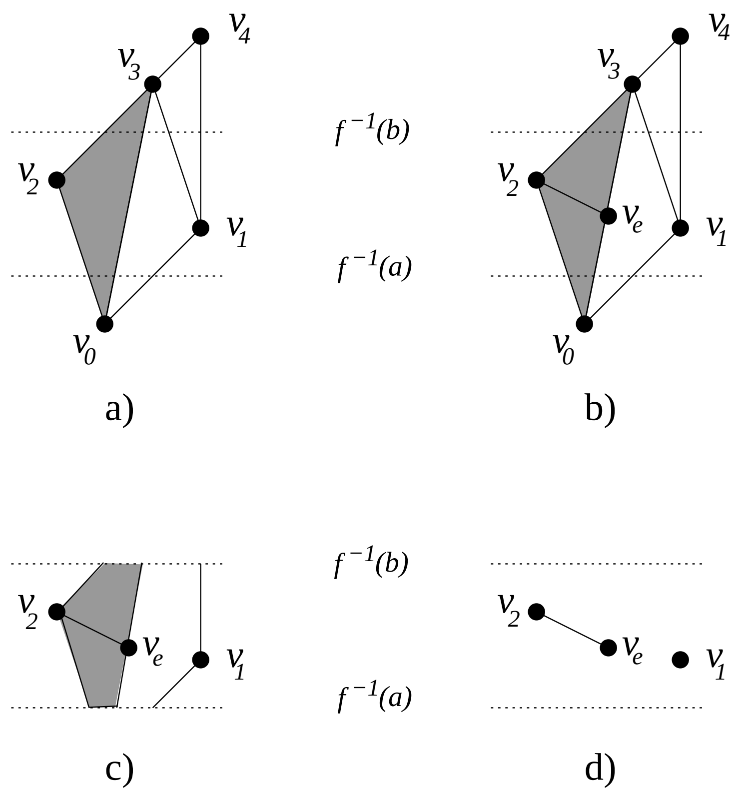

Figure 7: a) Simplicial complex Ω with edges $(v_0, v_1), (v_0, v_2)$, (v_0, v_3), (v_1, v_3), (v_1, v_4), (v_2, v_3), (v_3, v_4) and triangle (v_0, v_2, v_3). For each point (x, y) in the simplicial complex, $f(x, y)$ equals y. Edge (v_0, v_3) intersects both $F^a_{-\infty}$ and $\underline{F}^{\infty}_b$. b) Subdivision of edge (v_0, v_3) and triangle (v_0, v_2, v_3) by starring from v_e. c) Set $\underline{F}^b_a$. d) Deformation retract of $\underline{F}^b_a$ onto subcomplex Υ.

5.2 PL case

Assume there is some finite triangulation of Ω such that f and g are linear on each simplex of the triangulation. Functions f and g are LR (locally retractible), but not necessarily point destructible. A small perturbation of the scalar value at each triangulation vertex and the linear interpolation of those values over the triangulation simplices, gives new piecewise linear functions where no two vertices have the same scalar value. These new functions are point destructible and their vertices are located at the triangulation vertices.

Zomorodian and Carlsson in [13] show how to compute persistent Betti numbers for homology groups of filtered simplicial complexes over any field. However, spaces $\underline{F}^b_a$ and $\underline{F}^b_a \cup \{p\}$ are not closed. To compute their persistent Betti numbers, ξ^b_a and λ^p_a, we subdivide the simplicial complex and then extract a subcomplex with the same homology groups as $\underline{F}_a^b$ or $\underline{F}_a^b \cup \{p\}$.

Consider the space $\underline{F}_a^b$. Let $t \subseteq \Omega$ be a simplex which intersects this space. If t is contained in $F_{-\infty}^b$ or t is contained in $\underline{F}_a^\infty$, then $t \cap F_a^b$ can be collapsed to the face of t contained in $\underline{F}_a^b$. Thus if for every simplex $t \in \Omega$ either $t \subseteq F_{-\infty}^b$ or $t \subseteq \underline{F}_a^\infty$, then $\underline{F}_a^b$ can be collapsed to the subcomplex $\Upsilon = \{t \in \Omega : t \subseteq \underline{F}_a^b\}$. However, there may be some simplices $t \in \Omega$ for which $t \cap F_{-\infty}^a \neq \emptyset$ and $t \cap \underline{F}_b^\infty \neq \emptyset$. We subdivide Ω to eliminate such simplices.

For each edge $e \in \Omega$, where $e \cap F_{-\infty}^a \neq \emptyset$ and $e \cap \underline{F}_b^\infty \neq \emptyset$, add a vertex v_e in $e \cap \underline{F}_a^b$. Subdivide all the simplices t containing e by starring from v_e to all the vertices of t. Repeat this until every edge of the complex is either a subset of $F_{-\infty}^b$ or of $\underline{F}_a^\infty$. Let $\tilde{\Omega}$ be the resulting subdivision of Ω. Every edge in $\tilde{\Omega}$ is either a subset of $F_{-\infty}^b$ or of $\underline{F}_a^\infty$. (See Figure 7.)

We claim that for every simplex $t \in \tilde{\Omega}$, either $t \subseteq F_{-\infty}^b$ or $t \subseteq \underline{F}_a^\infty$. If not, then t has some vertex $v \in F_{-\infty}^a$ and some vertex $v' \in \underline{F}_b^\infty$. Edge $(v, v') \in \tilde{\Omega}$ intersects both $F_{-\infty}^a$ and $\underline{F}_b^\infty$, a contradiction. Set $\Upsilon = \{t \in \tilde{\Omega} : t \subseteq \underline{F}_a^b\}$ is a deformation retract of $\underline{F}_a^b$ and therefore has homology groups isomorphic to that of $\underline{F}_a^b$. Apply the algorithm from Zomorodian and Carlsson in [13] to compute the persistent Betti number β_a^b of Υ. The persistent Betti number β_a^b of Υ equals ξ_a^b, the number of generators in $H_k(f^{-1}(a))$ which destroyed in $\underline{F}_a^b$. A similar algorithm can be used to compute λ_a^b.

6 Maxima

In this section we establish a stronger result for critical points that are maxima on oriented manifolds. We show that the (γ, ∞)-neighborhoods of local maxima with interval persistence greater than γ are pairwise disjoint. This enables us to establish a matching of such critical points. Observe that Theorem 3 does not imply the disjointness of neighborhoods.

The idea of the proof is as follows. Consider two local maxima, $p_0, p_1 \in \Omega$, with interval persistence greater than γ. Let a_0 equal $f(p_0) - \gamma$ and a_1 equal $f(p_1) - \gamma$. Let σ_0 and σ_1 be the (γ, ∞)-neighborhoods of p_0 and p_1, respectively. Set σ_i is a connected component of $\underline{F}_{a_i}^\infty$. Without loss of generality, assume that $f(p_0)$ is less than or equal to $f(p_1)$.

Assume that σ_0 intersects σ_1. Since σ_0 intersects σ_1 and $a_0 \leq a_1$, set σ_1 is contained in σ_0. Therefore, p_1 is contained in σ_0.

Since p_0 has persistence greater than γ, point p_0 destroys some non-zero $h \in H_k(\partial \sigma_0)$. Since p_0 is a local maximum, k equals $d - 1$. Under appropriate assumptions, σ_0 is a manifold with boundary. We then show that $\underline{F}_{a_0}^{f(p_0)} \cup \{p_0\}$ contains σ_0. Since $f(p_0) \leq f(p_1)$, set $\underline{F}_{a_0}^{f(p_0)} \cup \{p_0\}$ does not contain p_1 and therefore point p_1 is not in σ_0, a contradiction. We conclude that σ_0 does not intersect σ_1.

We start the formal proof with a lemma about the homology groups destroyed by local maxima of d-manifolds. Lemmas 5 and 6 depend on some lemmas from homological algebra whose proofs are left to the appendix.

Lemma 5. *Let $\underline{F}_a^\infty$ be an oriented d-manifold with boundary and let σ be the connected component of $\underline{F}_a^\infty$ containing p. If p is a local maximum and p destroys $h \in H_k(f^{-1}(a))$, then k equals $d-1$ and p destroys the image of some $h' \in H_{d-1}(\partial\sigma)$ in $H_{d-1}(\underline{F}_a^{f(p)})$.*

Proof.

Part I: Show that k equals $d-1$.

Since $\underline{F}_a^\infty$ is an oriented d-manifold and p is a local maximum, some neighborhood N_p of p is homeomorphic to $\mathbb{R}^d$ and all points in $N_p - \{p\}$ have value less than $f(p)$. Let B be the unit ball in $\mathbb{R}^d$, and let B_p be its image under the homeomorphism from $\mathbb{R}^d$ to N_p. Since all points in $N_p - \{p\}$ have value less than $f(p)$, they are all in $\underline{F}_a^{f(p)}$.

Let $h' \in H_k(\underline{F}_a^{f(p)})$ be the non-zero image of $h \in H_k(f^{-1}(a))$ under the mapping $H_k(f^{-1}(a)) \rightarrow H_k(\underline{F}_a^{f(p)})$. The mapping $H_k(\underline{F}_a^{f(p)}) \rightarrow H_k(\underline{F}_a^{f(p)} \cup \{p\})$ sends h' to zero. The intersection of $\underline{F}_a^{f(p)}$ and B_p is $B_p - \{p\}$ and their union is $\underline{F}_a^{f(p)} \cup \{p\}$. By Corollary 8 (Appendix), h' is the image of some $h'' \in H_k(B_p - \{p\})$ under the mapping $H_k(B_p - \{p\}) \rightarrow H_k(\underline{F}_a^{f(p)})$. Since h' is non-zero, h'' is non-zero. Since $H_k(B_p - \{p\})$ is the zero group, for all $k \neq d-1$, element h'' must be in $H_{d-1}(B_p - \{p\})$. Therefore, h and h' are elements of $H_{d-1}(f^{-1}(a))$ and $H_{d-1}(\underline{F}_a^{f(p)})$, respectively, and so k equals $d-1$.

Part II: Show that p destroys the image of some $h' \in H_{d-1}(f^{-1}(\partial\sigma))$ in $H_{d-1}(\underline{F}_a^{f(p)})$.

Let Ω equal $f^{-1}(a)$ and let Ω' equal $\underline{F}_a^{f(p)} \cup \{p\}$. Let τ' be the connected component of Ω' containing p. Let τ equal $\tau' \cap \Omega$. Note that τ equals $\partial\sigma$.

The mapping $H_{d-1}(\Omega) \rightarrow H_{d-1}(\Omega')$ sends $h \in H_{d-1}(\Omega)$ to zero while the mapping $H_{d-1}(\Omega) \rightarrow H_{d-1}(\Omega' - \{p\})$ does not. By Lemma 11 (Appendix), the mapping $H_{d-1}(\tau) \rightarrow H_{d-1}(\Omega')$ sends some non-zero $h' \in H_{d-1}(\tau)$ to zero while the mapping $H_{d-1}(\tau) \rightarrow H_{d-1}(\Omega' - \{p\})$ does not. Thus p destroys the image of some $h' \in H_{d-1}(\partial\sigma) = H_{d-1}(\tau)$ in $H_{d-1}(\underline{F}_a^{f(p)})$. □

Let M be a connected, oriented d-manifold with non-empty boundary. We show that elements of $H_{d-1}(\partial M)$ are killed only by M, not by any subset of M.

Lemma 6. *Let M be a connected, oriented d-manifold with non-empty boundary. If $\partial M \subseteq M' \subseteq M$ and $H_{d-1}(\partial M) \rightarrow H_{d-1}(M')$ takes some non-zero $h \in H_{d-1}(\partial M)$ to zero, then M' equals M.*

Proof. Assume M' does not equal M. Let p be a point in $M - M'$. Let B be an open topological ball containing p whose closure does not intersect ∂M. There exists a deformation retract from $M - \{p\}$ to $M - B$. Thus $H_{d-1}(M - \{p\})$ is isomorphic to $H_{d-1}(M - B)$.

The mapping $H_{d-1}(\partial M) \rightarrow H_{d-1}(M') \rightarrow H_{d-1}(M - \{p\}) \rightarrow H_{d-1}(M - B)$ sends h to zero in $H_{d-1}(M - B)$. Let h_0 be the element of the homology group of $H_{d-1}(\partial M)$ generated by ∂M with orientation inherited from M. By Lemma 19 (Appendix), element h equals αh_0 for some non-zero α. Let h_B be the element of

$H_{d-1}(M-B)$ generated by ∂B with orientation inherited from $M - B$. Let h' be the image of h under the map $H_{d-1}(\partial M) \to H_{d-1}(\partial M \cup \partial B)$. The element h and hence h' is sent to zero in $H_{d-1}(M - B)$. By Lemma 19 (Appendix), element h' equals $\beta(h_0 + h_B)$ for some non-zero β. Thus αh_0 equals $\beta(h_0 + h_B)$. Since h_0 and h_B are linearly independent, α and β are both zero implying h is a zero element, a contradiction. It follows that M equals M'. □

Let

$$\sigma_p^f(\gamma) = (\gamma, \infty) - \text{neighborhood of } p \text{ for } f.$$

We prove that the neighborhoods $\sigma_p^f(\gamma)$ of maxima with interval persistence greater than γ are pairwise disjoint.

Theorem 4. *Let $f : \Omega \to \mathbb{R}$ be a continuous function such that $\underline{F}_a^\infty$ is a $(d-1)$-manifold with boundary for all but a finite number of a. If points $p_0, p_1 \in \Omega$ are local maxima with interval persistence greater than γ, then $\sigma_{p_0}^f(\gamma)$ does not intersect $\sigma_{p_1}^f(\gamma)$.*

Proof. Let $p_0, p_1 \in \Omega$ be local maxima with persistence γ_0, γ_1, both greater than γ. Without loss of generality, assume that $f(p_0) \leq f(p_1)$.

Assume that $\sigma_{p_0}^f(\gamma)$ intersects $\sigma_{p_1}^f(\gamma)$. Since $\underline{F}_a^\infty$ is a $(d-1)$-manifold for all but a finite number of a, there is some $\gamma' \geq \gamma$ such that $\gamma_0 > \gamma'$ and $\gamma_1 > \gamma'$ and $\underline{F}_{f(p_0)-\gamma'}^\infty$ is a $(d-1)$-manifold with boundary. Since $\sigma_{p_0}^f(\gamma)$ intersects $\sigma_{p_1}^f(\gamma)$, set $\sigma_{p_0}^f(\gamma')$ intersects $\sigma_{p_1}^f(\gamma')$.

Since $f(p_0) \leq f(p_1)$, set $\underline{F}_{f(p_0)-\gamma'}^\infty$ contains $\underline{F}_{f(p_1)-\gamma'}^\infty$. Since $\sigma_{p_0}^f(\gamma')$ intersects $\sigma_{p_1}^f(\gamma')$, set $\sigma_{p_0}^f(\gamma')$ contains $\sigma_{p_1}^f(\gamma')$. Thus $\sigma_{p_0}^f(\gamma')$ contains p_1.

By Lemma 3, point p_0 destroys some non-zero element of $H_k(f^{-1}(f(p_0) - \gamma'))$. By Lemma 5, k equals $d-1$ and p_0 destroys a non-zero elements of $H_{d-1}(\partial\sigma_{p_0}^f(\gamma'))$. In Lemma 6 putting M equal to $\sigma_{p_0}^f(\gamma')$ and M' equal to the connected component of $\underline{F}_{f(p_0)-\gamma'}^{f(p_0)} \cup \{p_0\}$ containing p_0, we conclude $\underline{F}_{f(p_0)-\gamma'}^{f(p_0)} \cup \{p_0\}$ equals $\sigma_{p_0}^f(\gamma')$ and thus contains p_1. However, since $f(p_0) \leq f(p_1)$, set $\underline{F}_a^{f(p_0)} \cup \{p_0\}$ does not contain p_1. Thus, $\sigma_{p_0}^f(\gamma)$ does not intersect $\sigma_{p_1}^f(\gamma)$. □

Our final theorem gives relationships between neighborhoods of local maxima of f and of g. See Figure 8 where $\delta = \gamma/4$.

Theorem 5. *Let $f, g : \Omega \to \mathbb{R}$ be continuous functions such that $\underline{F}_a^\infty$ and $\underline{G}_a^\infty$ are $(d-1)$-manifolds with boundary for all but a finite number of a and $|f - g| < \delta$. Let $p \in \Omega$ be a local maximum of f and let q and q' be local maxima of g where $q \neq q'$. Assume that p, q, q' have persistence greater than γ and $|f(p) - g(q)| < \delta$ and $|f(p) - g(q')| < \delta$.*

(i) If $\sigma_p^f(\gamma - 2\delta)$ intersects $\sigma_q^g(\gamma - 2\delta)$, then $\sigma_q^g(\gamma)$ contains $\sigma_p^f(\gamma - 2\delta)$.

(ii) If $\sigma_p^f(\gamma-2\delta)$ intersects $\sigma_q^g(\gamma-2\delta)$, then $\sigma_p^f(\gamma-2\delta)$ does not intersect $\sigma_{q'}^g(\gamma-2\delta)$.

Proof of (i). Let y be a point in $\sigma_p^f(\gamma-2\delta)\cap\sigma_q^g(\gamma-2\delta)$ and z be any point in $\sigma_p^f(\gamma-2\delta)$. Set $\sigma_p^f(\gamma-2\delta)$ is path connected, so there is a path $\zeta \subseteq \sigma_p^f(\gamma-2\delta)$ from y to z. Since $\zeta \subseteq \sigma_p^f(\gamma-2\delta)$, $f(x) \geq f(p)-\gamma+2\delta$ for every point $x \in \zeta$. Since $|f(x)-g(x)| < \delta$ for all $x \in \Omega$ and $|f(p)-g(q)| < \delta$, it follows that $g(x) \geq g(q)-\gamma$ for all $x \in \zeta$. Thus $\zeta \subseteq \sigma_q^g(\gamma)$ and z lies in $\sigma_q^g(\gamma)$. This holds for all $z \in \sigma_p^f(\gamma-2\delta)$ so $\sigma_q^g(\gamma)$ contains $\sigma_p^f(\gamma-2\delta)$. (See the neighborhoods of p_2 and q_2 in Figure 8.) □

Proof of (ii). By Theorem 4, $\sigma_q^g(\gamma)$ and $\sigma_{q'}^g(\gamma)$ are disjoint. By (i) above, if $\sigma_p^f(\gamma-2\delta)$ intersects $\sigma_{q'}^g(\gamma-2\delta)$, then $\sigma_{q'}^g(\gamma)$ contains $\sigma_p^f(\gamma-2\delta)$. Then, by condition of (ii) $\sigma_{q'}^g(\gamma)$ intersects $\sigma_q^g(\gamma-2\delta)$ and hence $\sigma_q^g(\gamma)$. However, $\sigma_q^g(\gamma)$ and $\sigma_{q'}^g(\gamma)$ are disjoint. Therefore, $\sigma_p^f(\gamma-2\delta)$ cannot intersect $\sigma_{q'}^g(\gamma-2\delta)$ (See Figure 8). □

7 Matching

We assume that $f, g : \Omega \to \mathbb{R}$ are continuous, point destructible, LR functions such that $\underline{F}_a^\infty$ and $\underline{G}_a^\infty$ are $(d-1)$-manifolds with boundary for all but a finite number of a. Let M_f and M_g be the set of local maxima of f and g, respectively, and let $M_f(\gamma) \subseteq M_f$ and $M_g(\gamma) \subseteq M_g$ be the set of local maxima of f and g, respectively, with persistence greater than γ. We would like to match points in $M_f(\gamma)$ with close points in $M_g(\gamma)$ in the sense of Theorem 3. However, there may be no such matching. In fact, f may contain a set of maxima with persistence a little bit above γ while nearby critical points in g all have persistence a bit below γ. Thus, $M_f(\gamma)$ can contain any number of points while $M_g(\gamma)$ is empty! Instead of matching $M_f(\gamma)$ and $M_g(\gamma)$ only with each other, we allow them to match with points with slightly less persistence.

We say that a *partial matching* of M_f with M_g *covers* $M_f(\gamma)$ and $M_g(\gamma)$ if all points in $M_f(\gamma)$ and $M_g(\gamma)$ are matched. A partial matching of M_f and M_g is *(α,β)-close* if for each pair (p,q) where $p \in M_f$ and $q \in M_g$, point q lies in $\sigma_p^f(\alpha)$ and point p lies in $\sigma_q^g(\alpha)$ and $|f(p)-g(q)| < \beta$.

Assume that $|f(x)-g(x)| < \gamma/4$ for all $x \in \Omega$. We will find a partial matching of M_f with M_g which covers $M_f(\gamma)$ and $M_g(\gamma)$ and is $(\gamma, \gamma/4)$-close.

The algorithm is as follows. For each point $p \in M_f(\gamma)$, we compute $\sigma_p = \sigma_p^f(\gamma/2)$. Similarly, for each $q \in M_g(\gamma)$, we compute $\sigma_q = \sigma_q^g(\gamma/2)$. By Theorem 5, each σ_p intersects at most one σ_q where $|f(p)-g(q)| < \gamma/4$ and vice versa. If σ_p intersects such a σ_q, then match p with q. If not, then match p with some $q' \in M_g(\gamma/2)$ lying in σ_p such that $|f(p)-g(q')| < \gamma/4$. (By Theorem 3 such a q' exists.) Similarly, if σ_q does not intersect any σ_p, match q with $p' \in M_f(\gamma/2)$ lying in σ_q such that $|f(p')-g(q)| < \gamma/4$.

We claim that algorithm MatchPersistentMax matches all maxima with persistence more than γ:

Proposition 1. *If $|f-g| \leq \gamma/4$, then* MatchPersistentMax(Ω, f, g, γ) *produces a partial matching of M_f with M_g which is a subset of $M_f(\gamma/2) \times M_g(\gamma/2)$ and which covers $M_f(\gamma)$ and $M_g(\gamma)$ and which is $(\gamma, \gamma/4)$-close.*

Proof. By construction, $p \in M_f(\gamma/2)$ and $q \in M_g(\gamma/2)$ for every matched pair (p,q).

MatchPersistentMax(Ω, f, g, γ)
/* $f, g : \Omega \to \mathbb{R}$ */

1 Compute sets $M_f(\gamma)$ and $M_g(\gamma)$;
2 For each point $p \in M_f(\gamma)$, compute $\sigma_p = \sigma_p^f(\gamma/2)$;
3 For each point $q \in M_g(\gamma)$, compute $\sigma_q = \sigma_q^g(\gamma/2)$;
4 For each point $p \in M_f(\gamma)$ and $q \in M_g(\gamma)$, if $\sigma_p \cap \sigma_q \neq \emptyset$ and $|f(p) - g(q)| < \gamma/4$, then match p with q;
5 For each unmatched $p \in M_f(\gamma)$, match p with $q' \in M_g(\gamma/2) \cap \sigma_p$ where $|f(p) - g(q')| < \gamma/4$;
6 For each unmatched $q \in M_g(\gamma)$, match q with $p' \in M_f(\gamma/2) \cap \sigma_q$ where $|f(p') - g(q)| < \gamma/4$.

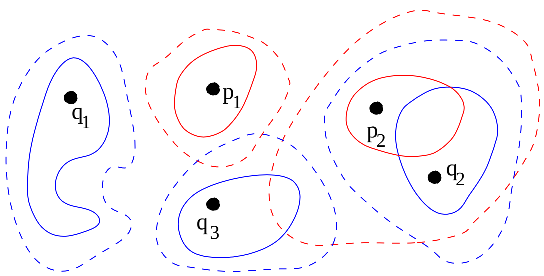

Figure 8: Local maxima and their neighborhoods. Solid lines around points p_i and q_i are neighborhoods $\sigma_{p_i}^f(\gamma/2)$ and $\sigma_{q_i}^g(\gamma/2)$. Dotted lines are neighborhoods $\sigma_{p_i}^f(\gamma)$ and $\sigma_{q_i}^g(\gamma)$. Neighborhood $\sigma_{p_2}^f(\gamma/2)$ intersects $\sigma_{q_2}^g(\gamma/2)$ so p_2 matches with q_2. Point p_1 matches with some point (not shown) from $M_g(\gamma/2)$ and points q_1 and q_3 match with points (not shown) from $M_f(\gamma/2)$.

By Theorem 4, the $\sigma_p = \sigma_p^f(\gamma/2)$, $p \in M_f(\gamma)$, are pairwise disjoint and the $\sigma_q = \sigma_q^g(\gamma/2)$, $q \in M_g(\gamma)$ are pairwise disjoint. By Theorem 5, each σ_p intersects at most one σ_q and vice versa. Thus Step 4 gives a one to one partial matching.

By Theorem 3, σ_p contains some point $q' \in M_g(\gamma/2)$ such that $|f(p) - g(q')| \leq \gamma/4$. Since (p, q') is not matched in Step 4, point q' is not in $M_g(\gamma)$. Thus point q' is not matched in Step 4. Since σ_p does not intersect any $\sigma_{p'}$, $p' \in M_f(\gamma)$, point q' is matched to at most one p in Step 5. Similarly, point p' in Step 6 is not matched in Steps 4 and 5 and is matched to at most one q in Step 6. Thus the matching is one to one and covers all of $M_f(\gamma)$ and $M_g(\gamma)$.

It remains to show that for each match (p, q), set $\sigma_p^f(\gamma)$ contains q and $\sigma_q^g(\gamma)$ contains p. If p and q are matched in Step 4, then $\sigma_p^f(\gamma/2)$ intersects $\sigma_q^g(\gamma/2)$. This holds true even if p and q are matched in Steps 5 or 6. By Theorem 5 with $\delta = \gamma/4$, $\sigma_p^f(\gamma)$ contains $\sigma_q^g(\gamma/2)$ which contains q and $\sigma_q^g(\gamma)$ contains $\sigma_p^f(\gamma/2)$ which contains p. Since points $p \in M_f$ and $q \in M_g$ are only matched if $|f(p) - g(q)| < \gamma/4$, the matching is $(\gamma, \gamma/4)$-close. $\square$

8 Discussions

Results on stability of topological persistence can be applied to shape distance functions. If we take a dense point sample from the boundary of a shape, the distance functions to the shape boundary and and its point sample are similar. Therefore, the results on topological persistence apply to the shape distance functions and their approximations by point samples. Previous works [2, 3, 4] have noted this application of topological persistence. It would be interesting to apply the results of this paper to these functions as well. Notably, our results in this paper have some connections to a shape matching algorithm proposed in [6]. According to our results, we can expect that distance functions of similar shapes have similar neighborhoods for maxima with large interval persistence. The algorithm in [6] uses maxima and their stable manifolds for matching. We suspect that these stable manifolds are playing the role of neighborhoods as suggested in this paper. Perhaps the performance of the matching algorithm in [6] now can be improved and better explained by our results. We plan to address this issue in future work.

This research brings up some other interesting questions. We have obtained a stronger result for maxima than other critical points. Is it possible to extend this stronger result to other critical points? We have given an algorithm to compute interval persistence. How can this algorithm be made more efficient? Is there an efficient algorithm along the line of Cohen-Steiner et al. [5]? Theorem 2 tells us that for most of the critical points we can use the linear time algorithm recently discovered by Cohen-Steiner et al. [5] for topological persistence. It would be interesting to see how one may compute the interval persistence of the critical points that remain unpaired by topological persistence.

Acknowledgement

We thank Prof. Michael Davis from OSU Dept. of Mathematics for providing the proof of Lemma 19. We also thank Frédéric Chazal and David Cohen-Steiner for helpful comments on earlier versions of this paper.

References

[1] P. K. Agarwal, H. Edelsbrunner, J. Harer, and Y. Wang. Extreme elevation on a 2-manifold. In *SCG '04: Proceedings of the twentieth annual symposium on Computational geometry*, pages 357–365, New York, NY, USA, 2004. ACM Press.

[2] G. Carlsson, A. Collins, L. Guibas, and A. Zomorodian. Persistence barcodes for shapes. In *Proc. 2nd Sympos. Geometry Process.*, pages 127–138, 2004.

[3] F. Chazal and A. Lieutier. Weak feature size and persistent homology: computing homology of solids in $\mathbb{R}^n$ from noisy data samples. In *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, pages 255–262, 2005.

[4] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. In *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, pages 263–271, 2005.

[5] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. Vines and vineyards by updating persistence in linear time. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, 2006, to appear.

[6] T. K. Dey, J. Giesen, and S. Goswami. Shape segmentation and matching with flow discretization. In *Proc. 8th Workshop on Algorithms Data Structures*, pages 25–36, 2003.

[7] H. Edelsbrunner, J. Harer, and A. Zomorodian. Hierarchical Morse Complexes for piecewise linear 2-manifolds. *Discrete and Computational Geometry*, 30:87–107, 2003.

[8] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.

[9] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[10] M. W. Hirsch. *Differential Topology*. Springer, 1976.

[11] J. Milnor. *Morse Theory*. Princeton University Press, 1963.

[12] J. J. Rotman. *An Introduction to Algebraic Topology*. Springer-Verlag, 1988.

[13] A. Zomorodian and G. Carlsson. Computing perisistent homology. *Discrete Comput. Geom.*, 33:249–274, 2005.

A Appendix: Homology

A sequence of groups G_i connected by homomorphisms form an *exact sequence* if any two consecutive homomorphisms in the sequence

$$\ldots \to G_i \stackrel{\ell_i}{\to} G_{i+1} \stackrel{\ell_{i+1}}{\to} G_{i+2} \to \ldots$$

satisfy the property that

$$\text{Im } \ell_i = \text{Ker } \ell_{i+1}.$$

Let $A, B \subset X$ so that X is the union of the interiors of A and B and $D = A \cap B$. The sequence

$$H_k(D) \stackrel{\Phi}{\to} H_k(A) \oplus H_k(B) \stackrel{\Psi}{\to} H_k(X) \stackrel{\partial}{\to} H_{k-1}(D)$$

is exact and is called the *Mayer-Vietoris sequence* [9, p. 149]. The map ∂ is the connecting homomorphism given by boundary maps [9, p. 116].

We need the following lemmas and corollaries about exact sequences.

Lemma 7. *Let*

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

be an exact sequence. If $H_k(A) \to H_k(A \cup B)$ maps $h_a \in H_k(A)$ to zero, then there exists some $h \in H_k(A \cap B)$ such that $H_k(A \cap B) \to H_k(A) \to H_k(A \cup B)$ maps h to h_a to zero.

Proof. The mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ sends $h_a \oplus 0 \in H_k(A) \oplus H_k(B)$ to zero. Since

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact, $h_a \oplus 0$ is the image of some $h \in H_k(A \cap B)$ under the mapping $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$. Thus h_a is the image of $h \in H_k(A \cap B)$ under the mapping $H_k(A \cap B) \to H_k(A)$. □

We get the following corollary for subsets A and B of a topological space X where X is the union of the interiors of A and B.

Corollary 8. *Let A and B be subsets of a topological space X where X is the union of the interiors of A and B. If $H_k(A) \to H_k(A \cup B)$ maps $h_a \in H_k(A)$ to zero, then there exists some $h \in H_k(A \cap B)$ such that $H_k(A \cap B) \to H_k(A) \to H_k(A \cup B)$ maps h to h_a to zero.*

Proof. By Mayer-Vietoris, the sequence

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact. The result follows by Lemma 7. □

Note that the interiors of A and B in Corollary 8 are taken with respect to X. Thus points on the boundary of X may lie on the interior of A or B.

Lemma 9. *Let*

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

be an exact sequence. If $H_k(A) \to H_k(A \cup B)$ maps $h_a \in H_k(A)$ to $h' \in H_k(A \cup B)$ and $H_k(B) \to H_k(A \cup B)$ maps $h_b \in H_k(B)$ to the same $h' \in H_k(A \cup B)$ then there exists some $h \in H_k(A \cap B)$ such that $H_k(A \cap B) \to H_k(A) \to H_k(A \cup B)$ maps h to h_a to h' and $H_k(A \cap B) \to H_k(B) \to H_k(A \cup B)$ maps h to h_b to h'.

Proof. The mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ sends $(h_a \oplus h_b) \in H_k(A) \oplus H_k(B)$ to $(h' - h') = 0$. Since

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact, $(h_a \oplus h_b)$ is the image of some $h \in H_k(A \cap B)$ under the mapping $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$. Thus h_a is the image of h under the mapping $H_k(A \cap B) \to H_k(A)$ and h_b is the image of h under the mapping $H_k(A \cap B) \to H_k(B)$. □

We have the following corollary for Lemma 9.

Corollary 10. *Let A and B be subsets of a topological space X where X is the union of the interiors of A and B. If $H_k(A) \to H_k(A \cup B)$ maps $h_a \in H_k(A)$ to $h' \in H_k(A \cup B)$ and $H_k(B) \to H_k(A \cup B)$ maps $h_b \in H_k(B)$ to the same $h' \in H_k(A \cup B)$ then there exists some $h \in H_k(A \cap B)$ such that $H_k(A \cap B) \to H_k(A) \to H_k(A \cup B)$ maps h to h_a to h' and $H_k(A \cap B) \to H_k(B) \to H_k(A \cup B)$ maps h to h_b to h'.*

Proof. By Mayer-Vietoris, the sequence

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact. The result follows by Lemma 9. □

The following lemma restricts the homology element destroyed by some point q to a connected component containing q. The lemma is used in the proof of Theorem 3 on the stability of critical points and in the proof of Lemma 5 on the local maxima of manifolds.

Lemma 11. *Let $\Omega \subseteq \Omega'$ be topological spaces, let q be a point in $\Omega' - \Omega$, let σ' be the pathwise connected component of Ω' containing q and let σ equal $\sigma' \cap \Omega$. If the mapping $H_k(\Omega) \to H_k(\Omega' - \{q\}) \to H_k(\Omega')$, sends non-zero $h \in H(\Omega)$ to non-zero $h' \in H_k(\Omega' - \{q\})$ to zero in $H_k(\Omega')$, then for some non-zero $h_\sigma \in H_k(\sigma)$ the mapping $H_k(\sigma) \to H_k(\Omega' - \{q\}) \to H_k(\Omega')$ sends h_σ to h' to zero.*

Proof. Let A equal $\Omega \cup (\Omega' - \sigma')$ and let B equal σ'. Note that $A \cup B$ equals Ω' while $A \cap B$ equals σ.

The mapping

$$H(\Omega) \to H(A) \to H(\Omega' - \{q\}) \to H(\Omega')$$

takes h to some $h_a \in H(A)$ to non-zero $h' \in H(\Omega' - \{q\})$ to zero in $H(\Omega')$. Since h' is non-zero, element h_a is also non-zero.

By Corollary 8, there is some $h_\sigma \in H_k(A \cap B) = H_k(\sigma)$ such that $H_k(\sigma) \to H_k(A) \to H_k(\Omega')$ takes h_σ to h_a to zero. Since $H_k(A) \to H_k(\Omega' - \{q\}) \to H_k(\Omega')$ sends h_a to h' to zero, the mapping $H_k(\sigma) \to H_k(\Omega' - \{q\}) \to H_k(\Omega')$ sends h_σ to h' to zero. Since h' is non-zero, element h_σ is non-zero. □

B Appendix: Critical points

Theorem 1 relates critical points for interval sets to critical points for sub-level sets. Whether a point is H_k-critical depends upon whether certain mappings of homology groups are isomorphisms. We present here the lemmas about isomorphisms of homology groups which are the basis of Theorem 1. We also prove that if point q destroys some element of H_k, then $f(q)$ is an H_k-critical value.

A short exact sequence is an exact sequence $0 \to H_k(X) \to H_k(Y) \to H_k(Z) \to 0$ where $H_k(X) \to H_k(Y)$ is an injection (one to one) and $H_k(Y) \to H_k(Z)$ is a surjection (onto) and the image of $H_K(X) \to H_k(Y)$ is the kernel of $H_k(Y) \to H_k(Z)$.

Lemma 12. *Let*

$$0 \to H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B) \to 0$$

be a short exact sequence of homology groups. Mapping $H_k(A \cap B) \to H_k(A)$ is an isomorphism if and only if $H_k(B) \to H_k(A \cup B)$ is an isomorphism.

Proof. Part I: Assume that $H_k(A \cap B) \to H_k(A)$ is an isomorphism. Show that $H_k(B) \to H_k(A \cup B)$ is an isomorphism.

Assume that $H_k(B) \to H_k(A \cup B)$ was not injective. By definition, there exists a non-zero $h_b \in H_k(B)$ such that $H_k(B) \to H_k(A \cup B)$ maps h_b to zero. Thus, the mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ maps $0 \oplus h_b$ to zero. Since

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact, $0 \oplus h_b$ is the image of some $h \in H_k(A \cap B)$ under the mapping $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$. Since $H_k(A \cap B) \to H_k(A)$ is an isomorphism, element h must equal zero. Since h_b is the image of zero under the mapping $H_k(A \cap B) \to H_k(B)$, element h_b must also equal zero, a contradiction. Thus $H_k(B) \to H_k(A \cup B)$ is injective.

Assume that $H_k(B) \to H_k(A \cup B)$ was not surjective. By definition, there exists a non-zero $h \in H_k(A \cup B)$ which is not in the image of $H_k(B) \to H_k(A \cup B)$. Since $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ is surjective, element h is the image of some $h_a \oplus h_b$ under the mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$. Since $H_k(A \cap B) \to H_k(A)$ is an isomorphism, element h_a is the image of some $h' \in H_k(A \cap B)$ under the mapping $H_k(A \cap B) \to H_k(A)$. Let h'_b be the image of h' under the mapping $H_k(A \cap B) \to H_k(B)$. Since

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact, the mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ sends $h_a \oplus h'_b$ to zero. Since this mapping takes $h_a \oplus h_b$ to h and $(h_a \oplus h_b) - (h_a \oplus h'_b)$ equals $0 \oplus (h_b - h'_b)$, this mapping takes $0 \oplus (h_b - h'_b)$ to h. Element h is the image of $h_b - h'_b$, contradicting the assumption. Thus $H_k(B) \to H_k(A \cup B)$ is surjective.

Part II: Assume that $H_k(B) \to H_k(A \cup B)$ is an isomorphism. Show that $H_k(A \cap B) \to H_k(A)$ is an isomorphism.

Assume that $H_k(A \cap B) \to H_k(A)$ is not injective. By definition, there exists an $h \in H_k(A \cap B)$ such that $H_k(A \cap B) \to H_k(A)$ maps h to zero. Since $H_k(A) \to H_k(A \cup B)$ sends zero to zero, the mapping $H_k(A \cap B) \to H_k(A \cup B)$ also maps h to zero. The composition $H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ maps h to $0 \oplus h_b$ to zero for some $h_b \in H_k(B)$. Since $H_k(B) \to H_k(A \cup B)$ is an isomorphism, element h_b equals zero. Thus the mapping $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$ sends h to $0 \oplus 0$, the zero element of $H_k(A) \oplus H_k(B)$. Since $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$ is injective, element h must be zero, a contradiction. Thus, $H_k(A \cap B) \to H_k(A)$ is injective.

Assume that $H_k(A \cap B) \to H_k(A)$ is not surjective. By definition, there exists an $h_a \in H_k(A)$ which is not in the image of $H_k(A \cap B) \to H_k(A)$. Thus, $h_a \oplus 0$ is not in the image of $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$. The mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ takes $h_a \oplus 0$ to some $h \in H_k(A \cup B)$. Since $h_a \oplus 0$ is not in the image of $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$ and

$$H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B)$$

is exact, element h is not equal to zero. Since $H_k(B) \to H_k(A \cup B)$ is an isomorphism, there exists an $h_b \in H_k(B)$ such that $H_k(B) \to H_k(A \cup B)$ maps h_b to h. The mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ takes $h_a \oplus h_b$ to $h - h = 0$. Thus, $h_a \oplus h_b$ is in the image of $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$ and h_a is in the image of $H_k(A \cap B) \to H_k(A)$, a contradiction. We conclude that $H_k(A \cap B) \to H_k(A)$ is surjective. $\square$

Lemma 13. *Let A and B be topological spaces such that*

$$\begin{array}{ccccc} H_k(A \cap B) & \longrightarrow & H_k(A) \oplus H_k(B) & \longrightarrow & H_k(A \cup B) \\ & & & & \downarrow \\ & & & & H_{k-1}(A \cap B) \end{array}$$

is exact for all $k \geq 0$. Mapping $H_k(A \cap B) \to H_k(A)$ is an isomorphism for all k if and only if mapping $H_k(B) \to H_k(A \cup B)$ is an isomorphism for all k.

Proof. Part I: Assume that $H_k(A \cap B) \to H_k(A)$ is an isomorphism for all k. Show that $H_k(B) \to H_k(A \cup B)$ is an isomorphism for all k.

Since $H_k(A \cap B) \to H_k(A)$ is injective for all k, mapping $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$ is also injective for all k. Since

$$H_k(A \cup B) \to H_{k-1}(A \cap B) \to H_{k-1}(A) \oplus H_{k-1}(B)$$

is exact and $H_{k-1}(A \cap B) \to H_{k-1}(A) \oplus H_{k-1}(B)$ is injective, the image of $H_k(A \cup B) \to H_{k-1}(A \cap B)$ is zero. Since

$$H_k(A) \oplus H_k(B) \to H_k(A \cup B) \to H_{k-1}(A \cap B)$$

is exact and the image of $H_k(A \cup B) \to H_{k-1}(A \cap B)$ is zero, $H_k(A \cup B)$ is the image of $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$. Thus, $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ is surjective and

$$0 \to H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B) \to 0$$

forms a short exact sequence. By Lemma 12, if $H_k(A \cup B) \to H_k(A)$ is an isomorphism, then $H_k(B) \to H_k(A \cup B)$ is an isomorphism.

Part II: Assume that $H_k(B) \to H_k(A \cup B)$ is an isomorphism for all k. Show that $H_k(A \cap B) \to H_k(A)$ is an isomorphism for all k.

Since $H_k(B) \to H_k(A \cup B)$ is surjective for all k, the mapping $H_k(A) \oplus H_k(B) \to H_k(A \cup B)$ is surjective for all k. Since

$$H_{k+1}(A) \oplus H_{k+1}(B) \to H_{k+1}(A \cup B) \to H_k(A \cap B)$$

is exact and $H_{k+1}(A) \oplus H_{k+1}(B) \to H_{k+1}(A \cup B)$ is surjective, the image of $H_{k+1}(A \cup B) \to H_k(A \cap B)$ is zero. Since

$$H_{k+1}(A \cup B) \to H_k(A \cap B) \to H_k(A) \oplus H_k(B)$$

is exact and the image of $H_{k+1}(A \cup B) \to H_k(A \cap B)$ is zero, the mapping $H_k(A \cap B) \to H_k(A) \oplus H_k(B)$ is injective and

$$0 \to H_k(A \cap B) \to H_k(A) \oplus H_k(B) \to H_k(A \cup B) \to 0$$

forms a short exact sequence. By Lemma 12, if $H_k(B) \to H_k(A \cup B)$ is an isomorphism. then $H_k(A \cup B) \to H_k(A)$ is an isomorphism. □

As a corollary, we get the relationship between critical points for interval sets and critical points for sub-level sets given by Lemma 1.

Proof of Lemma 1: Let A equal $F_{-\infty}^{f(p)}$ and B equal $F_a^{f(p)} \cup \{p\}$. By Mayer-Vietoris,

$$\begin{array}{ccccc} H_k(F_a^{f(p)}) & \to & H_k(A) \oplus H_k(B) & \to & H_k(F_{-\infty}^{f(p)} \cup \{p\}) \\ & & & & \downarrow \\ & & & & H_{k-1}(F_a^{f(p)}) \end{array}$$

is an exact sequence for all k. By Lemma 13, $H_k(F_a^{f(p)}) \to H_k(F_a^{f(p)} \cup \{p\})$ is an isomorphism for all integers $k \geq 0$ if and only if $H_k(F_{-\infty}^{f(p)}) \to H_k(F_{-\infty}^{f(p)} \cup \{p\})$ is an isomorphism for all integers $k \geq 0$. □

Finally, we prove that if point q destroys some element of H_k, then $f(q)$ is an H_k-critical value.

Lemma 14. *If point q destroys some non-zero $h \in H_k(f^{-1}(a))$, then $H_k(F_{a'}^{f(q)}) \to H_k(F_{a'}^{f(q)} \cup \{q\})$ is not an isomorphism for any a' where $a < a' < f(q)$.*

Proof. Let A equal $\underline{F}_a^{f(q)}$ and B equal $\underline{F}_{a'}^{f(q)} \cup \{q\}$. Note that $A \cup B$ equals $\underline{F}_a^{f(q)} \cup \{q\}$ and $A \cap B$ equals $\underline{F}_{a'}^{f(q)}$.

Let $h_a \in H_k(A)$ be the image of h under the mapping $H_k(f^{-1}(a)) \to H_k(A)$. Since q destroys h, element h_a is non-zero and the mapping $H_k(A) \to H_k(A \cup B)$ sends h_a to zero. By Corollary 8, there is some $h' \in H_k(A \cap B)$ such that $H_k(A \cap B) \to H_k(A) \to H_k(A \cup B)$ sends h' to h_a to zero. Since h_a is non-zero, element h' is non-zero. Thus the mapping $H_k(A \cap B) \to H_k(A \cup B)$ is not an isomorphism. □

C Appendix: Interval vs. topological persistence

In this section, we prove Theorem 2 showing that interval persistence generalizes topological persistence for Morse functions. Morse functions have the property that for all $a \in \mathbb{R}$, there exists some $\epsilon_0 > 0$ such that for all $\epsilon \leq \epsilon_0$, the set $f^{-1}(a)$ is a strong deformation retract of $N_\epsilon(f^{-1}(a)) = \{x : a - \epsilon < f(x) < a + \epsilon\}$. Note that if $f^{-1}(a)$ is a strong deformation retract of $N_\epsilon(f^{-1}(a))$, then $H_k(f^{-1}(a)) \to H_k(N_\epsilon(f^{-1}(a)))$ is an isomorphism for every integer $k \geq 0$.

As in Section 2.3, let δ be the minimum distance between any two critical values of f and let ϵ equal $\delta/2$. As defined in Section 2.1, $\beta_k(X)$ represent the k'th Betti number of X, i.e., the dimension of $H_k(X)$.

Since we consider homology groups defined only over fields, the homology groups have no torsion. Thus, if p is the only critical point of f with critical value $b = f(p)$, then, for some integer k, $\beta_k(F_{-\infty}^{b+\epsilon})$ equals $\beta_k(F_{-\infty}^{b-\epsilon}) \pm 1$ and $\beta_j(F_{-\infty}^{b+\epsilon})$ equals $\beta_j(F_{-\infty}^{b-\epsilon})$ for all $j \neq k$. The proof (omitted) follows from the theorem [10, Theorem 3.1] that there exists a k-cell, ω, in Ω, such that $\omega \cap F_{-\infty}^{b-\epsilon} = \partial\omega$, and there is a deformation retract from $F_{-\infty}^{b+\epsilon}$ onto $F_{-\infty}^{b-\epsilon} \cup \omega$.

Since the homology groups have no torsion, if $\beta_j(F_{-\infty}^{b-\epsilon})$ equals $\beta_j(F_{-\infty}^{b+\epsilon})$ then $H_j(F_{-\infty}^{b-\epsilon}) \to H_j(F_{-\infty}^{b+\epsilon})$ is an isomorphism. If $\beta_j(F_{-\infty}^{b-\epsilon})$ equals $\beta_j(F_{-\infty}^{b+\epsilon})+1$ and only a single non-degenerate critical point has critical value b, then $H_j(F_{-\infty}^{b-\epsilon}) \to H_j(F_{-\infty}^{b+\epsilon})$ is a surjection. If $\beta_j(F_{-\infty}^{b-\epsilon})$ equals $\beta_j(F_{-\infty}^{b+\epsilon}) - 1$ and only a single non-degenerate critical point has critical value b, then $H_j(F_{-\infty}^{b-\epsilon}) \to H_j(F_{-\infty}^{b+\epsilon})$ is an injection.

In Section 2.6, we defined μ_a^b as

$$(\beta_{a+\epsilon}^{b-\epsilon} - \beta_{a+\epsilon}^{b+\epsilon}) - (\beta_{a-\epsilon}^{b-\epsilon} - \beta_{a-\epsilon}^{b+\epsilon}).$$

Two critical points p and q with $f(p) < f(q)$ are paired if and only if $\mu_{f(p)}^{f(q)}$ is positive. We call (p, q) a pair of critical points and $(f(p), f(q))$ a pair of critical values.

For the proof of Theorem 2, we need a lemma which shows that if (a, b) is a pair of critical values and no two critical points share the same critical values, then μ_a^b equals one and $\mu_{a'}^b$ equals zero for all critical values $a' \neq a$. Fix the homology group index k. For any $x, y \in \mathbb{R}$ where $x < y$, let W_x^y be the image of $H_k(F_{-\infty}^x)$ under the mapping $H_k(F_{-\infty}^x) \to H_k(F_{-\infty}^y)$, i.e., W_x^y is a persistent homology group. The persistent Betti number β_x^y is the dimension of W_x^y.

Lemma 15. *Let f be a Morse function where no two critical points share the same critical value. If (a, b) are a pair of critical values, then*

1. *μ_a^b equals one;*
2. *$\mu_{a'}^b$ equals zero for all critical values $a' \neq a$;*
3. *$\beta_{a''}^{b-\epsilon} - \beta_{a''}^{b+\epsilon}$ equals zero for all $a'' < a$.*

Proof. Let x, x', y, y' be reals such that $x' \leq x < y \leq y'$. We prove that $\beta_x^y \geq \beta_{x'}^{y'}$. Intuitively, any homology group which persists from x' to y' surely persists from x to y. More formally, we define an injective mapping Φ which takes each element of $W_{x'}^{y'}$ to a distinct element of W_x^b.

Define the mapping $\Phi : W_{x'}^{y'} \to W_x^y$ as follows. Every $h_y' \in W_{x'}^{y'} \subseteq H_k(F_{-\infty}^{y'})$, is the image of some $h_x' \in H_k(F_{-\infty}^{x'})$ under the mapping $H_k(F_{-\infty}^{x'}) \to H_k(F_{-\infty}^{y'})$. Let h_x be the image of h_x' under the mapping $H_k(F_{-\infty}^{x'}) \to H_k(F_{-\infty}^x)$ and let h_y the image of h_x under the mapping $H_k(F_{-\infty}^x) \to H_k(F_{-\infty}^y)$. The mappings are

$$\begin{array}{cccc} h_x' \in & h_x \in & h_y \in & h_y' \in \\ H_k(F_{-\infty}^{x'}) \longrightarrow & H_k(F_{-\infty}^x) \longrightarrow & H_k(F_{-\infty}^y) \longrightarrow & H_k(F_{-\infty}^{y'}). \end{array}$$

Define $\Phi(h'_y) = h_y$. Note that h'_y is the image of h_y under the mapping $H_k(F_{-\infty}^{y'}) \to H_k(F_{-\infty}^{y})$, so if h_y is zero, then h'_y is zero. Thus Φ is injective and β_x^y is greater than or equal to $\beta_{x'}^{y'}$. In particular, $\beta_{a+\epsilon}^{b-\epsilon} \geq \beta_{a+\epsilon}^{b+\epsilon}$ and $\beta_{a-\epsilon}^{b-\epsilon} \geq \beta_{a-\epsilon}^{b+\epsilon}$.

Next we show that $\beta_{x'}^{y} - \beta_{x'}^{y'} \leq \beta_x^y - \beta_x^{y'}$. The difference $\beta_{x'}^{y} - \beta_{x'}^{y'}$ represents the rank of the the homology groups which exist at x' and are destroyed between x and y'. Similarly, $\beta_x^y - \beta_x^{y'}$ represents the rank of the the homology groups which exist at x and are destroyed between y and y'. Since y is greater than x, any group existing at x' and destroyed between y and y' also exists at x. More formally, we define surjective mappings $\Psi : W_x^y \to W_x^{y'}$ and $\Psi' : W_{x'}^{y} \to W_{x'}^{y'}$ whose kernels represent the specified homology groups and show that the kernel of Ψ' is a subgroup of the kernel of Ψ.

As before, $x \leq y \leq y'$. Define the mapping $\Psi : W_x^y \to W_x^{y'}$ as follows. Let h_y be an element of $W_x^y \subseteq H_k(F_{-\infty}^{y})$. Since h_y is in W_x^y, element h_y is the image of some $h_x \in H_k(F_{-\infty}^{x})$ under the mapping $H_k(F_{-\infty}^{x}) \to H_k(F_{-\infty}^{y})$. Let h'_y be the image of h_y under the mapping $H_k(F_{-\infty}^{y}) \to H_k(F_{-\infty}^{y'})$. The mappings are

$$\begin{array}{ccccccc} h'_x \in & & h_x \in & & h_y \in & & h'_y \in \\ H_k(F_{-\infty}^{x'}) & \longrightarrow & H_k(F_{-\infty}^{x}) & \longrightarrow & H_k(F_{-\infty}^{y}) & \longrightarrow & H_k(F_{-\infty}^{y'}). \end{array}$$

The composition mapping $H_k(F_{-\infty}^{x}) \to H_k(F_{-\infty}^{y}) \to H_k(F_{-\infty}^{y'})$ takes h_x to h_y to h'_y. Thus h'_y is in $W_x^{y'}$. Define $\Psi(h_y) = h'_y$.

We claim that Ψ is surjective. Let $\hat{h}'_y$ be an element of $W_x^{y'}$. Since $\hat{h}'_y$ is in $W_x^{y'}$, element $\hat{h}'_y$ is the image of some $\hat{h}_x \in H_k(F_{-\infty}^{x})$ under the mapping $H_k(F_{-\infty}^{x}) \to H_k(F_{-\infty}^{y})$. Let $\hat{h}_y$ be the image of $\hat{h}_x$ under the mapping $H_k(F_{-\infty}^{x}) \to H_k(F_{-\infty}^{y})$. Element $\hat{h}_y$ is in W_x^y and $\Psi(\hat{h}_y)$ equals $\hat{h}'_y$ so Ψ is surjective.

The difference $\beta_x^y - \beta_x^{y'}$ is the difference between the rank of W_x^y and $W_x^{y'}$. Since $\Psi : W_x^y \to W_x^{y'}$ is surjective, this difference is simply the rank of the kernel of Ψ.

Define the mapping $\Psi' : W_{x'}^{y} \to W_{x'}^{y'}$ in the same manner as Ψ. Note that W_x^y and $W_{x'}^{y}$ are subgroups of $H_k(F_{-\infty}^{y})$. We claim that kernel of Ψ' is a subgroup of the kernel of Ψ. Let $h_y \in W_{x'}^{y} \subseteq H_k(F_{-\infty}^{y})$ be an element of the kernel of Ψ'. Since h_y is an element of $W_{x'}^{y}$, element h_y is the image of some $h'_x \in H_k(F_{-\infty}^{x'})$ under the mapping $H_k(F_{-\infty}^{x'}) \to H_k(F_{-\infty}^{y})$. Let h_x be the image of h'_x under the mapping $H_k(F_{-\infty}^{x'}) \to H_k(F_{-\infty}^{x})$. Since the mapping $H_k(F_{-\infty}^{x}) \to H_k(F_{-\infty}^{y}) \to H_k(F_{-\infty}^{y'})$ maps h_x to h_y to zero, element h_y is in the kernel of Ψ and the kernel of Ψ' is a subgroup of the kernel of Ψ. Since the ranks of the kernels of Ψ and Ψ' are $\beta_x^y - \beta_x^{y'}$ and $\beta_{x'}^{y} - \beta_{x'}^{y'}$, respectively, it follows that $\beta_x^y - \beta_x^{y'} \geq \beta_{x'}^{y} - \beta_{x'}^{y'}$.

Since no critical points share a critical value, $\beta_{b-\epsilon}^{b-\epsilon} - \beta_{b+\epsilon}^{b+\epsilon} \leq 1$ and the mapping $H_k(F_{-\infty}^{b-\epsilon}) \to H_k(F_{-\infty}^{b+\epsilon})$ is either an injection or a surjection. If $H_k(F_{-\infty}^{b-\epsilon}) \to H_k(F_{-\infty}^{b+\epsilon})$ is an injection, then $\beta_{b-\epsilon}^{b+\epsilon}$ equals $\beta_{b-\epsilon}^{b-\epsilon}$ and so $\beta_{b-\epsilon}^{b-\epsilon} - \beta_{b-\epsilon}^{b+\epsilon}$ equals zero. If $H_k(F_{-\infty}^{b-\epsilon}) \to H_k(F_{-\infty}^{b+\epsilon})$ is a surjection, then $\beta_{b-\epsilon}^{b+\epsilon}$ equals $\beta_{b+\epsilon}^{b+\epsilon}$ and so $\beta_{b-\epsilon}^{b-\epsilon} - \beta_{b-\epsilon}^{b+\epsilon} = \beta_{b-\epsilon}^{b-\epsilon} - \beta_{b+\epsilon}^{b+\epsilon} \leq 1$. In both cases, $\beta_{b-\epsilon}^{b-\epsilon} - \beta_{b-\epsilon}^{b+\epsilon}$ is at most one. Thus for all $x < b - \epsilon$,

$$\beta_x^{b-\epsilon} - \beta_x^{b+\epsilon} \leq \beta_{b-\epsilon}^{b-\epsilon} - \beta_{b-\epsilon}^{b+\epsilon} \leq 1.$$

In particular, $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon} \leq 1$ and $\beta^{b-\epsilon}_{a+\epsilon} - \beta^{b+\epsilon}_{a+\epsilon} \leq 1$.

By definition,

$$\mu^b_a = (\beta^{b-\epsilon}_{a+\epsilon} - \beta^{b+\epsilon}_{a+\epsilon}) - (\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}).$$

Since $\beta^{b-\epsilon}_{a-\epsilon} \geq \beta^{b+\epsilon}_{a-\epsilon}$ and $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon} \leq 1$, the value $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}$ is either zero or one. Similarly, the value $\beta^{b-\epsilon}_{a+\epsilon} - \beta^{b+\epsilon}_{a+\epsilon}$ is either zero or one.

Since (a, b) is a pair of critical values, μ^b_a is positive. The only way μ^b_a can be positive is if $\beta^{b-\epsilon}_{a+\epsilon} - \beta^{b+\epsilon}_{a+\epsilon}$ equals one and $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}$ equals zero. Substituting one and zero in the formula for μ^b_a shows that μ^b_a equals one.

Since $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}$ equals zero, $\beta^{b-\epsilon}_{a'+\epsilon} - \beta^{b+\epsilon}_{a'+\epsilon} \leq \beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon} = 0$ must also be zero for all $a' + \epsilon \leq a - \epsilon$. Thus, $\mu^b_{a'}$ must be zero for all $a' < a$. On the other hand, if $\mu^b_{a'}$ equaled one for some $a' > a$, then this argument would show that μ^b_a must equal zero. Thus for each b there is at most one non-zero μ^b_a.

Similarly, if $a'' \leq a - \epsilon$, then $\beta^{b-\epsilon}_{a''} - \beta^{b+\epsilon}_{a''} \leq \beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}$. Since $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}$ equals zero, $\beta^{b-\epsilon}_{a''} - \beta^{b+\epsilon}_{a''}$ equals zero. If $a - \epsilon < a'' < a$, then $\beta^{b-\epsilon}_{a''}$ equals $\beta^{b-\epsilon}_{a-\epsilon}$ and $\beta^{b+\epsilon}_{a''}$ equals $\beta^{b+\epsilon}_{a-\epsilon}$, so again $\beta^{b-\epsilon}_{a''} - \beta^{b+\epsilon}_{a''}$ equals zero. □

Proof of Theorem 2. Let (p, q) where $f(p) < f(q)$ be a pair of critical points with topological persistence $f(q) - f(p)$. By Lemma 15, $\mu^{f(q)}_{f(p)}$ equals one. Let k be the index of the homology groups H_k used in determining the μ^b_a. We wish to show that point q has interval persistence $f(q) - f(p)$.

Part I: Show that the interval persistence of q is at least $f(q) - f(p)$.

Let a equal $f(p)$ and b equal $f(q)$. By Lemma 15, μ^b_a equals one. Since the persistent Betti numbers are non-negative, $\beta^{b-\epsilon}_{a+\epsilon} - \beta^{b+\epsilon}_{a+\epsilon}$ must equal one and $\beta^{b-\epsilon}_{a-\epsilon} - \beta^{b+\epsilon}_{a-\epsilon}$ must equal zero. Since $(\beta^{b-\epsilon}_{a+\epsilon} - \beta^{b+\epsilon}_{a+\epsilon})$ equals one, there is some $h \in H_k(F^{a+\epsilon}_{-\infty})$ which maps to zero under $H_k(F^{a+\epsilon}_{-\infty}) \to H_k(F^{b+\epsilon}_{-\infty})$ but not under $H_k(F^{a+\epsilon}_{-\infty}) \to H_k(F^{b-\epsilon}_{-\infty})$. Since the only critical value between $b - \epsilon$ and $b + \epsilon$ is at b, mapping $H_k(F^{a+\epsilon}_{-\infty}) \to H_k(F^b_{-\infty} \cup f^{-1}(b))$ sends h to zero. Since q is the only critical point with critical value b, mapping $H_k(F^{a+\epsilon}_{-\infty}) \to H_k(F^b_{-\infty} \cup \{q\})$ sends h to zero.

Let a' be any real between a and $a + \epsilon$, i.e., $a < a' < a + \epsilon$. We claim that there is some $\hat{h} \in H_k(f^{-1}(a'))$ whose image is h under the mapping $H_k(f^{-1}(a')) \to H_k(F^{a+\epsilon}_{-\infty}))$ and which is destroyed at point q. Since f is a Morse function, there is some ϵ' such that $a < a' - \epsilon' < a' < a' + \epsilon' < a + \epsilon$ and $f^{-1}(a')$ is a strong deformation retract of $N_{\epsilon'}(f^{-1}(a'))$. Thus, $H_k(f^{-1}(a')) \to H_k(N_{\epsilon'}(f^{-1}(a')))$ is an isomorphism.

Since there are no critical values between a and $a+\epsilon$, $H_k(F^{a'+\epsilon'}_{-\infty}) \to H_k(F^{a+\epsilon}_{-\infty})$ is an isomorphism. Let h' be the element of $H_k(F^{a'+\epsilon'}_{-\infty})$ whose image is $h \in H_k(F^{a+\epsilon}_{-\infty})$ under this mapping. The mapping $H_k(F^{a'+\epsilon'}_{-\infty}) \to H_k(F^{a'+\epsilon'}_{-\infty} \cup f^{-1}(b))$ takes h' to zero. The intersection of $H_k(F^{a'+\epsilon'}_{-\infty})$ and $H_k(F^b_{a'-\epsilon'} \cup f^{-1}(b)$ is $N_{\epsilon'}(f^{-1}(a'))$ and their union is $F^b_{-\infty} \cup f^{-1}(b)$. By Corollary 8, h' is the image of some $h'' \in H_k(N_{\epsilon'}(f^{-1}(a')))$ under the mapping $H_k(N_{\epsilon'}(f^{-1}(a'))) \to H_k(F^{a'+\epsilon'}_{-\infty})$. Since the mapping $H_k(f^{-1}(a')) \to H_k(N_{\epsilon'}(f^{-1}(a')))$ is an isomorphism, h'' is the image of some $\hat{h}$ under this mapping.

The mapping

$$\begin{array}{ccc} H_k(f^{-1}(a')) \longrightarrow & H_k(N_{\epsilon'}(f^{-1}(a'))) & \\ & \downarrow & \\ & H_k(F^{a'+\epsilon'}_{-\infty}) & \longrightarrow H_k(F^{a+\epsilon}_{-\infty}) \end{array}$$

takes $\hat{h}$ to h'' to h' to h. Since h is destroyed by q, so is $\hat{h}$. Thus, for any a' between a and $a+\epsilon$, some element of $H_k(f^{-1}(a'))$ is destroyed by q. Since the interval persistence of q is

$$\sup\{f(q) - a' : q \text{ destroys some non-zero } h \in H_k(f^{-1}(a'))\}.$$

and $\sup\{f(q) - a' : a < a' < a + \epsilon\}$ equals $f(q) - a$, the interval persistence is at least $f(q) - a = f(q) - f(p)$.

Part II: Show that the interval persistence of q is at most $f(q) - f(p)$.

Again let a equal $f(p)$ and let b equal $f(q)$. If the interval persistence of q is more than $f(q) - f(p)$ then for some $a' < a$, some non-zero $h \in H_k(f^{-1}(a'))$ is destroyed by q. Let a'' be a real between a' and a, i.e., $a' < a'' < a$. Let h'_a be the image of h under the mapping $H_k(f^{-1}(a')) \to H_k(\underline{F}^{a''}_{a'})$. The mappings

$$H_k(f^{-1}(a')) \to H_k(\underline{F}^{a''}_{a'}) \to H_k(\underline{F}^{b}_{a'}) \to H_k(\underline{F}^{b}_{a'} \cup \{q\})$$

take h to h'_a to some $h' \in H_k(\underline{F}^{b}_{a'})$ to $0 \in H_k(\underline{F}^{b}_{a'} \cup \{q\})$. Since h is destroyed by q, element h' is non-zero. Thus h'_a is non-zero and h'_a is also destroyed at q.

Consider the image h''_a of h'_a under the mapping $H_k(F^{a''}_{a'}) \to H_k(F^{a''}_{-\infty})$. If h''_a is non-zero, then h''_a is destroyed at q. Thus, $\beta^{b-\epsilon}_{a''} - \beta^{b+\epsilon}_{a''}$ is positive. However, since $a'' < a$, $\beta^{b-\epsilon}_{a''} - \beta^{b+\epsilon}_{a''}$ equals zero by Lemma 15, a contradiction.

Now assume that h''_a is zero. Since μ^b_a equals one, $H_k(F^{b-\epsilon}_{-\infty}) \to H_k(F^{b+\epsilon}_{-\infty})$ is not an isomorphism. Since f is a Morse function, $H_j(F^{b-\epsilon}_{-\infty}) \to H_j(F^{b+\epsilon}_{-\infty})$ is an isomorphism for all $j \neq k$.

Since the image of h'_a under the mapping $H_k(F^{a''}_{a'}) \to H_k(F^{a''}_{-\infty})$ is $h''_a = 0$ and the image of h'_a under the mapping $H_k(F^{a''}_{a'}) \to H_k(F^{b+\epsilon}_{a'})$ is zero, the image of h'_a under the mapping $H_k(F^{a''}_{a'}) \to H_k(F^{a''}_{-\infty} \oplus H_k(F^{b+\epsilon}_{a'})$ is $0 \oplus 0$. Since the mapping

$$H_{k+1}(F^{b+\epsilon}_{-\infty}) \to H_k(F^{a''}_{a'}) \to H_k(F^{a''}_{-\infty}) \oplus H_k(F^{b+\epsilon}_{a'})$$

is exact by Mayer-Vietoris, element h'_a must be in the image of some $h_b \in H_{k+1}(F^{b+\epsilon}_{-\infty})$ under the mapping $H_{k+1}(F^{b+\epsilon}_{-\infty}) \to H_k(F^{a''}_{a'})$.

On the other hand, $H_{k+1}(F^{b-\epsilon}_{-\infty}) \to H_{k+1}(F^{b+\epsilon}_{-\infty})$ is an isomorphism. Thus h_b is the image of some $h'_b \in H_{k+1}(F^{b-\epsilon}_{-\infty})$. The composed mappings $H_{k+1}(F^{b-\epsilon}_{-\infty}) \to H_{k+1}(F^{b+\epsilon}_{-\infty}) \to H_k(F^{a''}_{a'})$ take h'_b to h_b to non-zero h'_a. Since h'_a is destroyed by q, the image of h'_a under mapping $H_k(F^{a''}_{a'}) \to H_k(F^{b-\epsilon}_{a'})$ is non-zero. Thus the image of h'_a under the mapping $H_k(F_{a'}^{a''}) \to H_k(F_{a'}^{b-\epsilon}) \oplus H_k(F_{-\infty}^{a''})$ is non-zero. By Mayer-Vietoris, the mapping

$$H_{k+1}(F_{-\infty}^{b-\epsilon}) \to H_k(F_{a'}^{a''}) \to H_k(F_{-\infty}^{a''}) \oplus H_k(F_{a'}^{b-\epsilon})$$

is exact. Since h'_a is in the image of h'_b under the mapping $H_{k+1}(F_{-\infty}^{b-\epsilon}) \to H_k(F_{a'}^{a''})$, the image of h'_a under the mapping $H_k(F_{a'}^{a''}) \to H_k(F_{a'}^{b-\epsilon}) \oplus H_k(F_{-\infty}^{a''})$ is zero, a contradiction. Thus h''_a cannot be zero and the interval persistence of q is at most $f(q) - f(p)$. □

D Appendix: Properties of LR maps

Lemma 16 says that if f is LR, then we can replace $\{x : a \leq f(x) \leq b\}$ by suitably chosen small neighborhoods without changing its homology.

Lemma 16. *If continuous function $f : \Omega \to \mathbb{R}$ is LR, then for every $a, b \in \mathbb{R}$ where $a < b$, there exists an ϵ_0 such that for all $\epsilon < \epsilon_0$, set $\{x : a \leq f(x) \leq b\}$ is a strong deformation retract of $\{x : a \leq f(x) < b + \epsilon\}$ and set $\{x : a \leq f(x) < b\}$ is a strong deformation retract of $\{x : a - \epsilon < f(x) < b\}$.*

Proof. Assume function $f : \Omega \to \mathbb{R}$ is LR. Choose $\epsilon_1 > 0$ such that $f^{-1}(a)$ is a strong deformation retract of $\{x : a - \epsilon < f(x) < a + \epsilon\}$ for all $\epsilon < \epsilon_1$. Similarly, choose $\epsilon_2 > 0$ such that $f^{-1}(b)$ is a strong deformation retract of $\{x : b - \epsilon < f(x) < b + \epsilon\}$ for all $\epsilon < \epsilon_2$. Let ϵ_0 be the minimum of ϵ_1 and ϵ_2 and $(b - a)/2$.

For any $\epsilon \leq \epsilon_0$, let ϕ_ϵ^a be the mapping from $\{a : a - \epsilon < f(x) < a + \epsilon\} \times I$ to $f^{-1}(a)$ representing the strong deformation retract of $\{a : a - \epsilon < x < a + \epsilon\}$ to $f^{-1}(a)$. Let ϕ_ϵ^b be the mapping from $\{b : b - \epsilon < f(x) < b + \epsilon\} \times I$ to $f^{-1}(b)$ representing the strong deformation retract of $\{b : b - \epsilon < x < b + \epsilon\}$ to $f^{-1}(b)$. Define

$$\phi_\epsilon(x, t) = \begin{cases} \phi_\epsilon^a(x, t) & \text{for } a - \epsilon < x < a, \\ x & \text{for } a \leq x \leq b, \\ \phi_\epsilon^b(x, t) & \text{for } b < x < b + \epsilon. \end{cases}$$

ϕ_ϵ is constant on $\{x : a \leq x \leq b\}$ and continuously deforms $\{x : a - \epsilon < x < a\}$ and $\{x : b < x < b + \epsilon\}$ onto $\{x : a \leq x \leq b\}$. Thus $\{x : a \leq f(x) \leq b\}$ is a strong deformation retract of $\{x : a \leq f(x) < b + \epsilon\}$ and $\{x : a \leq f(x) < b\}$ is a strong deformation retract of $\{x : a - \epsilon < f(x) < b\}$. □

Let non-zero $h \in H_k(f^{-1}(a))$ be destroyed by $\underline{F}_a^\infty$. If f is LR, then h is destroyed by $\{x : a \leq f(x) \leq b\}$ for some $b \geq a$. Equivalently, the image of h in $H_k(\underline{F}_a^b)$ is destroyed by $f^{-1}(b)$.

Lemma 17. *Let $f : \Omega \to \mathbb{R}$ be a continuous, LR function. For any non-zero $h \in H_k(f^{-1}(a))$, if $\underline{F}_a^\infty$ destroys h, then for some $b \geq a$, the image of h under the mapping $H_k(f^{-1}(a)) \to H_k(\underline{F}_a^b)$ is destroyed by $f^{-1}(b)$.*

Proof. Since $H_k(f^{-1}(a)) \to H_k(\underline{F}_a^\infty)$ sends h to zero, element h is the boundary of some chain $C \subseteq \underline{F}_a^\infty$. Chain C is compact. (See [12, p. 71].) Thus $\{f(x) : x \in C\}$ is compact and has a maximum value b'. Since $C \subseteq \underline{F}_a^{b'}$, the mapping $H_k(f^{-1}(a)) \to H_k(\underline{F}_a^{b'})$ sends h to zero.

Let b equal $\inf\{\tilde{b} : H_k(f^{-1}(a)) \to H_k(\underline{F}_a^{\tilde{b}})$ sends h to zero$\}$. Note that $b \leq b'$. Since f is LR, $H_k(f^{-1}(a))$ is isomorphic to $H_k(F_{a-\epsilon}^{a+\epsilon})$ for sufficiently small ϵ and thus $H_k(f^{-1}(a)) \to H_k(\underline{F}_a^{a+\epsilon})$ does not send h to zero. Thus b is strictly greater than a.

Let h' be the image of h under the mapping $H_k(f^{-1}(a)) \to H_k(\underline{F}_a^b)$. If h' were zero, then h would be the boundary of some chain $C' \subseteq \underline{F}_a^b$. Since C' is compact, chain C' would also be a subset of $\underline{F}_a^{\tilde{b}}$ for some $\tilde{b} < b$, contradicting the choice of b. Thus h' is non-zero.

We show that h' is destroyed by $f^{-1}(b)$. Since f is LR, there is some $\epsilon_0 > 0$ such that $H_k(\underline{F}_a^b \cup f^{-1}(b))$ is isomorphic to $H_k(\underline{F}_a^{b+\epsilon})$ for all $\epsilon \leq \epsilon_0$. If $H_k(\underline{F}_a^b) \to H_k(\underline{F}_a^b \cup f^{-1}(b))$ does not map h' to zero, then $H_k(\underline{F}_a^b) \to H_k(\underline{F}_a^{b+\epsilon})$ does not map h' to zero for all $\epsilon \leq \epsilon_0$, and b does not equal $\inf\{\tilde{b} : H_k(f^{-1}(b)) \to H_k(\underline{F}_a^{\tilde{b}})$ sends h to zero$\}$. Thus, $H_k(\underline{F}_a^b) \to H_k(\underline{F}_a^b \cup f^{-1}(b))$ maps h' to zero and h' is destroyed by $f^{-1}(b)$. □

The proof of Lemma 2 follows from Lemma 17.

Proof of Lemma 2. By Lemma 17, there exists some $b > a$ such that the image h' of h under the mapping $H_k(f^{-1}(a)) \to H_k(\underline{F}_a^b)$ is destroyed by $f^{-1}(b)$. Since f is point destructible, there is some point $q \in f^{-1}(b)$ which destroys h' and h. □

Proof of Lemma 4. Let b equal $f(q)$. The intersection of $\underline{F}_{a_1}^b$ and $X \cup \{q\}$ is $X = \{x \in \Omega : a_2 \leq g(x)$ and $f(x) < f(q)\}$. and their union is $\underline{F}_{a_1}^b \cup \{q\}$. By Corollary 8, h is the image of some $h_x \in H_k(X)$ under the mapping $H_k(X) \to H_k(\underline{F}_{a_1}^b)$ and h is destroyed by q.

We now prove that h is the image of some $h_2 \in H_k(g^{-1}(a_2))$. Since g is LR, there exists some $\epsilon_1 > 0$ such that $H_k(g^{-1}(a_2)) \to H_k(N_{\epsilon'}(g^{-1}(a_2)))$ is an isomorphism for all $\epsilon' \leq \epsilon_1$. By Lemma 16, there exists some $\epsilon_2 > 0$ such that $H_k(X) \to H_k(X \cup N_{\epsilon'}(g^{-1}(a_2)))$ is an isomorphism for all $\epsilon' \leq \epsilon_2$. Since $g^{-1}(a_2) \subseteq F_{a_1}^b$, there is some ϵ_3 such that $N_{\epsilon_3}(g^{-1}(a_2)) \subseteq F_{a_1}^b$. Let ϵ be the smaller of ϵ_1, ϵ_2 and ϵ_3.

Let Y equal $N_\epsilon(g^{-1}(a_2))$. Let $Z = \{x \in \Omega : a_2 - \epsilon < g(x)$ and $f(x) < f(q)\}$ and $Z' = \{x \in \Omega : a_1 \leq f(x)$ and $g(x) < a_2 + \epsilon\}$ The following commutative diagram gives the relevant mappings between homology groups:

$$
\begin{array}{ccccc}
 & & & & h_1 \in H_k(f^{-1}(a_1)) \\
 & & & & \downarrow \\
h_2 \in H_k(g^{-1}(a_2)) & = & h_y \in H_k(Y) & \longrightarrow & H_k(Z') \ni h'_z \\
\downarrow & & \downarrow & & \downarrow \\
h_x \in H_k(X) & = & h_z \in H_k(Z) & \longrightarrow & H_k(\underline{F}^b_{a_1}) \ni h
\end{array}
$$

Let h_z be the image of h_x under the mapping $H_k(X) \to H_k(Z)$. Let h'_z be the image of h_1 under the mapping $H_k(f^{-1}(a_1)) \to H_k(Z')$. Element h is the image of $h_x \in H_k(X)$ and $h_1 \in H_k(f^{-1}(a_1))$ and so is the image of both h_z and h'_z under the respective mappings $H_k(Z) \to H_k(X)$ and $H_k(Z') \to H_k(X)$. By Corollary 10, there is some $h_y \in H_k(Y)$ such that $H_k(Y) \to H_k(Z)$ maps h_y to h_z and $H_k(Y) \to H_k(Z')$ maps h_y to h'_z. Since $H_k(g^{-1}(a_2)) \to H_k(Y)$ is an isomorphism, there is some $h_2 \in H_k(g^{-1}(a_2))$ whose image is h_y under the mapping $H_k(g^{-1}(a_2)) \to H_k(Y)$. The mapping $H_k(g^{-1}(a_2)) \to H_k(Y) \to H_k(Z) \to H_k(F^b_{a_1})$ takes h_2 to h_y to h_z to h. □

Proof of Lemma 3. Set g equal to f and apply Lemma 4. □

Finally, we prove that piecewise linear functions on compact spaces are LR.

Lemma 18. *A piecewise linear function $f : \Omega \to \mathbb{R}$ on the compact space Ω is LR.*

Proof. By definition, Ω has a finite triangulation $\mathcal{T}$ such that f is linear on each simplex in $\mathcal{T}$. Fix $a \in \mathbb{R}$. Since $\mathcal{T}$ has a finite number of simplices, there is some ϵ such that $\{x : a - \epsilon \leq x < a\}$ and $\{x : a < x \leq a + \epsilon\}$ do not contain any vertices of $\mathcal{T}$. Let Ω' equal $\{x \in \Omega : a - \epsilon \leq x \leq a + \epsilon\}$. Let $\mathcal{T}'$ be the triangulation of Ω' induced by intersecting the simplices of $\mathcal{T}$ by $\{x \in \Omega : a - \epsilon \leq x \leq a\}$ and $\{x \in \Omega : a \leq x \leq a + \epsilon\}$ and subdividing the resulting convex polytopes without adding any new vertices. Note that the vertices of $\mathcal{T}'$ all lie on $f^{-1}(a - \epsilon)$ or $f^{-1}(a)$ or $f^{-1}(a + \epsilon)$.

Let V be the vertices of $\mathcal{T}'$. Each vertex $v \in V$ which does not lie on $f^{-1}(a)$ lies on a unique edge e_v of $\mathcal{T}$. Since sets $\{x : a - \epsilon < x < a\}$ and $\{x : a < x < a + \epsilon\}$ do not contain any vertices of $\mathcal{T}$, edge e_v intersects $f^{-1}(a)$. Since f is linear, e_v intersects $f^{-1}(a)$ at only one point. Let $\tilde{v}$ be the point $e_v \cap f^{-1}(a)$. For each vertex $v \in V$ which lies on $f^{-1}(a)$, let $\tilde{v}$ equal v.

Define the map $\phi : V \times [0, 1] \to N_\epsilon(a)$ as $\phi(v, s) = (1 - s)v + s\tilde{v}$. Note that this map is constant on all points in $f^{-1}(a)$.

Extend map ϕ linearly over each simplex in $\mathcal{T}'$ forming a map over $\Omega' \times [0, 1]$. More specifically, each point $x \in \Omega'$ lies in some simplex $t' \in \mathcal{T}'$ whose vertices are $V \cap t'$. Point x is a convex combination, $\sum_{v \in V \cap t'} \alpha_v v$, of the points $V \cap t'$. Define

$\phi(x,s) = \sum_{v \in V \cap t'} \alpha_v \phi(v,s)$. Since ϕ is constant on all vertices in $f^{-1}(a)$, map ϕ is also constant on all simplices in $f^{-1}(a)$ and thus constant on $f^{-1}(a)$. Thus, $\phi(x,0) = x$ and $\phi(x,1) \in f^{-1}(a)$ for all $x \in \Omega$ and $\phi(x,s) = x$ for all $x \in f^{-1}(a)$ and $f^{-1}(a)$ is a strong deformation retract of Ω'. Restricting ϕ to the open set $N_\epsilon(f^{-1}(a)) \subseteq \Omega'$, shows that $f^{-1}(a)$ is a strong deformation retract of $N_\epsilon(f^{-1}(a))$. Thus, f is locally retractible. □

E Appendix: Maxima

The following homology lemma is used in the proof of the pairwise disjointness of (γ, ∞)-neighborhoods of local maxima.

Lemma 19. *Let M be a connected, oriented d-manifold with non-empty boundary. Let $D_1, D_2, \ldots, D_k$ be the connected components of ∂M with orientation inherited from M. If $a_1 D_1 + a_2 D_2 + \ldots a_k D_k$ generate the zero element of $H_{d-1}(M)$, then $a_1 = a_2 = \ldots = a_k$.*

Proof. The sequence $H_d(M) \to H_d(M, \partial M) \to H_{d-1}(\partial M) \to H_{d-1}(M)$ is exact [9, Theorem 2.16, p. 117]. Since the boundary of M is not empty, the homology group $H_d(M)$ is zero. The homology group of $H_d(M, \partial M)$ is $\mathbb{G}$, the ground ring of the homology group. The map $H_d(M, \partial M) \to H_{d-1}(\partial M)$ is the connecting homomorphism. It maps $\mathbb{G}$ to $h \in H_{d-1}(\partial M)$ which is generated by $D_1 + D_2 + \ldots + D_k$. Since the mapping is exact, the image of $\mathbb{G}h$ under the mapping $H_{d-1}(\partial M) \to H_{d-1}(M)$ is zero. Moreover, only elements in $\mathbb{G}h$ map to zero. Thus, if $a_1 D_1 + a_2 D_2 + \ldots a_k D_k$ generate the zero element of $H_{d-1}(M)$, then $a_1 D_1 + a_2 D_2 + \ldots a_k D_k$ must generate an element of $\mathbb{G}h$ in $H_{d-1}(\partial M)$ and so $a_1 = a_2 = \ldots = a_k$. □