Stability of Topological Persistence for Domains *

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 level sets $\{x \in \Omega : f(x) = c\}$ 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 domain points 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 NMR 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 regions bounded by 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?

Edelsbrunner, Letscher and Zomorodian in [5] introduced the notion of *topological persistence*. As $c \in \mathbb{R}$ increases, topological features appear and disappear in the 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. in [5] use homology groups over $\mathbb{Z}/2\mathbb{Z}$ to define topological features. Carlsson and Zomorodian [9] showed how topological persistence could be computed for homology groups over any fields.

At the core of various application areas such as shape matching and visualization is the problem of characterizing and comparing scalar fields. Topological persistence gives one approach to comparing such 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 [3] 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 values are the range scalar values where the topological structure of the level sets changes. Cohen-Steiner et. al. 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"?

In this paper we revisit topological persistence and establish a stability result in terms of the critical points in

*Research partly supported by NSF CARGO grant DMS-0310642 and ARO grant DAAD19-02-1-0347.

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

the domain. There are two obstacles that we must overcome. First, if we look only at the topological structure of $\{x \in \Omega : f(x) \leq c\}$, then there is no such stability result.

Consider Figure 1. Let f be a function defined on the surface in $\mathbb{R}^3$ shown in the figure where $f(x)$ is the z-coordinate of the point x. The set $\{x \in \Omega : f(x) \leq f(r)\}$ is homeomorphic to a cylinder with circles c_1 and c_2 bounding each end. The first homology group (H_1) is generated by the circle c_1. (It is also generated by circle c_2.) Assume that the two maxima p and q have z-coordinates that are almost equal except that $f(p) < f(q)$. Since $f(p) < f(q)$, the first homology group becomes trivial in the set $\{x \in \Omega : f(x) \leq f(p)\}$. Loosely speaking, the cycle generated by c_1 at r gets destroyed at p. According to the framework of [5], we get a persistent value pair $[f(r), f(p)]$.

Now consider a slightly perturbed f denoted as $\hat{f}$. Set $\hat{f}$ equal to f everywhere except in the vicinity of p and q where $\hat{f}$ is perturbed so that $\hat{f}(q) < \hat{f}(p)$. Let p' and q' be new maxima close to p and q respectively for $\hat{f}$. The cycle generated by c_1 at r gets destroyed at q' in $\{x \in \Omega : \hat{f}(x) \leq \hat{f}(q')\}$. We get a persistent value pair $[f(r), \hat{f}(q')]$ for $\hat{f}$. Since $f(p)$ is close to $\hat{f}(q')$, the two persistent value pairs, $[f(r), f(p)]$ and $[\hat{f}(r), \hat{f}(q')]$, are close, confirming the Cohen-Steiner et al. result. However, the points p and q' are not close in any sense in the domain.

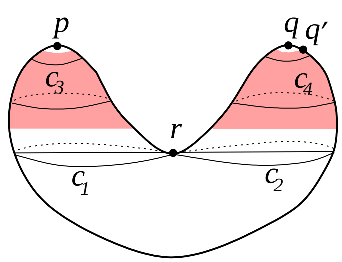

Figure 1: Set $\{x : f(x) \leq f(r)\}$ is homeomorphic to a cylinder with ends at circles c_1 and c_2. Circle c_1 generates the first homology group of $\{x : f(x) \leq f(r)\}$. Since $f(p) < f(q)$, this homology group is destroyed at p in $\{x : f(x) \leq f(p)\}$. If $\hat{f}$ is a slight perturbation of f where $\hat{f}(q') < \hat{f}(p)$, then this homology group is destroyed at q' in $\{x : \hat{f}(x) \leq \hat{f}(q')\}$. Points p and q' are far apart in the domain.

Instead of considering only sets $\{x \in \Omega : f(x) \leq c\}$ which are bounded from above by a single level set, we will consider *interval sets* that are bounded from above and below by level sets. The sets we use are $\{x \in \Omega : a \leq f(x) < b\}$. This is one crucial deviation we make from the set up in earlier works [3, 5, 9]. It also leads to slightly different definitions for critical points and persistence.

Returning to Figure 1, the first homology group of the set $\{x : f(r) \leq f(x) < f(p)\}$ (open at the top) has two distinct generators, one given by circle c_1 and one by circle c_2. Point p destroys the homology group generated by c_1 while point q destroys the one generated by c_2. In the perturbed field $\hat{f}$ where $\hat{f}(q') < \hat{f}(p')$, a similar thing happens. Thus p and p' (also q and q') are critical points for f and $\hat{f}$ respectively and destroy homology groups with approximately the same persistence.

The second problem in stability of critical points is that corresponding critical points can be arbitrarily far apart. Consider functions f and $\hat{f}$ in Figure 2 where $|f(x) - \hat{f}(x)| < \delta$ for all $x \in \Omega$. The maxima, p and q, of f and $\hat{f}$, respectively, can be made arbitrarily far apart even as δ is made arbitrarily small.

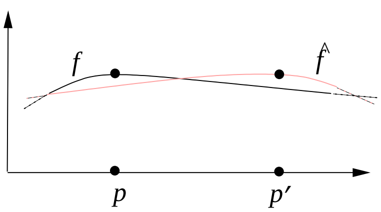

Figure 2: The maximum p for a real valued function f has moved by large distance even for an arbitrarily close approximant $\hat{f}$.

Instead of using a metric in the domain, we use the range to determine neighborhoods of points. A neighborhood of point p is the connected component of $\{x \in \Omega : f(p) - \gamma_1 \leq x \leq f(p) + \gamma_2\}$ containing p. A point which is in this neighborhood for small values of γ_1 and γ_2 is "close" to p. Note that points p and p' in Figure 2 are "close" in this sense, but points p and q' in Figure 1 remain far apart. We show that if p destroys a persistent homology group in f, then the neighborhood of p, $\{x \in \Omega : f(p) - \gamma_1 \leq x \leq f(p) + \gamma_2\}$, contains a point q which destroys a persistent homology group in $\hat{f}$. The values of γ_1 and γ_2 depend upon the persistence of the homology group and the difference δ between f and $\hat{f}$.

Theorem 1, one of our main results, states that every destroying critical point of f contains in its "neighborhood" a similar destroying critical point of $\hat{f}$. However, to construct a matching of critical points, we need each destroying critical point of $\hat{f}$ to be in only one such neighborhood. We establish this stronger result for critical points which are local maxima of functions on manifolds. More specifically, we show that the neighborhoods of local maxima with large persistence are pairwise disjoint.

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 [6]. We will use the *singular homology*. 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 [5, 9]. As discussed in Carlsson and Zomorodian [9], computing the persistent homology groups over non-fields is an unsolved and perhaps intractable problem. 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_k, between their homology groups, $H_k(X) \stackrel{f_k}{\to} H_k(Y)$. This property is carried over the composition of maps, that is, $(f \circ g)_k = f_k \circ g_k$. In our case, the maps between spaces will be *inclusions* maps. This means, if $X \subseteq Y$, we will consider the map $H_k(X) \stackrel{\ell}{\to} H_k(Y)$ where ℓ is induced by the inclusion map $\iota\colon X \to Y$. From now on, we take the liberty of dropping the subscript k from $H_k(X)$ when it is clear from the context. For $X \subseteq Y$ the *relative* homology group of Y with respect to X is given by $H(Y, X) = H(Y)/\ell(H(X))$ where $H(X) \stackrel{\ell}{\to} H(Y)$ is the homomorphism induced by inclusion $\iota\colon X \to Y$.

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

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

We will use a specific type of sequence called Mayer-Vietoris sequence which is known to be exact.

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* [6, p. 149]. The map ∂ is the connecting homomorphism given by boundary maps [6, p. 116].

2.2 Notation

We use the following notation to define the region bounded by $f^{-1}(a)$ and $f^{-1}(b)$. 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 ∞.

2.3 Destruction

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

If $X \subseteq Z \subseteq \Omega$ and $Y \subseteq \Omega$, then we say that Y *destroys* the image of $h_x \in H(X)$ in $H(Z)$, if $h_z \in H(Z)$ is the image of h_x under the mapping $H(X) \to H(Z)$ and Y destroys h_z. We encounter this situation repeatedly where X is some level set $f^{-1}(a)$ and Y is a point. For brevity, we say that point q destroys $h \in f^{-1}(a)$ if point q destroys the image of h in $H(\underline{F}_a^{f(q)})$.

A function $f : \Omega \to \mathbb{R}$ is *point destructible* if whenever $h \in H(\underline{F}_a^b)$ is destroyed by $f^{-1}(b)$, then h is destroyed by some point $q \in f^{-1}(b)$.

2.4 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 $a \in \mathbb{R}$, let $N_\epsilon(f^{-1}(a)) = \{x \in \Omega : a - \epsilon < f(x) < a + \epsilon\}$ denote the *open ϵ-neighborhood* of $f^{-1}(a)$. The first property we require is that the topology of $f^{-1}(a)$ is similar (isotopic) 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 continous $\phi : X \times I \to X$ 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 $t \in \mathbb{R}$. 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 $H(X)$ is isomorphic to $H(Y)$.

We say that 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 $\epsilon \le \epsilon_0$, the set $f^{-1}(a)$ is a strong deformation retract of $N_\epsilon(f^{-1}(a)) = \{x : a - \epsilon < f(x) < a + \epsilon\}$.

Piecewise linear functions on finite simplicial complexes are LR. If a continous function is Morse, then the function is LR. (See Milnor [7, pp. 12–20] for a proof. Milnor actually proves that $\{x : f(x) \le a\}$ is a deformation retract of $\{x : f(x) \le a + \epsilon\}$ but his proof also shows that $f^{-1}(a)$ is a strong deformation retract of $\{x : a - \epsilon < f(x) < a + \epsilon\}$.) Some properties of LR functions will be used in our proofs. We state them in Lemma 1 and Lemma 2. The proofs are given in the appendix.

Lemma 1 says that if f is LR, then we can replace $\{x : a \le f(x) \le b\}$ by suitably chosen small neighborhoods without changing its homology.

Lemma 1. *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 $\{x : a \le f(x) \le b\}$ is a strong deformation retract of $\{x : a \le f(x) < b + \epsilon\}$ and set $\{x : a \le f(x) < b\}$ is a strong deformation retract of $\{x : a - \epsilon < f(x) < b\}$.*

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

Lemma 2. *Let $f : \Omega \to \mathbb{R}$ be a continous, LR function. For any non-zero $h \in H(f^{-1}(a))$, if $H(\underline{F}_a^\infty)$ destroys h, then for some $b \ge a$, the image of h under the mapping $H(f^{-1}(a)) \to H(\underline{F}_a^b)$ is destroyed by $f^{-1}(b)$.*

Lemma 2 has the following corollary.

Corollary 3. *Let $f : \Omega \to \mathbb{R}$ be a continuous, point destructible, LR function. For any non-zero $h \in H(f^{-1}(a))$, if $H(\underline{F}_a^\infty)$ destroys h, then some $q \in \Omega$ destroys h.*

Proof. By Lemma 2, there exists some $b > a$ such that the image h' of h under the mapping $H(f^{-1}(a)) \to H(\underline{F}_a^b)$ 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. □

3 Persistence

Intuitively, the persistence of a point $p \in \Omega$ is the "age" of the "oldest" homology element destroyed by p. Formally, the *persistence* of point $p \in \Omega$ is

$$\Pi_f(p) = \limsup\{f(p) - a : p \text{ destroys some non-zero } h \in H(f^{-1}(a))\}.$$

We use $\limsup$ in place of $\max$ because it is possible that p destroys non-zero elements of $H(f^{-1}(a + \epsilon))$ for any $\epsilon > 0$ but not elements of $H(f^{-1}(a))$.

If point p has persistence γ, and $f(p) - \gamma < a < f(p)$, does p destroy some element of $H(f^{-1}(a))$? As we show below, the answer is yes.

We will need to use subsets of Ω which are bounded by two different functions in our discussion of persistence. The set $X = \{x : a \le f(x) \text{ and } g(x) < b\}$ is *properly bounded* by f and g at a and b respectively, if the level sets $f^{-1}(a)$ and $g^{-1}(b)$ are disjoint and X is non-empty. Notice that X includes $f^{-1}(a)$ at the bottom but not $g^{-1}(b)$ at the top.

In Lemma 4 we will establish a result relating three functions and the spaces delimited by their level sets. Later we will set these level sets only to those of two functions f and g that are used to establish the stability result.

Let $f_1, f_2, f_3 : \Omega \to \mathbb{R}$ be three continuous LR functions. Let q be a point in Ω and let $a_3 = f_3(q)$. Let $X = \{x \in \Omega : a_1 \le f_1(x) \text{ and } f_3(x) < a_3\}$ and $X' = \{x \in \Omega : a_2 \le f_2(x) \text{ and } f_3(x) < a_3\}$ be properly bounded subsets of Ω where $X' \subseteq X$ and $f_1^{-1}(a_1) \cap f_2^{-1}(a_2) = \emptyset$, see Figure 3. The essence of the following lemma is that if q destroys a generator in $H(f_1^{-1}(a_1))$, then it necessarily destroys some generator in $H(f_2^{-1}(a_2))$. Some subtle aspect of this destruction is narrated in the caption of Figure 3.

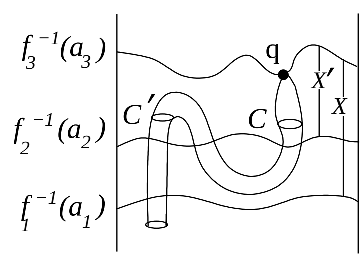

Figure 3: Sets $X = \{x \in \Omega : a_1 \le f_1(x) \text{ and } f_3(x) < a_3\}$ and $X' = \{x \in \Omega : a_2 \le f_2(x) \text{ and } f_3(x) < a_3\}$. Cycle C generates $h_x \in H(X)$ and $h'_x \in H(X')$ where both h_x and h'_x are destroyed at q. Elements h_x and h'_x are the images of some $h_1 \in H(f_1^{-1}(a_1))$ and some $h_2 \in H(f_2^{-1}(a_2))$, respectively. The mapping $H(X') \to H(X)$ sends h'_x to h_x. Cycle C' generates $h_x \in H(X)$ and $h''_x \in H(X')$ where h_x is destroyed at q but h''_x is not, even though the mapping $H(X') \to H(X)$ sends h''_x to h_x.

Lemma 4. *If q destroys some non-zero $h_x \in H(X)$ and h_x is the image of some $h_1 \in H(f^{-1}(a_1))$, then h_x is the image of* some *non-zero $h'_x \in H(X')$ which is destroyed by q. Moreover, h'_x is the image of some non-zero $h_2 \in H(f^{-1}(a_2))$.*

Proof. Let $h_x \in H(X)$ be the image of some $h_1 \in H(f^{-1}(a_1))$ under the mapping $H(f^{-1}(a_1)) \to H(X)$ where h_x is destroyed by q. By Mayer-Vietoris, the sequence

$$H(X') \to H(X' \cup \{q\}) \oplus H(X) \to H(X \cup \{q\})$$

is exact. Since the induced mapping $H(X) \to H(X \cup \{q\})$ sends h_x to 0, the mapping $H(X' \cup \{q\}) \oplus H(X) \to H(X \cup \{q\})$ sends $(0 \oplus h_x)$ to 0. Since the sequence is exact, there is some h'_x whose image is $(0 \oplus h_x)$ under the mapping $H(X') \to H(X' \cup \{q\}) \oplus H(X)$. Thus h_x is the image of h'_x and h'_x is destroyed by q.

We now prove that h'_x is the image of some $h_2 \in H(f_2^{-1}(a_2))$. Since f_2 is LR, there exists some $\epsilon_1 > 0$ such that $H(f_2^{-1}(a_2)) \to H(N_{\epsilon'}(f_2^{-1}(a_2)))$ is an isomorphism for all $\epsilon' \leq \epsilon_1$. By Lemma 1, $H(X') \to H(X' \cup N_{\epsilon'}(f_2^{-1}(a_2)))$ is an isomorphism for all $\epsilon' \leq \epsilon_2$. Since $f_1^{-1}(a_1) \cap f_2^{-1}(a_2) = \emptyset$, there is some ϵ_3 such that $N_{\epsilon_3}(f_2^{-1}(a_2)) \subseteq H(X)$. Let ϵ be the smaller of ϵ_1, ϵ_2 and ϵ_3.

Let Y equal $N_\epsilon(f_2^{-1}(a_2))$. Let $Z = \{x \in \Omega : a_1 \leq f_1(x)$ and $f_2(x) < a_2 + \epsilon\}$ and $Z' = \{x \in \Omega : a_2 - \epsilon < f_2(x)$ and $f_3(x) < a_3\}$. The following commutative diagram gives the relevant mappings between homology groups:

$$\begin{tikzcd}
 & & \begin{matrix} h_1 \in \\ H(f_1^{-1}(a_1)) \end{matrix} \arrow[d] \\
h_2 \in H(f_2^{-1}(a_2)) \arrow[r, equal] \arrow[d] & \begin{matrix} h_y \in \\ H(Y) \end{matrix} \arrow[r] \arrow[d] & H(Z) \ni h_z \arrow[d] \\
h'_x \in H(X') \arrow[r, equal] & \begin{matrix} h'_z \in \\ H(Z') \end{matrix} \arrow[r] & H(X) \ni h_x
\end{tikzcd}$$

Element h_x is the image of $h_1 \in H(f_1^{-1}(a_1))$. Let h_z be the image of h_1 under the mapping $H(f^{-1}(a_1)) \to H(Z)$. Let h'_z be the image of h'_x under the mapping $H(X') \to H(Z')$. Element h_x is the image of both h_z and h'_z under the respective mappings $H(Z) \to H(X)$ and $H(Z') \to H(X)$.

By Mayer-Vietoris, the sequence

$$H(Y) \to H(Z) \oplus H(Z') \to H(X)$$

is exact. Since the mapping $H(Z) \oplus H(Z') \to H(X)$ sends $(h_z \oplus h'_z)$ to $(h_x - h_x) = 0$, element $(h_z \oplus h'_z)$ must be in the image of some $h_y \in H(Y)$. Since $H(f_2^{-1}(a_2)) \to H(Y)$ is an isomorphism, there is some $h_2 \in H(f_2^{-1}(a_2))$ whose image is h_y under the mapping $H(f_2^{-1}(a_2)) \to H(Y)$.

All the mappings of homology groups are induced by the inclusion mapping and thus the diagram is commutative. Since $H(X') \to H(Z')$ is an isomorphism, the inverse mapping $H(Z') \to H(X')$ takes h'_z to h'_x. The mapping $H(f_2^{-1}(a_2)) \to H(Y) \to H(Z') \to H(X')$ sends element $h_2 \in H(f_2^{-1}(a_2)$ to $h'_x \in H(X')$. Thus $h'_x \in H(X')$ is the image of some non-zero $h_2 \in H(f_2^{-1}(a_2)$. □

Setting $f_1 = f_2 = f_3$ gives the following corollary:

Corollary 5. *Let $f : \Omega \to \mathbb{R}$ be a continuous, LR function. If $q \in \Omega$ has persistence γ, then for every a where $f(a) - \gamma < a < f(q)$, point q destroys some element of $H(f^{-1}(a))$.*

4 Stability

In this section we prove one of our main results, Theorem 1. Let f and g be two functions defined on Ω. We say $|f - g| < \delta$ if $|f(x) - g(x)| < \delta$ for all $x \in \Omega$. We show that if q is a destructor for f with persistence $\gamma > 2\delta$, there is a point q' which is a destructor for g where q and q' lie in the same connected component of the space $\underline{F}_{f(q)-\gamma}^{f(q)+2\delta}$. (Wherever we use the term connected component, we always mean *path connected*.) 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. [3] but also in the domain.

We will need the concept of *chains* and *cycles* that define the homology groups. Chains are formal sums of maps from standard simplices into the domain Ω and cycles are chains which have no boundary. The boundary of a chain is always a cycle. For details, see Hatcher [6].

Lemma 6. *If non-zero $h \in H(\underline{F}_a^b)$ is destroyed by point q and cycle $C \subseteq \underline{F}_a^b$ generates h, then C and q lie in the same connected component of* $\mathrm{cl}(\underline{F}_a^b)$.

Proof. Since q destroys h, cycle C is the boundary of some chain $D \subseteq \underline{F}_a^b \cup \{q\}$. The chain D must contain point q or else C would be the boundary of a chain in $\underline{F}_a^b$ and h would be 0. Since all points in D other than q lie in $\underline{F}_a^b$, point q is in $\mathrm{cl}(\underline{F}_a^b)$. Since D connects C and q in $\mathrm{cl}(\underline{F}_a^b)$, cycle C and point q lie in the same component of $\mathrm{cl}(\underline{F}_a^b)$. □

Theorem 1. *Let $f, g : \Omega \to \mathbb{R}$ be continuous, point destructible, LR functions on Ω where $|f - g| < \delta$. If q destroys non-zero $h_f \in H(f^{-1}(a))$ and $|f(q) - a| > 2\delta$*

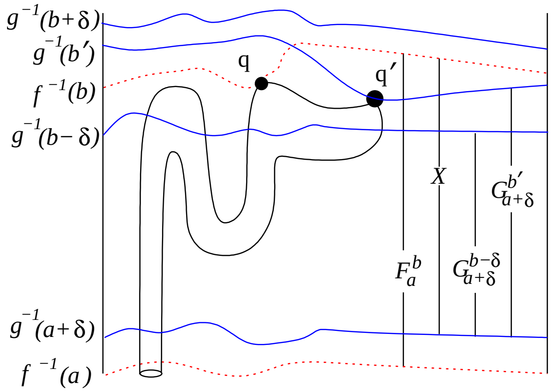

Figure 4: Sets $\underline{F}_a^b$, $X = \{x \in \Omega : a + \delta \leq g(x)$ and $f(x) < b\}$, $\underline{G}^{b-\delta}_{a+\delta}$ and $\underline{G}^{b'}_{a+\delta}$.

and σ is the connected component of $\mathrm{cl}(\underline{F}_a^{f(q)+2\delta})$ *containing q, then σ contains a point q' for g which destroys some $h_g \in H(g^{-1}(a+\delta))$ and $f(q) - \delta \leq g(q') \leq f(q) + \delta$.*

Proof. 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 $b - a > 2\delta$, space X is properly bounded. Note that $\underline{G}^{b-\delta}_{a+\delta} \subseteq X \subseteq \underline{F}_a^b$ and so there exist homomorphisms $H(\underline{G}^{b-\delta}_{a+\delta}) \to H(X) \to H(\underline{F}_a^b)$ induced by inclusions.

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

$$\begin{tikzcd}
h_g \in H(g^{-1}(a+\delta)) \arrow[d] & & h_f \in H(f^{-1}(a)) \arrow[d] \\
H(\underline{G}^{b-\delta}_{a+\delta}) \arrow[r] \arrow[d] & h_x \in H(X) \arrow[r] \arrow[d] & H(\underline{F}_a^b) \ni h'_f \arrow[d] \\
h'_g \in H(\underline{G}^{b'}_{a+\delta}) \arrow[d] & H(X \cup \{q\}) \arrow[r] \arrow[dl] & H(\underline{F}_a^b \cup \{q\}) \\
H(\underline{G}^{\infty}_{a+\delta}) & &
\end{tikzcd}$$

The value b' will be defined below.

Let $h'_f \in H(\underline{F}_a^b)$ be the image of h_f under the mapping $H(f^{-1}(a)) \to H(\underline{F}_a^b)$. Since h'_f is destroyed at $f^{-1}(b)$, element h'_f is non-zero. By Lemma 4, element h'_f is the image of some element $h_x \in H(X)$ which is destroyed by q and is the image of some $h_g \in H(g^{-1}(a+\delta))$. Since h'_f is non-zero, elements h_x and h_g are non-zero.

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

By Corollary 3, 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 $H(g^{-1}(a+\delta)) \to H(\underline{G}^{f(q')}_{a+\delta})$ is destroyed by q'.) Let b' equal $f(q')$. Since $\underline{G}^{b-\delta}_{a+\delta}$ is a subset of X, the image of h_g under the mapping $H(g^{-1}(a+\delta)) \to H(\underline{G}^{b-\delta}_{a+\delta})$ is non-zero and so $b' \geq b - \delta$. Let h'_g be the image of h_g under the mapping $H(g^{-1}(a+\delta)) \to H(\underline{G}^{b'}_{a+\delta})$.

Element h_g is generated by some cycle C in $g^{-1}(a+\delta)$. Since the image of h_g is $h_x \in H(X)$ under the mapping $H(g^{-1}(a+\delta) \to H(X)$, cycle C also generates h_x. Similarly, cycle C generates $h'_g \in H(\underline{G}^{b'}_{a+\delta})$ and $h'_f \in H(\underline{F}_a^b)$.

Since h_x is destroyed by q, cycle C is the boundary of some chain $D \subseteq X \cup f^{-1}(b)$. Since $|f - g| \leq \delta$, set $X \cup f^{-1}(b)$ is a subset of $\underline{G}^{b+\delta}_{a+\delta} \cup g^{-1}(b+\delta)$ and so chain D is a subset of $\underline{G}^{b+\delta}_{a+\delta} \cup g^{-1}(b+\delta)$. Since $h'_g \in H(\underline{G}^{b'}_{a+\delta})$ is non-zero, space $\underline{G}^{b+\delta}_{a+\delta} \cup g^{-1}(b+\delta)$ cannot be a supspace of $\underline{G}^{b'}_{a+\delta}$. Thus b' is at most $b + \delta$.

As noted above, cycle C generates $h'_g \in H(\underline{G}^{b'}_{a+\delta})$ and $h'_f \in H(\underline{F}_a^b)$. Point q destroys h'_f which is generated by C. By Lemma 6, point q must lie in the connected component of $\mathrm{cl}(\underline{F}_a^b)$ containing C. Similarly, since q' destroys h'_g, point q' must lie in the connected component of $\mathrm{cl}(\underline{G}^{b'}_{a+\delta})$ containing C. Since $\mathrm{cl}(\underline{F}_a^b) \subseteq \mathrm{cl}(\underline{F}_a^{b+2\delta})$ and $\mathrm{cl}(\underline{G}^{b'}_{a+\delta}) \subseteq \mathrm{cl}(\underline{G}^{b+\delta}_{a+\delta}) \subseteq \mathrm{cl}(\underline{F}_a^{b+2\delta})$, points q and q' must lie in the same connected component σ of $\mathrm{cl}(\underline{F}_a^{b+2\delta})$. ◻

5 Computing persistence

Theorem 1 can be used to compare two real valued functions f and g defined on a topological space Ω. The key computation to apply Theorem 1 is:

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

We use Betti numbers and their persistent counterparts to compute (i). Recall that, for a homology group $H(X)$, the Betti number is $\beta(X) = \dim H(X)$. It gives the number of generators in $H(X)$. The persistent Betti numbers relate the homology classes of one space into the other. For $Y \subseteq X$, let H^Y_X be the image of the map $H(Y) \to H(X)$ induced by inclusion $Y \to X$. Define $\beta(Y, X) = \dim H^Y_X$. In words, $\beta(Y, X)$ counts the number of non-zero generators of $H(Y)$ that remain so in the larger space X.

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 $H(f^{-1}(a))$ gets destroyed by p. Let

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

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

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

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

Notice that Theorem 1 can be applied to any point $p \in \Omega$ and any value $a < f(p)$. However, for a canonical computation one can focus on the critical points of the functions. We define the critical points of f as follows.

Definition 1. *A point $p \in \Omega$ is critical for $f : \Omega \to \mathbb{R}$ if $H(F_a^b) \to H(F_a^b \cup \{p\})$ is not an isomorphism for some $a, b \in \mathbb{R}$ where either $f(p) = a$ or $f(p) = b$.*

The above definition of critical points is similar to that of Cohen-Steiner et al. [3] with a distinction that the relevant space is an interval set whose lower level may not be at $-\infty$. Also, notice that the space F_a^b could be above or below the level of $f(p)$.

Compute the critical points of f according to the above definition. Let $p_0, p_1..p_k$ be the 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 persistence $\Pi_f(p_j)$ 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 persistence $\Pi_f(p_j)$ is at least $f(p_j) - f(p_{i-1})$. Thus we compute $\Pi_f(p_j)$ 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 persistence of a critical point q of g is measured similarly by $\Pi_g(q)$.

To compare f and g, we check if any critical point p of f has $\Pi_f(p)$ greater than a user supplied parameter τ. If so, we search for a critical point q of g in the connected component of $\mathrm{cl}(\underline{F}_{f(p)-\tau}^{f(p)+\tau})$ so that $\Pi_g(q) > \tau$ and $|f(p) - g(q)| \leq \frac{\tau}{2}$. If $\tau > 2\delta$, such a q exists by Theorem 1.

5.1 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. The critical points of f and g are located at the triangulation vertices. 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 which are point destructible.

Carlsson and Zomorodian in [9] show how to compute persistent Betti numbers for homology groups of filtered simplicial complexes over any field. However, spaces $\underline{F}_a^b$ and $\underline{F}_a^b \cup \{p\}$ are not closed. To compute their persistent Betti numbers, β_a^b and λ_a^p, we collapse them onto closed sets which are simplicial complexes.

Consider the space $\underline{F}_a^b$. Let $t \subseteq \Omega$ be a simplex with a point in this space. If each vertex v of t either has $f(v) \geq b$ or $a \leq f(v) < b$, the subset $t \cap \underline{F}_a^b$ can be collapsed to the face of t made by the vertices whose values lie in $[a, b)$. This cannot be done for simplices that cut across the levels of a and b. These simplices have vertices with values above b and also below a. For such a simplex we take an edge $e = \{u, v\}$ where $f(u) < a$ and $f(v) \geq b$ and consider a point x on this edge where $a < f(x) < b$. We divide t by starring from x to all its vertices. After subdividing all such simplices we obtain a subdivision $\tilde{\Omega}$ of Ω which has no simplex cutting across the interval $[a, b]$. Consider the simplicial complex made by the collection of simplices in $\tilde{\Omega}$ that have all vertices with values in $[a, b)$. The underlying space of this simplicial complex is a deformation retract of $\underline{F}_a^b$ and therefore has homology groups isomorphic to that of $\underline{F}_a^b$.

6 Maxima

We show that the neighborhoods of local maxima with large persistence are pairwise disjoint. This enables us to establish a matching of such critical points.

The idea of the proof is as follows. Consider two local maxima, $p, p' \in \Omega$, where $f(p) \leq f(p')$. If p destroys non-zero $h \in H_k(f^{-1}(a))$, then k equals $d-1$. Let σ be the connected component of $\underline{F}_a^\infty$ containing p. If σ is a manifold with boundary, then $\underline{F}_a^{f(p)} \cup \{p\}$ must contain σ. Since $f(p) \leq f(p')$, set $\underline{F}_a^{f(p)} \cup \{p\}$ does not contain p' and therefore point p' is not in σ.

Lemma 7. *Let $\underline{F}_a^{\infty}$ be an oriented d-manifold with boundary. If p is a local maximum and p destroys non-zero $h \in H_k(\underline{F}_a^{f(p)})$, then k equals $d-1$.*

Proof. See appendix. □

As previously noted, we always mean path connected when referring to connected components.

Lemma 8. *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 [6, 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 $a_1 = a_2 = \ldots = a_k$. □

Lemma 9. *Let M be a connected, oriented d-manifold with non-empty boundary. If $\partial M \subseteq M' \subseteq M$ and $H_{d-1}(\partial M) \to H_{d-1}(M')$ takes 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 $H(M - \{p\})$ is isomorphic to $H(M - B)$.

The mapping $H_{d-1}(\partial M) \to H_{d-1}(M') \to H_{d-1}(M - \{p\}) \to H_{d-1}(M - B)$ 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 8, 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 8, 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 $\sigma_p^f(\gamma)$ represent the connected component of $\underline{F}_{f(p)-\gamma}^{\infty}$ containing p. We prove that the neighborhoods $\sigma_p^f(\gamma)$ of points with persistence greater than γ are pairwise disjoint. (See Figure 5.)

Theorem 2. *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 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 4, point p_0 destroys some non-zero element of $H_k(f^{-1}(f(p_0) - \gamma'))$. By Lemma 7, k equals $d-1$. By Lemma 10 (appendix), point p_0 destroys a non-zero elements of $H_{d-1}(\partial\sigma_{p_0}^f(\gamma'))$. In Lemma 9 putting $M = \sigma_{p_0}^f(\gamma')$ and M' equal the connected component of $\underline{F}_{f(p_0)-\gamma'}^{f(p_0)}$ containing p_0, we conclude $\underline{F}_{f(p_0)-\gamma'}^{f(p_0)} \cup \{p_0\}$ contains $\sigma_{p_0}^f(\gamma')$ 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.

Theorem 3. *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 maxima of f and let q and q' be local maxima of g such 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 5.) □

Proof of (ii). By Theorem 2, $\sigma_q^g(\gamma)$ 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 $\sigma_p^f(\gamma - 2\delta)$. Similarly, if $\sigma_p^f(\gamma - 2\delta)$ intersects $\sigma_{q'}^g(\gamma - 2\delta)$, then $\sigma_{q'}^g(\gamma)$ contains $\sigma_p^f(\gamma - 2\delta)$. However, $\sigma_q^g(\gamma)$ and $\sigma_{q'}^g(\gamma)$ are disjoint. Thus, if $\sigma_p^f(\gamma - 2\delta)$ intersects $\sigma_q^g(\gamma - 2\delta)$, then $\sigma_p^f(\gamma - 2\delta)$ does not intersect $\sigma_q^g(\gamma - 2\delta)$. (See Figure 5.) □

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 1. However, there may be no such matching. In fact, f may contain a set of critical points 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. As before, let $\sigma_p^f(\gamma)$ and $\sigma_q^g(\gamma)$ be the connected components of $\underline{F}_{f(p)-\gamma}^\infty$ and $\underline{G}_{g(q)-\gamma}^\infty$ containing p and q, respectively. 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 3, 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 $|f(p) - f(q')| < \gamma/4$. (By Theorem 1 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 $|f(q') - f(p)| < \gamma/4$.

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

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 $|f(p) - f(q')| < \gamma/4$;
6. For each unmatched $q \in M_g(\gamma)$, match q with $p' \in M_f(\gamma/2) \cap \sigma_q$ where $|g(q) - g(p')| < \gamma/4$.

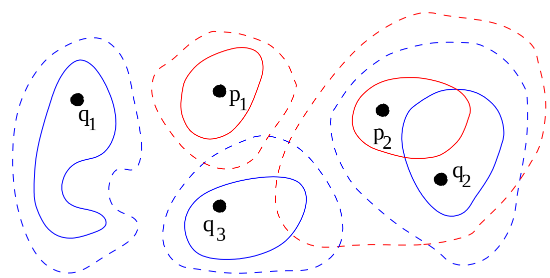

Figure 5: Local maxima and their neighborhoods. Solid lines around points p_i and q_i are neighborhoods $\sigma_{p_i}^f(\gamma/2)$ and $\sigma_{q_i}^g(\gamma/2)$. Dotted lines are neighborhoods $\sigma_{p_i}^f(\gamma)$ and $\sigma_{q_i}^g(\gamma)$. Neighborhood $\sigma_{p_2}^f(\gamma/2)$ intersects $\sigma_{q_2}^g(\gamma/2)$ 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 $M_f(\gamma/2)$.

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 covers $M_f(\gamma)$ and $M_g(\gamma)$ such that every matched pair (p, q) where $p \in M_f$ and $q \in M_g$ is $(\gamma, \gamma/4)$-close.*

Proof. By Theorem 2, 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 3, each σ_p intersects at most one σ_q and vice versa. Thus Step 4 gives a one to one partial matching.

By Theorem 1, σ_p contains some point $q' \in M(\gamma/2)$ 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 and $\sigma_q^g(\gamma/2)$. This holds true even if p and q are matched in Steps 5 or 6. By Theorem 3 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. □

8 Discussions

Results on stability of topological persistence can be used in shape matching. 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, if we have two similar shapes, the distance functions defined by their point samples are similar. As observed in previous works [1, 2], the results on persistence apply to such functions. Our results in this paper have some notable connections to a shape matching algorithm proposed by Dey et al. [4]. According to our results, we can expect that similar shapes have similar structures for maxima with large persistence in terms of the interval sets. The algorithm in [4] uses maxima and their stable manifolds for matching. We suspect that these stable manifolds are playing the role of connected components as suggested in this paper. Perhaps the performance of the matching algorithm in [4] now can be improved and better explained by our results. We plan to address this issue in future work.

Persistence diagrams [3] can be used to match shapes based on the critical values of the distance function. Does adding critical points increase the discrimination of this matching? Figure 6 provides such an example. The two unsimilar shapes in this figure has matching persistence diagrams. The distance function for each shape contains three local maxima and two saddle points. For each shape, two of the local maxima and one of the saddle points has long persistence while the other local maxima and saddle point have short persistence. (Persistence of saddle points which “create homology” groups is defined in [3].) Thus the persistence diagrams match.

Now assume that the two shapes are registered so that each q_i lies on top of p_i. Even with such a registration, the local maxima p_1 does not match with q_1 since p_1 and q_1 have different persistence and p_1 does not match with q_2 or q_3 since they are not in the neighborhood of p_1 in the domain. Thus, our matching based on closeness of local maxima in both the range and domain distinguishes these shapes.

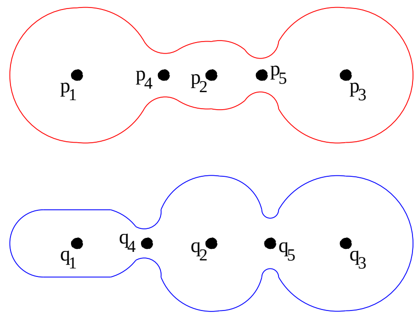

Figure 6: Two shapes with matching persistence diagram. Matching based on closeness in both the range and domain distinguishes these shapes. Local maxima p_1, p_3, q_1, q_3 and saddle points p_5, q_5 have long persistence. Local maxima p_2, q_1 and saddle points p_4, q_4 have short persistence.

Acknowledgement: We thank Prof. Michael Davis from OSU Dept. of Mathematics for providing the proof of Lemma 8.

References

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

[2] 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.

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

[4] 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.

[5] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci.*, pages 454–463, 2000.

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

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

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

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

9 Appendix

Proof of Lemma 1. 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 $\epsilon < \epsilon_1$ and $a + \epsilon < b$. 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 $\epsilon < \epsilon_2$ and $a < b-\epsilon$. Let ϵ_0 be the minimum of ϵ_1 and ϵ_2 and $b - a$.

For any $\epsilon \leq \epsilon_0$, let ϕ^a_ϵ 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^b_ϵ 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

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

ϕ_ϵ is constant on $\{x : a \leq x \leq b\}$ and continously 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 $\{x : a - \epsilon < f(x) < b\}$. □

Proof of Lemma 2. Since $H(f^{-1}(a)) \to H(\underline{F}^\infty_a)$ sends h to zero, element h is the boundary of some chain $C \subseteq \underline{F}^\infty_a$. Chain C is compact. (See [8, p. 71].) Thus $\{f(x) : x \in C\}$ is compact and has a maximum value b'. Since $C \subseteq \underline{F}^{b'}_a$, the mapping $H(f^{-1}(a)) \to H(\underline{F}^{b'}_a)$ sends h to zero.

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

Let h' be the image of h under the mapping $H(f^{-1}(a)) \to H(\underline{F}^b_a)$. If h' were zero, then h would be the boundary of some chain $C' \subseteq \underline{F}^b_a$. Since C' is compact, chain C' would also be a subset of $\underline{F}^{\tilde{b}}_a$ 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 $H(\underline{F}^b_a \cup f^{-1}(b))$ is isomorphic to $H(\underline{F}^{b+\epsilon}_a)$ for all $\epsilon \leq \epsilon_0$. If $H(\underline{F}^b_a) \to H(\underline{F}^b_a \cup f^{-1}(b))$ does not map h' to zero, then $H(\underline{F}^b_a) \to H(\underline{F}^{b+\epsilon}_a)$ does not map h' to zero for all $\epsilon \leq \epsilon_0$, and b does not equal $\liminf\{\tilde{b} : H(f^{-1}(b)) \to H(\underline{F}^{\tilde{b}}_a)$ sends h to zero$\}$. Thus, $H(\underline{F}^b_a) \to H(\underline{F}^b_a \cup f^{-1}(b))$ maps h' to zero and h' is destroyed by $f^{-1}(b)$. □

Proof of Lemma 7. Since $\underline{F}^\infty_a$ 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 $\underline{F}^{f(p)}_a$.

By the Mayer-Vietoris Theorem, the sequence

$$H_k(B_p-\{p\}) \to H_k(B_p) \oplus H_k(\underline{F}^{f(p)}_a) \to H_k(\underline{F}^{f(p)}_a \cup \{p\})$$

is exact. Since the mapping $H_k(\underline{F}^{f(p)}_a) \to H_k(\underline{F}^{f(p)}_a \cup \{p\})$ sends h to zero, the mapping $H_k(B_p) \oplus H_k(\underline{F}^{f(p)}_a) \to H_k(\underline{F}^{f(p)}_a \cup \{p\})$ sends $(0 \oplus h)$ to zero. Since the sequence is exact, element $(0 \oplus h)$ is the image of some non-zero $h' \in H_k(B_p)$ under the mapping $H_k(B_p - \{p\}) \to H_k(B_p) \oplus H_k(\underline{F}^{f(p)}_a)$. 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 is an element of $H_{d-1}(\underline{F}^{f(p)}_a)$ and so k equals $d - 1$. □

Lemma 10. *Let $\Omega \subseteq \Omega'$ be topological spaces, let $\sigma'_1, \ldots, \sigma'_m$ be the pathwise connected components of Ω', and let σ equal $\sigma'_i \cap \Omega$ for each i. If the mapping $H_k(\Omega) \to H_k(\Omega')$ sends non-zero $h \in H_k(\Omega)$ to zero, then there exists some non-zero $h_i \in H_k(\sigma_i)$ such that the mapping $H_k(\sigma_i) \to H_k(\Omega')$ sends h_i to zero. Moreover, if point p is an element of $\Omega' - \Omega$, and the mapping $H_k(\sigma_i) \to H_k(\Omega' - \{p\})$ does not send h_i to zero, then p is an element of σ'_i.*

Proof. Since the σ_i are pairwise disjoint, the homology group $H_k(\Omega)$ is isomorphic to $H_k(\sigma_1) \oplus \cdots \oplus H_k(\sigma_m)$. (See [8, Theorem 4.13, p. 69].) This isomorphism takes h to $(h_1 \oplus \cdots \oplus h_m)$ where $h_i \in H_k(\sigma_i)$. At least one of these h_i must be non-zero.

Since the σ'_i are pairwise disjoint, the homology group $H_k(\Omega')$ is isomorphic to $H_k(\sigma'_1) \oplus \cdots \oplus H_k(\sigma'_m)$. The mapping $H_k(\Omega) \to H_k(\Omega')$ sends h to zero, so the mapping

$$H_k(\sigma_1) \oplus \cdots \oplus H_k(\sigma_m) \to H_k(\sigma'_1) \oplus \cdots \oplus H_k(\sigma'_m)$$

sends $(h_1 \oplus \cdots \oplus h_m)$ to $(0 \oplus \cdots \oplus 0)$. Thus, the mapping $H_k(\sigma_i) \to H_k(\sigma'_i) \to H_k(\Omega')$, sends non-zero $h_i \in H_k(\sigma_i)$ to zero.

Let point p be an element of $\Omega' - \Omega$ where the mapping $H_k(\sigma_i) \to H_k(\Omega' - \{p\})$ sends h_i to some non-zero $h' \in H_k(\Omega' - \{p\})$. If $p \notin \sigma'_i$, then $\sigma'_i \subset \Omega' - \{p\}$ and the mapping $H_k(\sigma_i) \to H_k(\sigma'_i) \to H_k(\Omega' - \{p\})$ sends h_i to zero. Thus σ'_i contains p. □