Discrete Comput Geom 37:103–120 (2007)
DOI: 10.1007/s00454-006-1276-5

Discrete & Computational
Geometry

Stability of Persistence Diagrams*

David Cohen-Steiner,[1] Herbert Edelsbrunner,[2] and John Harer[3]

[1]INRIA, 2004 Route des Lucioles,
BP 93, 06904 Sophia-Antipolis, France
dcohen@sophia.inria.fr

[2]Department of Computer Science, Duke University,
Durham, NC 27708, USA
edels@cs.duke.edu
and
Geomagic,
Research Triangle Park, NC 27709, USA

[3]Department of Mathematics, Duke University,
Durham, NC 27708, USA
harer@math.duke.edu

Abstract. The persistence diagram of a real-valued function on a topological space is a multiset of points in the extended plane. We prove that under mild assumptions on the function, the persistence diagram is stable: small changes in the function imply only small changes in the diagram. We apply this result to estimating the homology of sets in a metric space and to comparing and classifying geometric shapes.

1. Introduction

In this paper we consider real-valued functions on topological spaces and use the concept of persistence to study their qualitative and quantitative behavior. More specifically, we encode the topological characteristics of a function in what we call its persistence diagram and study the stability of this encoding.

Motivation. Topological spaces and functions on them are common types of data in all disciplines of the natural sciences and engineering and their computational treatment is

* The NSF partially supported the first two authors under Grant CCR-00-86013 and the third author under Grant DMS-01-07621. The DARPA partially supported all three authors under Grant HR0011-05-1-0007.

of central concern in support areas, such as visualization. However, much of the data that is currently available is too large and detailed for direct human consumption. Moreover, measurement errors and discretization problems inherent to any acquisition process add irrelevant complexity to the data. A crucial problem is therefore the assessment of the importance of a feature, one goal being the emphasis of dominant features and the suppression of all others. This problem immediately raises two questions, namely: "What is a feature?" and "How do we measure the relative importance of features?" We argue that the importance and the stability of a feature are overlapping, if not identical, concepts, as importance can be quantified in terms of the amount of change necessary to eliminate a feature.

Results and Prior Work. The main result of this paper is the stability of the persistence diagram of a function on a topological space. The persistence diagram, introduced by Edelsbrunner et al. [11], is a point set in the extended plane that encodes the difference in the homology of the sub-level sets of the function. Each point corresponds to a feature and quantifies its importance by the absolute difference between the point's two coordinates. Measuring the distance between two functions, f and g, by the L_∞-norm and that between the corresponding persistence diagrams, $\mathrm{D}(f)$ and $\mathrm{D}(g)$, by the bottleneck distance, the stability result bounds the latter distance by the former:

$$d_{\mathrm{B}}(\mathrm{D}(f), \mathrm{D}(g)) \leq \|f - g\|_\infty. \tag{1}$$

The assumptions required for this result are mild and are satisfied by Morse functions on compact manifolds, piecewise linear functions on simplicial complexes, and more. The bottleneck distance is based on a bijection between the points and is therefore always at least the Hausdorff distance between the two diagrams. A result similar to ours in the special case of degree 0 homology has recently been obtained by d'Amico et al. [2]. To treat our more general case, we use diagram-chasing methods from algebraic topology to prove (1) for the Hausdorff distance and approximation by piecewise linear functions together with linear interpolation to strengthen (1) from the Hausdorff to the bottleneck distance.

The authors believe that (1) can be used to shed new light on natural phenomena through improvements of our analysis capabilities. To provide evidence for this claim, we apply the inequality to two specific problems. The first is the estimation of the homology of a closed subset of a metric space from a finite point sample. A few years ago, Robins proved an algebraic tool we call the Quadrant Lemma [19]. We use it to show that under some assumptions on the sampling density, the persistent homology of the point sample, for parameters related to the sampling density, is the same as the homology of the subset. Somewhat surprisingly, this result does not require the full power of our stability result, not even the Hausdorff version, which we prove using the Quadrant Lemma but also the more powerful Box Lemma. The same result on homology estimation has independently been obtained by Chazal and Lieutier [6]. Their methods are limited to subspaces of Euclidean space, but extend beyond homology to fundamental groups. The second problem is the comparison and classification of geometric shapes. Due to its practical importance, it has been studied extensively in a number of areas including morphology [3] and image processing [20]. Recently, Carlsson et al. introduced barcodes, which are persistence diagrams (drawing points in the plane as intervals) for

the curvature function of a certain derived space of the shape [4]. We use (1) to make concrete statements about the stability of these barcodes. As an additional application, not described in this paper, we established that under fairly mild assumptions it is possible to estimate the total mean curvature of a smooth surface in $\mathbb{R}^3$ by the discrete analog of that measure defined for a piecewise linear approximation of the smooth surface [7].

Outline. Section 2 introduces the mathematical concepts used in this paper. Section 3 proves the stability of persistence diagrams, focusing on the Hausdorff distance in Section 3.2 and on the bottleneck distance in Section 3.3. Section 4 presents two applications of the stability result. Section 5 concludes this paper.

2. Background and Definitions

In this section we review background from topology, in particular homology groups, and we introduce the setting for our results.

Homology and Tame Functions. We refer to [17] for an introduction to homology that is both mathematically rigorous and accessible to non-specialists. Given a topological space $\mathbb{X}$ and an integer k, we denote the kth *singular homology group* of $\mathbb{X}$ by $\mathsf{H}_k(\mathbb{X})$, and the kth *Betti number* by $\beta_k(\mathbb{X}) = \dim \mathsf{H}_k(\mathbb{X})$. In this paper we work with modulo 2 coefficients, so that homology groups are vector spaces over $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z}$. We recall that a continuous function f between two topological spaces $\mathbb{X}$ and $\mathbb{Y}$ induces linear maps $f_k\colon \mathsf{H}_k(\mathbb{X}) \to \mathsf{H}_k(\mathbb{Y})$ between the homology groups. Also, if $f\colon \mathbb{X} \to \mathbb{Y}$ and $g\colon \mathbb{Y} \to \mathbb{Z}$ are two continuous functions, then the linear map induced by the composition is the composition of the induced linear maps: $(g \circ f)_k = g_k \circ f_k$. In what follows we only consider the special case in which $\mathbb{X}$ is a subspace of $\mathbb{Y}$ and f is the inclusion of $\mathbb{X}$ into $\mathbb{Y}$.

The results of this paper apply to a fairly general class of functions which we refer to as tame. We begin by extending the classical notion of critical values to real functions on topological spaces, without further restriction.

Definition. Let $\mathbb{X}$ be a topological space and let f be a real function on $\mathbb{X}$. A *homological critical value* of f is a real number a for which there exists an integer k such that for all sufficiently small $\varepsilon > 0$ the map $\mathsf{H}_k(f^{-1}(-\infty, a-\varepsilon]) \to \mathsf{H}_k(f^{-1}(-\infty, a+\varepsilon])$ induced by inclusion is not an isomorphism.

In words, the homological critical values are the levels where the homology of the sublevel sets changes. If f is a Morse function on a smooth manifold, then Morse theory implies that its homological critical values coincide with its classical critical values, i.e., its values at critical points [16]. For (generic) PL functions defined on simplicial complexes, homological critical values form a subset of the function values at the vertices. Both examples are special cases of Morse functions on Whitney-stratified spaces [14], which include a large class of piecewise smooth functions defined on smooth manifolds. For such a function, homological critical values form a subset of the critical values of the restriction of the function to the strata.

Definition. A function $f\colon \mathbb{X} \to \mathbb{R}$ is *tame* if it has a finite number of homological critical values and the homology groups $\mathsf{H}_k(f^{-1}(-\infty, a])$ are finite-dimensional for all $k \in \mathbb{Z}$ and $a \in \mathbb{R}$.

In particular, Morse functions on compact manifolds are tame, as well as PL functions on finite simplicial complexes and, more generally, Morse functions on compact Whitney-stratified spaces. We end this paragraph with an observation about homological critical values. Let $\mathbb{X}$ be a topological space and let $f\colon \mathbb{X} \to \mathbb{R}$. Assuming a fixed integer k, we write $F_x = \mathsf{H}_k(f^{-1}(-\infty, x])$, and for $x < y$, we let $f_x^y\colon F_x \to F_y$ be the map induced by inclusion of the sub-level set of x in that of y.

Critical Value Lemma. *If some closed interval* $[x, y]$ *contains no homological critical value of* f, *then* f_x^y *is an isomorphism for every integer* k.

Proof. Letting $m = (x + y)/2$, we have $f_x^y = f_m^y \circ f_x^m$. If f_x^y is not an isomorphism then at least one of f_x^m and f_m^y is not an isomorphism either. By induction we obtain a decreasing sequence of intervals whose intersection is a homological critical value inside $[x, y]$, contradicting our assumption. □

Persistence Diagrams. Using the same notation as above, we write $F_x^y = \operatorname{im} f_x^y$ for the image of F_x in F_y. By convention, we set $F_x^y = \{0\}$ whenever x or y is infinite. The groups F_x^y, called *persistent homology group* in [11], are key objects in the study of topological persistence. Whereas the groups F_x tell us about the topology of the sub-level sets of f, persistent homology groups contain information about the topological relationships between these sub-level sets.

We now show that the set of all persistent homology groups of a tame function can be encoded in a planar drawing, which we call a persistence diagram. As we will see later, persistence diagrams are but another representation of the k-intervals introduced in [11] and extended in [23]. Let $f\colon \mathbb{X} \to \mathbb{R}$ be a tame function, let $(a_i)_{i=1..n}$ be its homological critical values, and let $(b_i)_{i=0..n}$ be an interleaved sequence, namely, $b_{i-1} < a_i < b_i$ for all i. We set $b_{-1} = a_0 = -\infty$ and $b_{n+1} = a_{n+1} = +\infty$. For two integers $0 \le i < j \le n+1$, we define the *multiplicity* of the pair (a_i, a_j) by

$$\mu_i^j = \beta_{b_{i-1}}^{b_j} - \beta_{b_i}^{b_j} + \beta_{b_i}^{b_{j-1}} - \beta_{b_{i-1}}^{b_{j-1}},$$

where $\beta_x^y = \dim F_x^y$ denote *persistent Betti numbers* for all $-\infty \le x \le y \le +\infty$. To visualize this definition, consider β_x^y as the value of a function β at the point $(x, y) \in \bar{\mathbb{R}}^2$, where $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. Then μ_i^j is the alternating sum of β on the corners of the box $[b_{i-1}, b_i] \times [b_{j-1}, b_j]$, depicted in Fig. 1. Observe that if x and x' lie in the open interval (a_i, a_{i+1}) and y and y' lie in (a_{j-1}, a_j), then $\beta_x^y = \beta_{x'}^{y'}$. Indeed, it follows from the Critical Value Lemma that F_x^y and $F_{x'}^{y'}$ are isomorphic. The multiplicities μ_i^j are thus well-defined and we will see later that they are always non-negative. We now introduce the main object of study.

Definition. The *persistence diagram* $\mathrm{D}(f) \subset \bar{\mathbb{R}}^2$ of f is the set of points (a_i, a_j), counted with multiplicity μ_i^j for $0 \le i < j \le n + 1$, union all points on the diagonal, counted with infinite multiplicity.

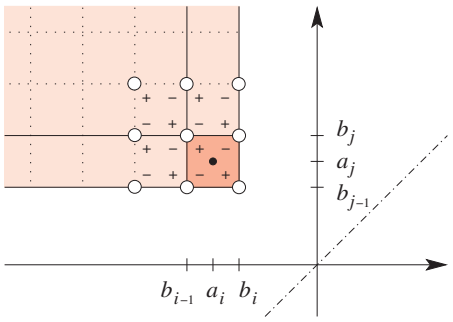

Fig. 1. The multiplicity of the point (a_i, a_j) is the alternating sum of persistent Betti numbers at the corners of the lower right square. When adding other multiplicities, cancellations between plus and minus signs occur.

We write $\sharp(A)$ for the *total multiplicity* of a multiset A which, by definition, is the sum of multiplicities of the elements in A. For example, the total multiplicity of the persistence diagram minus the diagonal is

$$\sharp(\mathrm{D}(f) - \Delta) = \sum_{i<j} \mu_i^j.$$

We call this number the *size* of the persistence diagram. Points with multiplicity zero are not counted and can therefore be discarded.

Basic Properties and Interpretation. By construction, persistence diagrams satisfy the k-Triangle Lemma [11], which we now restate. It is convenient to have short notation for the closed upper left quadrant defined by a point (x, y), $Q_x^y = [-\infty, x] \times [y, \infty]$.

***k*-Triangle Lemma.** *Let f be a tame function and suppose $x < y$ are different from the homological critical values of f. Then the total multiplicity of the persistence diagram within the upper left quadrant is $\sharp(\mathrm{D}(f) \cap Q_x^y) = \beta_x^y$.*

Proof. We may assume without loss of generality that $x = b_i$ and $y = b_{j-1}$. By definition, the total multiplicity in the upper left quadrant is

$$\begin{aligned}
\mu &= \sum_{k \le i \le j \le \ell} \mu_k^\ell \\
&= \sum_{k \le i \le j \le \ell} (\beta_{b_{k-1}}^{b_\ell} - \beta_{b_k}^{b_\ell} + \beta_{b_k}^{b_{\ell-1}} - \beta_{b_{k-1}}^{b_{\ell-1}}) \\
&= \beta_{b_{-1}}^{b_{n+1}} - \beta_{b_i}^{b_{n+1}} + \beta_{b_i}^{b_{j-1}} - \beta_{b_{-1}}^{b_{j-1}}.
\end{aligned}$$

Indeed, all other terms cancel, as indicated in Fig. 1. However, the remaining terms vanish, except for the third, which is equal to β_x^y. □

The fact that persistence diagrams satisfy the k-Triangle Lemma implies that they are equivalent to the pairing defined in [11] for filtrations of simplicial complexes. More

precisely, let $\emptyset = K_0 \subset K_1 \subset \cdots \subset K_m = K$ be a filtration of a simplicial complex K such that K_{i+1} differs from K_i by a single simplex σ_i. The complexes in this filtration are the sub-level sets of the function whose value on the interior of σ_i is i. Here is the connection: the persistence diagram of this function is the set of points whose coordinates are the pairs of indices computed by the persistence algorithm [11], together with points at infinity that correspond to indices left unpaired by the algorithm, together with the diagonal. This means that each off-diagonal point in the persistence diagram can be interpreted as the life-span of a topological feature, as explained in [11]. This fact can also be seen from the definition of multiplicities. We explain this while temporarily simplifying notation to $F_i = F_{b_i}$ and $\beta_i^j = \beta_{b_i}^{b_j}$. The multiplicity can be written as the difference between two differences: $\mu_i^j = (\beta_i^{j-1} - \beta_i^j) - (\beta_{i-1}^{j-1} - \beta_{i-1}^j)$. The first term, β_i^{j-1}, can be interpreted as the number of independent homology classes in F_{j-1} born before F_i. The first difference, $\beta_i^{j-1} - \beta_i^j$, thus counts the classes in F_{j-1} born before F_i that die before F_j. Similarly, the second difference, $\beta_{i-1}^{j-1} - \beta_{i-1}^j$, counts the classes in F_{j-1} born before F_{i-1} that die before F_j. It follows that μ_i^j counts the classes born between F_{i-1} and F_i that die between F_{j-1} and F_j.

For more general functions, persistence diagrams above the diagonal coincide with the (multi)sets of $\mathcal{P}$-intervals described in [23], except that we picture them as points in the extended plane rather than intervals. The advantage of this representation will be obvious. While the size of a persistence diagram can be quadratic in the number of homological critical values in the worst case, it is linear in several important cases. We already mentioned the case of filtrations obtained by adding one simplex at a time. Other examples with linear size persistence diagrams are Morse functions on smooth manifolds, and PL functions on simplicial complexes in which each vertex belongs to at most some constant number of simplices. In these cases, persistence diagrams provide a compact encoding of the persistent homology groups.

3. Stability

In this section we state and prove the main result of this paper. The proof is done in two steps, establishing the result for the Hausdorff distance in Section 3.2 and strengthening it to the bottleneck distance in Section 3.3.

3.1. *Statement of Theorem*

We need some definitions. For points $p = (p_1, p_2)$ and $q = (q_1, q_2)$ in $\bar{\mathbb{R}}^2$, let $\|p - q\|_\infty$ be the maximum of $|p_1 - q_1|$ and $|p_2 - q_2|$. Similarly for functions f and g, let $\|f - g\|_\infty = \sup_x |f(x) - g(x)|$. Let X and Y be multisets of points.

Definition. The *Hausdorff distance* and the *bottleneck distance* between X and Y are

$$d_{\mathrm{H}}(X, Y) = \max\left\{\sup_x \inf_y \|x - y\|_\infty, \sup_y \inf_x \|y - x\|_\infty\right\},$$
$$d_{\mathrm{B}}(X, Y) = \inf_\gamma \sup_x \|x - \gamma(x)\|_\infty,$$

where $x \in X$ and $y \in Y$ range over all points and γ ranges over all bijections from X to Y. Here we interpret each point with multiplicity k as k individual points and the bijection is between the resulting sets.

The prime example of multisets we consider are persistence diagrams. A bijection between two diagrams has three types of point pairs: both off the diagonal, one off the diagonal and the other on the diagonal, and both on the diagonal. The most important type is the first, matching features between the two functions, and the least important is the last, completing the matching in a way that does not affect the bottleneck distance. Since the bottleneck distance satisfies one more constraint, namely a bijection between the points, we have $d_{\mathrm{H}}(X, Y) \leq d_{\mathrm{B}}(X, Y)$. Recalling that a topological space is *triangulable* if there is a (finite) simplicial complex with homeomorphic underlying space, we now state the main result of this paper, which may be referred to as the Bottleneck Stability Theorem for Persistence Diagrams.

Main Theorem. *Let $\mathbb{X}$ be a triangulable space with continuous tame functions $f, g\colon \mathbb{X} \to \mathbb{R}$. Then the persistence diagrams satisfy $d_{\mathrm{B}}(\mathrm{D}(f), \mathrm{D}(g)) \leq \|f - g\|_\infty$.*

In words, persistence diagrams are stable under possibly irregular perturbations of small amplitude. This is illustrated in Fig. 2 where the surplus critical values of one function define points of the persistence diagram near the diagonal. As shown by Zomorodian and Carlsson [23] in a different language, persistence diagrams completely describe the homology groups of sub-level sets of a function and the maps induced by inclusion between them, up to isomorphism. They are thus a detailed representation of the topological features of a function that is stable and, in fact, Lipschitz. Moreover, this representation is meaningful, since each point in the persistence diagram of a function corresponds to a topological event in the filtration associated with that function.

The bottleneck distance between two persistence diagrams can be computed by adapting standard maximum matching algorithms for bipartite graphs; see [8, Chapter 26] or [15]. Since the bottleneck distance is bounded from below by the Hausdorff distance, the claim in the Main Theorem is also true for the Hausdorff distance, which is easier to compute. Indeed, we just need to find the smallest ε such that squares of side-length 2ε placed with their centers at the points of one diagram cover all off-diagonal points

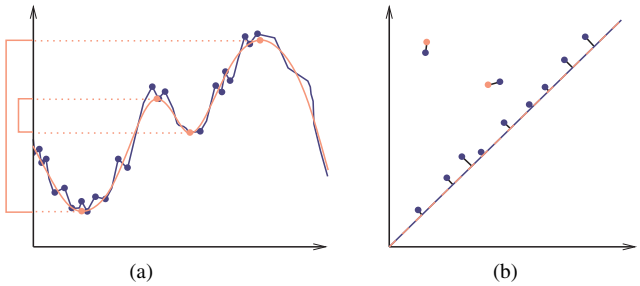

Fig. 2. (a) Two close functions, one with many and the other with just four critical values. (b) The persistence diagrams of the two functions, and the bijection between them.

of the other diagram, and vice versa with the diagrams exchanged. We note also that the stability of critical value pairs is in sharp contrast to the lack of stability of critical values and, for Morse and PL functions, critical points and critical point pairs. Critical values are destroyed by cancellations and created by their inverses. Also, the location of critical points is unstable in regions where the function is nearly constant. Even when critical points happen to be stable, pairs of critical points change when critical values go through interchanges. These changes prohibit any stability results for all three concepts.

3.2. *Proof of Hausdorff Stability*

We state and prove a preliminary result, the Box Lemma, which implies the stability of persistence diagrams for the Hausdorff distance. This result will be used in Section 3.3 to prove the stronger statement that persistence diagrams are stable for the bottleneck distance.

Relations between Quadrants. Let f and g be two tame functions defined on a topological space $\mathbb{X}$. For all $x \in \mathbb{R}$, we let $F_x = \mathsf{H}_k(f^{-1}(-\infty, x])$ and $G_x = \mathsf{H}_k(g^{-1}(-\infty, x])$. Also, for all $x < y$, we denote by $f_x^y\colon F_x \to F_y$ and $g_x^y\colon G_x \to G_y$ the maps induced by inclusions, and by $F_x^y = \operatorname{im} f_x^y$ and $G_x^y = \operatorname{im} g_x^y$ the corresponding persistent homology groups. Writing $\varepsilon = \|f - g\|_\infty$, we have $f^{-1}(-\infty, x] \subseteq g^{-1}(-\infty, x + \varepsilon]$ for all $x \in \mathbb{R}$. We denote the map induced by this inclusion by $\varphi_x\colon F_x \to G_{x+\varepsilon}$. The symmetric inclusion in which f and g are exchanged induces another map $\psi_x\colon G_x \to F_{x+\varepsilon}$. Given $b < c$, the maps described above fit into the following two diagrams:

$$\begin{array}{ccc} F_{b-\varepsilon} & \xrightarrow{f_{b-\varepsilon}^{c+\varepsilon}} & F_{c+\varepsilon} \\ {\scriptstyle\varphi_{b-\varepsilon}}\downarrow & & \uparrow{\scriptstyle\psi_c} \\ G_b & \xrightarrow{g_b^c} & G_c \end{array} \qquad \begin{array}{ccc} F_{b+\varepsilon} & \xrightarrow{f_{b+\varepsilon}^{c+\varepsilon}} & F_{c+\varepsilon} \\ {\scriptstyle\psi_b}\uparrow & & \uparrow{\scriptstyle\psi_c} \\ G_b & \xrightarrow{g_b^c} & G_b^c \end{array}$$

Since the inclusion maps commute, so do the induced maps. Considering the first diagram, we get $f_{b-\varepsilon}^{c+\varepsilon} = \psi_c \circ g_b^c \circ \varphi_{b-\varepsilon}$. Let now $\xi \in F_{b-\varepsilon}^{c+\varepsilon}$. By definition, $\xi = f_{b-\varepsilon}^{c+\varepsilon}(\eta)$ for some $\eta \in F_{b-\varepsilon}$. Hence $\xi = \psi_c(\zeta)$, with $\zeta = g_b^c(\varphi_{b-\varepsilon}(\eta)) \in G_b^c$. It follows that $F_{b-\varepsilon}^{c+\varepsilon}$ is a subset of the image of G_b^c under ψ_c. Considering the second diagram, we see that $\psi_c(G_b^c)$ equals $\psi_c \circ g_b^c(G_b)$, which in turn equals $f_{b+\varepsilon}^{c+\varepsilon} \circ \psi_b(G_b) \subseteq F_{b+\varepsilon}^{c+\varepsilon}$. We state these two findings for later reference:

$$F_{b-\varepsilon}^{c+\varepsilon} \subseteq \psi_c(G_b^c) \subseteq F_{b+\varepsilon}^{c+\varepsilon}. \tag{2}$$

The first inclusion implies $\dim F_{b-\varepsilon}^{c+\varepsilon} \leq \dim G_b^c$, which is a result that already appears in [19]. Applying the k-Triangle Lemma, we get a first inequality between accumulated multiplicities within the two persistence diagrams. To synchronize the statement of the inequality with that of the next, we let $Q = Q_b^c$ and $Q_\varepsilon = Q_{b-\varepsilon}^{c+\varepsilon}$.

Quadrant Lemma. $\sharp(\mathrm{D}(f) \cap Q_\varepsilon) \leq \sharp(\mathrm{D}(g) \cap Q)$.

In words, the total multiplicity of $\mathrm{D}(g)$ inside the upper left quadrant with corner (b, c) is bounded from below by the total multiplicity of $\mathrm{D}(f)$ inside the quadrant shrunk by ε. Of course, the inequality is symmetric in f and g. Strictly speaking, the above discussion proves the claimed inequality only for the case when b, c are not homological critical values of g and $b - \varepsilon, c + \varepsilon$ are not homological critical values of f. However, if they are then we can enlarge the quadrants with a sufficiently small real number $0 < \delta < \varepsilon$ such that

$$\begin{aligned} \sharp(\mathrm{D}(f) \cap Q_\varepsilon) &= \sharp(\mathrm{D}(f) \cap Q_{b-\varepsilon+\delta}^{c+\varepsilon-\delta}), \\ \sharp(\mathrm{D}(g) \cap Q) &= \sharp(\mathrm{D}(g) \cap Q_{b+\delta}^{c-\delta}), \end{aligned}$$

and the above argument applies directly because the modified coordinates are not homological critical values.

Images, Kernels, and Quotients. The Quadrant Lemma is too weak for our purposes. To prepare a similar result for nested boxes, we introduce vector spaces that correspond to rectangular regions in $\bar{\mathbb{R}}^2$ defined by up to four constraints. Using the k-Triangle Lemma, we express the dimensions of these vector spaces by the total multiplicities of the corresponding regions. Let $w < x < y < z$ be four real numbers, all different from homological critical values of $f\colon \mathbb{X} \to \mathbb{R}$. We recall that the dimension of the homology group F_x is the total multiplicity of the upper left quadrant with corner (x, x) (not including the corner itself), and the dimension of the persistent homology group F_x^y is the total multiplicity of the upper left quadrant with corner (x, y); see Fig. 3(a),(b). Restricting $f_y^z\colon F_y \to F_z$ to the vector space F_x^y gives a surjection $f_x^{y,z}\colon F_x^y \to F_x^z$. Writing $F_x^{y,z}$ for the kernel of this map, we have $\dim F_x^{y,z} = \dim F_x^y - \dim F_x^z$ [22,

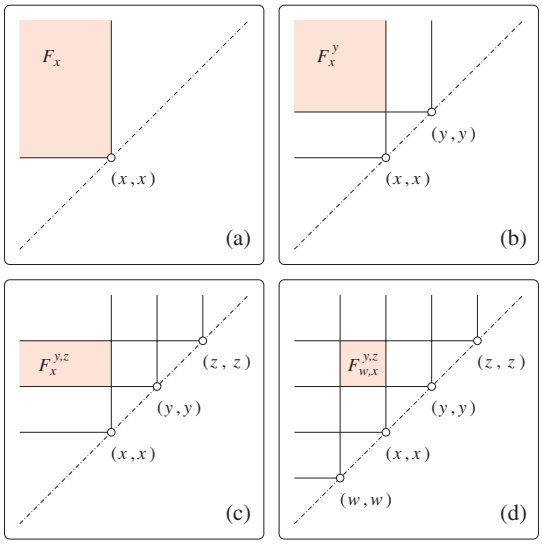

Fig. 3. (a) Homology group of the sub-level set $f^{-1}(-\infty, x]$. (b) Image of F_x in F_y. (c) Kernel of surjection $F_x^y \to F_x^z$. (d) Quotient of $F_x^{y,z}$ and $F_w^{y,z}$.

Chapter 3]. This is the total multiplicity of the shaded three-sided rectangle depicted in Fig. 3(c). Also, $F^y_w \subseteq F^y_x$ since any element of F^y_w, being the image of some element $\xi \in F_w$ by f^y_w, is also the image of $f^x_w(\xi)$ by f^y_x. Thus the map $f^{y,z}_w$ is just the restriction of $f^{y,z}_x$: $F^y_x \to F^z_x$ to F^y_w. As a consequence, the kernel $F^{y,z}_w$ of the former map is included in the kernel $F^{y,z}_x$ of the latter map. We can therefore consider the quotient space $F^{y,z}_{w,x} = F^{y,z}_x / F^{y,z}_w$. Its dimension is the difference between the dimensions of the two kernels, $\dim F^{y,z}_{w,x} = \dim F^{y,z}_x - \dim F^{y,z}_w$. Equivalently, it is the total multiplicity of the shaded rectangular box $[w, x] \times [y, z]$ depicted in Fig. 3(d).

An Inequality for Nested Boxes. We use the above definitions to prove a crucial improvement of the Quadrant Lemma. For $a < b < c < d$, let $R = [a, b] \times [c, d]$ be a box in $\bar{\mathbb{R}}^2$ and let $R_\varepsilon = [a + \varepsilon, b - \varepsilon] \times [c + \varepsilon, d - \varepsilon]$ be the box obtained by shrinking R at all four sides.

Box Lemma. $\sharp(\mathrm{D}(f) \cap R_\varepsilon) \leq \sharp(\mathrm{D}(g) \cap R)$.

Proof. As explained above, we may assume that a, b, c, d are not homological critical values of g and $a + \varepsilon, b - \varepsilon, c + \varepsilon, d - \varepsilon$ are not homological critical values of f. Furthermore, we may assume $a + \varepsilon < b - \varepsilon$ and $c + \varepsilon < d - \varepsilon$, else there is nothing to show. We approach the inequality by interpreting the total multiplicity of a persistence diagram within a box as the dimension of a vector space, as explained in the previous paragraph. More precisely, consider the vector spaces whose dimensions give the total multiplicities within the two boxes

$$\dim F^{c+\varepsilon,d-\varepsilon}_{a+\varepsilon,b-\varepsilon} = \sharp(\mathrm{D}(f) \cap R_\varepsilon), \tag{3}$$

$$\dim G^{c,d}_{a,b} = \sharp(\mathrm{D}(g) \cap R). \tag{4}$$

We prove the claimed inequality by finding a surjection from a subspace of the latter vector space to the former. The main tool used to relate these vector spaces is the commutative diagram shown in Fig. 4, which has a vector space for each corner of the two boxes.

To define the relevant subspace of $G^{c,d}_{a,b}$, we introduce subspaces of G^c_b and G^c_a. First, we let E^c_b be the preimage, by the restriction of ψ_c to G^c_b, of the kernel of u_3 (see Fig. 4), that is, $E^c_b = \psi_c^{-1}(F^{c+\varepsilon,d-\varepsilon}_{b-\varepsilon}) \cap G^c_b$. Note that by (2), the image of G^c_b under ψ_c contains $F^{c+\varepsilon}_{b-\varepsilon}$, so the restriction of ψ_c to E^c_b, which we denote by s_3, has the kernel of u_3 as its image. We also consider the intersection $E^c_a = G^c_a \cap E^c_b$. We will see below that E^c_b / E^c_a is a subspace of $G^{c,d}_{a,b}$, from which a surjection to $F^{c+\varepsilon,d-\varepsilon}_{a+\varepsilon,b-\varepsilon}$ can be constructed. Continuing with the description of the diagram, the maps r_1, r_2, r_3 and r_4 are just inclusions between vector spaces. Furthermore, u_1 is the restriction of $g^{c,d}_a$ to E^c_a and u_4 is the restriction of $g^{c,d}_b$ to E^c_b. The map s_2 is the restriction of ψ_c to E^c_a, and we get $\psi_c(G^c_a) \subseteq F^{c+\varepsilon}_{a+\varepsilon}$ from (2), which implies that the image of s_2 is contained in the same vector space, as required. Finally, s_1 is the restriction of $\varphi_{d-\varepsilon}$ to $F^{d-\varepsilon}_{b-\varepsilon}$, and we get $\varphi_{d-\varepsilon}(F^{d-\varepsilon}_{b-\varepsilon}) \subseteq G^d_b$ from (2) (with F and G interchanged), which implies that the image of s_1 is contained in G^d_b, as required. The diagram in Fig. 4 is therefore valid and it obviously commutes. Hence, $u_4 = s_1 \circ u_3 \circ s_3$, which implies $E^c_b = \ker u_4$ because

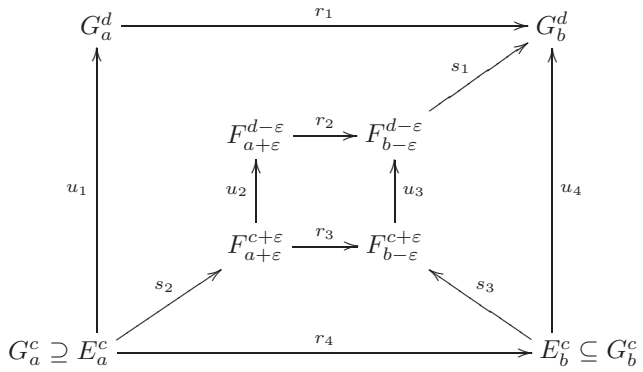

Fig. 4. Commutative diagram patterned after the two nested boxes $R_\varepsilon \subseteq R$. We shorten notation by writing $u_2 = f_{a+\varepsilon}^{c+\varepsilon,d-\varepsilon}$ and $u_3 = f_{b-\varepsilon}^{c+\varepsilon,d-\varepsilon}$.

$u_3 \circ s_3$ is zero. Furthermore, $r_1 \circ u_1 = u_4 \circ r_4$, which implies $E_a^c = \ker u_1$ because $u_4 \circ r_4$ is zero and r_1 is an inclusion. We express these relations with redundant notation, writing $E_b^c = E_b^{c,d} \subseteq G_b^{c,d}$ and $E_a^c = E_a^{c,d} \subseteq G_a^{c,d}$. Since $E_a^{c,d} = E_b^{c,d} \cap G_a^{c,d}$, the quotient $E_{a,b}^{c,d} = E_b^{c,d}/E_a^{c,d}$ is just the set of cosets of elements in $E_b^{c,d} \subseteq G_b^{c,d}$ modulo $G_a^{c,d}$, so $E_{a,b}^{c,d} \subseteq G_{a,b}^{c,d}$. In particular,

$$\dim E_{a,b}^{c,d} \leq \dim G_{a,b}^{c,d}. \tag{5}$$

We are now ready for the final argument relating the two quotients. Recall that $E_{a,b}^{c,d} = \ker u_4 / \ker u_1$ and consider $F_{a+\varepsilon,b-\varepsilon}^{c+\varepsilon,d-\varepsilon} = \ker u_3 / \ker u_2$. By construction, $s_3(\ker u_4) = \ker u_3$. To show that s_3 induces a surjection between the quotients, it thus remains to prove that $s_3(\ker u_1) = s_2(\ker u_1)$ is included in $\ker u_2$. However, this is clear because $r_2 \circ u_2 \circ s_2(\xi) = u_3 \circ s_3 \circ r_4(\xi) = 0$, for every $\xi \in \ker u_1$, and r_2 is an injection. As a consequence,

$$\dim F_{a+\varepsilon,b-\varepsilon}^{c+\varepsilon,d-\varepsilon} \leq \dim E_{a,b}^{c,d}. \tag{6}$$

We get the claimed inequality by concatenating (3), (6), (5), (4), in this sequence. □

A direct consequence of the Box Lemma is that the Hausdorff distance between $\mathrm{D}(f)$ and $\mathrm{D}(g)$ is not larger than ε. Indeed, if (x, y) is a point of $\mathrm{D}(f)$, then there must be a point of $\mathrm{D}(g)$ at distance less than or equal to ε from (x, y) since the total multiplicity of $\mathrm{D}(g)$ inside the square $[x - \varepsilon, x + \varepsilon] \times [y - \varepsilon, y + \varepsilon]$ is at least one.

3.3. *Proof of Bottleneck Stability*

The Hausdorff distance between two persistence diagrams never exceeds the bottleneck distance because it is oblivious to multiplicities and clusters of points. In this subsection we strengthen the stability result to the bottleneck distance, thus completing the proof of the Main Theorem. This strengthening is crucial for some of the applications, including the inequalities proved in [7].

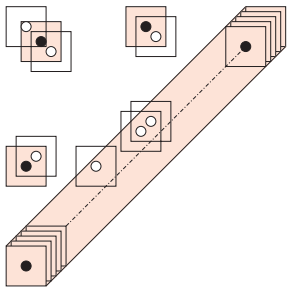

Fig. 5. The shaded squares are centered at the (black) points of $\mathrm{D}(f)$. The white squares are centered at the (white) points of $\mathrm{D}(g)$.

An Easy Special Case. Before proving the stability for bottleneck distance in the general case, we discuss a special case that permits an easy proof. Given a tame function $f\colon \mathbb{X} \to \mathbb{R}$, we consider the minimum distance between two different off-diagonal points or between an off-diagonal point and the diagonal:

$$\delta_f = \min\{\|p - q\|_\infty \mid \mathrm{D}(f) - \Delta \ni p \neq q \in \mathrm{D}(f)\}.$$

If we draw squares of radius $\varepsilon = \delta_f/2$ around the points of $\mathrm{D}(f)$ we get a thickened diagonal and a finite collection of squares that are disjoint from each other and from the thickened diagonal; see Fig. 5. We call another tame function $g\colon \mathbb{X} \to \mathbb{R}$ *very close* to f if $\|f - g\|_\infty < \delta_f/2$. We now prove the Main Theorem under the additional assumption of very close functions.

Easy Bijection Lemma. *Let $f, g\colon \mathbb{X} \to \mathbb{R}$ be tame functions and let g be very close to f. Then the persistence diagrams satisfy $d_{\mathrm{B}}(\mathrm{D}(f), \mathrm{D}(g)) \leq \|f - g\|_\infty$.*

Proof. Writing μ for the multiplicity of the point p in $\mathrm{D}(f) - \Delta$ and $\square_\varepsilon$ for the square with center p and radius $\varepsilon = \|f - g\|_\infty$, we get

$$\mu \leq \sharp(\mathrm{D}(g) \cap \square_\varepsilon) \leq \sharp(\mathrm{D}(f) \cap \square_{2\varepsilon})$$

from the Box Lemma. Since $2\varepsilon < \delta_f$, p is the only point of $\mathrm{D}(f)$ in $\square_{2\varepsilon}$, which implies $\sharp(\mathrm{D}(g) \cap \square_\varepsilon) = \mu$. We can therefore map all points of $\mathrm{D}(g) \cap \square_\varepsilon$ to p. After repeating this step for all off-diagonal points of $\mathrm{D}(f)$, the only points of $\mathrm{D}(g)$ that remain without image have distance more than ε from $\mathrm{D}(f) - \Delta$. Because the Hausdorff distance between $\mathrm{D}(f)$ and $\mathrm{D}(g)$ is at most ε, these points of $\mathrm{D}(g)$ are at distance at most ε from the diagonal. Mapping them to their respective closest points on Δ yields a bijection between the multisets $\mathrm{D}(f)$ and $\mathrm{D}(g)$, keeping in mind that the points on Δ have infinite multiplicity. Since the bijection moves points by at most ε, this concludes the proof. □

We prove the Main Theorem by composing many bijections of the type described above, thus constructing a bijection for the general case.

The Case of Piecewise Linear Functions. We now prove the Main Theorem for two piecewise linear functions $\hat{f}$ and $\hat{g}$ defined on a simplicial complex K. A *convex combination* of $\hat{f}$ and $\hat{g}$ is a function $h_\lambda = (1-\lambda)\hat{f} + \lambda\hat{g}$ for which $\lambda \in [0,1]$. The one-parameter family of convex combinations forms a linear interpolation between the two piecewise linear functions, starting at $h_0 = \hat{f}$ and ending at $h_1 = \hat{g}$.

Interpolation Lemma.

$$d_{\mathrm{B}}(\mathrm{D}(\hat{f}), \mathrm{D}(\hat{g})) \leq \|\hat{f} - \hat{g}\|_\infty.$$

Proof. We decompose the linear interpolation into sufficiently small steps so we can use the Easy Bijection Lemma to get a bijection for each step. Let $c = \|\hat{f} - \hat{g}\|_\infty$ and note that for each $\lambda \in [0,1]$, h_λ is tame and $\delta(\lambda) = \delta_{h_\lambda}$ is positive. It follows that the set C of open intervals $J_\lambda = (\lambda - \delta(\lambda)/4c, \lambda + \delta(\lambda)/4c)$ forms an open cover of the interval $[0,1]$. Consider now a minimal subcover C' of C. Since $[0,1]$ is compact, C' is finite. Let $\lambda_1 < \lambda_2 < \cdots < \lambda_n$ be the midpoints of the intervals in C'. Since C' is minimal, any two consecutive intervals, J_{λ_i} and $J_{\lambda_{i+1}}$, have a non-empty intersection. Hence,

$$\begin{aligned}\lambda_{i+1} - \lambda_i &\leq (\delta(\lambda_i) + \delta(\lambda_{i+1}))/4c \\ &\leq \max\{\delta(\lambda_i), \delta(\lambda_{i+1})\}/2c.\end{aligned}$$

By definition of c, $\|h_{\lambda_i} - h_{\lambda_{i+1}}\|_\infty = c(\lambda_{i+1} - \lambda_i)$. As a consequence, $\|h_{\lambda_i} - h_{\lambda_{i+1}}\|_\infty \leq \max\{\delta(\lambda_i), \delta(\lambda_{i+1})\}/2$, which implies that h_{λ_i} is very close to $h_{\lambda_{i+1}}$ or the other way around. We can thus apply the Easy Bijection Lemma, which yields that the bottleneck distance between $\mathrm{D}(h_{\lambda_i})$ and $\mathrm{D}(h_{\lambda_{i+1}})$ is bounded from above by $\|h_{\lambda_i} - h_{\lambda_{i+1}}\|_\infty$ for $1 \leq i \leq n-1$. Putting $\lambda_0 = 0$ and $\lambda_{n+1} = 1$, we see that the previous inequality holds also for $i = 0$ and for $i = n$ because h_0 is very close to h_{λ_1} and h_1 is very close to h_{λ_n}. Using the triangle inequality, we get

$$\begin{aligned}d_{\mathrm{B}}(\mathrm{D}(\hat{f}), \mathrm{D}(\hat{g})) &\leq \sum_{i=0}^{n} d_{\mathrm{B}}(\mathrm{D}(h_{\lambda_i}), \mathrm{D}(h_{\lambda_{i+1}})) \\ &\leq \sum_{i=0}^{n} \|h_{\lambda_i} - h_{\lambda_{i+1}}\|_\infty.\end{aligned}$$

However, since the h_{λ_i} sample the linear interpolation from $\hat{f}$ to $\hat{g}$, the latter sum equals $\|\hat{f} - \hat{g}\|_\infty$, which concludes the proof. □

Finale. We are now ready to combine the accumulated technical results to complete the proof of the Main Theorem. Recall that we assume a triangulable topological space $\mathbb{X}$ and two continuous tame functions $f, g\colon \mathbb{X} \to \mathbb{R}$. By definition of triangulability, there is a (finite) simplicial complex L and a homeomorphism $\Phi\colon L \to \mathbb{X}$. We note that the persistence diagram is invariant under this change of variables, that is, $f \circ \Phi\colon L \to \mathbb{R}$ is tame and has the same persistence diagram as f. Let $\delta > 0$ be sufficiently small. Since f and g are continuous and L is compact, there exists a subdivision K of L such that

$$\begin{aligned}|f \circ \Phi(u) - f \circ \Phi(v)| &\leq \delta, \\ |g \circ \Phi(u) - g \circ \Phi(v)| &\leq \delta,\end{aligned}$$

whenever u and v are points of a common simplex in K. Let now $\hat{f}, \hat{g}\colon \operatorname{Sd} K \to \mathbb{R}$ be the piecewise linear interpolations of $f \circ \Phi$ and $g \circ \Phi$ on K. By construction of K, these functions satisfy $\|\hat{f} - f \circ \Phi\|_\infty \leq \delta$ and $\|\hat{g} - g \circ \Phi\|_\infty \leq \delta$.

We finish the argument using the triangle inequality to bound $d_{\mathrm{B}}(\mathrm{D}(f), \mathrm{D}(g))$ from above by the sum of the bottleneck distances between the persistence diagrams of adjacent functions in the sequence $f, \hat{f}, \hat{g}, g$. For the middle pair we get

$$\begin{aligned} d_{\mathrm{B}}(\mathrm{D}(\hat{f}), \mathrm{D}(\hat{g})) &\leq \|\hat{f} - \hat{g}\|_\infty \\ &\leq \|f - g\|_\infty + 2\delta \end{aligned}$$

using the Interpolation Lemma, the fact that $\hat{f}$ and $\hat{g}$ differ by at most δ from $f \circ \Phi$ and $g \circ \Phi$, and $\|f - g\|_\infty = \|f \circ \Phi - g \circ \Phi\|_\infty$, in this order. To derive a bound for the first pair in the sequence we assume $\delta < \delta_f/2$ so we get a bijection from the Easy Bijection Lemma. Since the change of variables does not affect the persistence diagram, we get

$$d_{\mathrm{B}}(\mathrm{D}(f), \mathrm{D}(\hat{f})) = d_{\mathrm{B}}(\mathrm{D}(f \circ \Phi), \mathrm{D}(\hat{f})) \leq \delta.$$

Similarly, we get δ as an upper bound for the third pair assuming δ is smaller than $\delta_g/2$. In total we have

$$d_{\mathrm{B}}(\mathrm{D}(f), \mathrm{D}(g)) \leq \|f - g\|_\infty + 4\delta.$$

However, this is true for every positive δ, which we can make as small as we like. The inequality therefore holds also without the term 4δ, which is the claimed inequality in the Main Theorem.

4. Applications

By applying our results to different functions, we get several corollaries, some of which we now describe.

Homology from Point Samples. We first address a problem studied by Robins [19] and de Silva and Carlsson [9], namely, estimating the homology groups of a closed subset X of a metric space M from a set of possibly inaccurate point samples. For smooth surfaces embedded in $\mathbb{R}^3$, this can be done by applying a surface reconstruction algorithm and returning the homology groups of the output. A subset of the available algorithms guarantee correct reconstruction, e.g., [1] and [10], implying correct Betti numbers under some assumptions on the input. In the smooth but possibly higher-dimensional setting, Niyogi et al. [18] show how to build a homotopy equivalent complex, which suffices for homology estimation. This approach does not extend to singular spaces, for which provably correct reconstruction algorithms are currently not available. We need definitions to describe our approach to the problem. Let $d^X\colon M \to \mathbb{R}$ be defined by mapping each point $p \in M$ to its distance from X.

Definition. The *homological feature size* of X, denoted by $\operatorname{hfs} X$, is the smallest positive homological critical value of d^X.

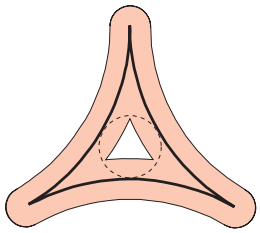

Fig. 6. The bold curve has a positive homological feature size, equal to the radius of the dotted circle, but it has zero reach because the three cusp points have a vanishing local feature size.

When the ambient metric space is the Euclidean space, the homological feature size of X is closely related to the weak feature size of its complement, $\text{wfs}\,(M - X)$, recently introduced in [6]. In particular, results in [6] imply $\text{wfs}\,(M - X) \leq \text{hfs}\,X$. For a surface $S \subseteq \mathbb{R}^3$, both concepts are related to the *local feature size*, $\text{lfs}\colon S \to \mathbb{R}$, defined by mapping each point p to its distance from the medial axis [1]. The minimum of lfs(p), over all points $p \in S$, is sometimes referred to as the *reach* of S. Classical results on parallel bodies imply that $\text{hfs}\,S$ is at least as large as the reach. A non-smooth object can well have a non-zero homological feature size but its reach is necessarily zero; see Fig. 6. For instance, semi-algebraic sets always have a positive weak feature size [13] and therefore a positive homological feature size. This property turns out to be essential in our approach. Suppose we estimate the homology of X from another closed subset P approximating X, which may be a finite set of points. For any two numbers $x < y$, let X_x^y and P_x^y be the persistent kth homology groups of d^X and d^P associated with x and y. To state our result, let $X^{+\delta}$ be the *parallel body* consisting of all points in M at distance less than δ from X.

Homology Inference Theorem. *For all real numbers ε with $d_{\text{H}}(X, P) < \varepsilon < \text{hfs}\,X/4$ and all sufficiently small $\delta > 0$, the dimensions of the homology group of $X^{+\delta}$ and $P_\varepsilon^{3\varepsilon}$ are either both infinite or both finite and equal.*

Proof. Note that $\|d^X - d^P\|_\infty = d_{\text{H}}(X, P)$, by definition of the Hausdorff distance. Hence, $\|d^X - d^P\|_\infty < \varepsilon$. Our assumptions do not imply that d^X and d^P are tame, but we can still apply the inequality between persistent Betti numbers implied by the first inclusion in (2), whose proof makes no use of the tameness assumption. This yields

$$\dim X_\delta^{4\varepsilon+\delta} \leq \dim P_{\varepsilon+\delta}^{3\varepsilon+\delta} \leq \dim X_{2\varepsilon+\delta}^{2\varepsilon+\delta}.$$

Choosing δ such that $4\varepsilon + \delta < \text{hfs}\,X$, the interval $[\delta, 4\varepsilon + \delta]$ contains no homological critical value of d^X. It follows that $X_\delta^{4\varepsilon+\delta}$ and $X_{2\varepsilon+\delta}^{2\varepsilon+\delta}$ have the same dimension implying that both inequalities above are equalities. Furthermore, $\dim \mathsf{H}_k(X^{+\delta}) = \dim X_\delta^\delta = \dim X_\delta^{4\varepsilon+\delta}$, again because there are no homological critical values in $[\delta, 4\varepsilon + \delta]$. □

Perhaps unexpectedly, the homology groups of $X^{+\delta}$ can be different from those of X, even when X has a positive homological feature size and δ is arbitrarily small. An example of such an X is described in [6] and [21]. However, this kind of pathological

behavior cannot happen for absolute neighborhood retracts [21, Chapter 1], which include most practically encountered sets. We note that the dimensions of the homology groups of $X^{+\delta}$ are the dimensions of the Čech cohomology vector spaces in algebraic topology [21, Chapter 6]. From a practical point of view, the Homology Inference Theorem gives an algorithm for estimating the homology of a closed subset of $\mathbb{R}^3$ from a set of samples, provided estimates of the Hausdorff distance and homological feature size are known. Indeed, the dimensions of the persistent homology groups that appear in the theorem are efficiently computable using the persistence algorithm on the filtration of alpha complexes [11], [12]. Finally we mention that Chazal and Lieutier [6] have independently obtained similar but finer results for a more restricted setting. For instance, they can determine the fundamental group of a closed set in Euclidean space from an approximating closed set. This group captures topological information not captured by the homology groups of the set.

Stable Signatures of Shapes. To decide whether two shapes are similar is useful in a variety of settings, including drug design, face recognition, forensic comparison, and sourcing of standard components. In all these applications we need not only fast comparison algorithms but also fast search methods in shape databases. One approach to this problem is to associate with each shape a simpler object, or *signature*, in such a way that two congruent shapes have the same signature. Two shapes can then be compared by comparing their signatures. The term signature originates in the work of Cartan, whose idea was to associate to each curve $s \mapsto f(s)$ in the plane the curve $s \mapsto (\kappa(s), \kappa'(s))$, where $\kappa(s)$ denotes the curvature at the point $f(s)$. This new curve is invariant under rigid motion and contains enough information to reconstruct the original curve, up to rigid motion. However, due to its extreme sensitivity to noise, this particular signature is difficult to use in practice.

In sharp contrast, the persistence diagram of the distance function is a signature that is stable under perturbations in the Hausdorff sense, which is fairly weak. Indeed, as mentioned above, the maximum difference between the distance functions of two shapes is equal to the Hausdorff distance between the shapes. The Main Theorem therefore implies that if two shapes have the Hausdorff distance smaller than ε, up to rigid motion, the bottleneck distance between the persistence diagrams of their distance functions is smaller than ε. In contrast to Cartan's signature, the persistence diagram handles point sets in a reliable way since sufficiently dense samplings of a shape are good approximations for the Hausdorff distance. On the other hand, it is not detailed enough to determine the original shape, not even up to rigid motion. In particular, two shapes whose persistence diagrams are close are not necessarily approximately congruent.

To fill this gap, we may consider persistence diagrams of functions different from distance. An example is the approach introduced in [4]. In the case of a smooth hypersurface, S, it maps each point $p \in S$ and each unit length vector v in the tangent space of S at p to the absolute normal curvature at p in the direction v, $(p, v) \mapsto |\kappa(p, v)|$. The persistence diagram of this function defined on the unit tangent bundle of S, provides a signature of S. It is called a *barcode* in [4], where persistence diagrams are represented as sets of intervals. In order to compare two barcodes B and B', it is suggested to use the following metric, which we describe in terms of the corresponding persistence diagrams

D and D':

$$d(D, D') = \inf_{\gamma} \sum_{p} \|p - \gamma(p)\|_1,$$

where $\|p - q\|_1 = |p_1 - q_1| + |p_2 - q_2|$ and γ ranges over all bijections between D and D'. Although this distance function may be relevant in practice, it is unclear whether it makes barcodes stable under reasonable metrics for the class of hypersurfaces. However, if the distance between barcodes is measured by the bottleneck distance, then the Main Theorem implies that barcodes are continuous when the set of hypersurfaces is endowed with the C^2-topology.

Barcode Theorem. *If a diffeomorphism $\Phi\colon \mathbb{R}^d \to \mathbb{R}^d$ has derivatives up to second order close to those of a rigid motion, then the bottleneck distance between the barcodes of a hypersurface $S \subseteq \mathbb{R}^d$ and its perturbation $\Phi(S)$ is small.*

Indeed, applying a diffeomorphism that is C^2-close to a rigid motion to a hypersurface does not change its normal curvatures by much. A similar result can be obtained for arbitrary co-dimension submanifolds. Computing the bottleneck distance between the barcodes of two shapes is thus a relevant approach to deciding whether these two shapes are approximately congruent.

5. Discussion

We have shown that persistence diagrams provide a stable representation of the topological properties of a broad class of functions, and have described several applications of this result. Many questions remain unanswered, some of which probe the limits of the result. Is it possible to extend the Main Theorem to functions that are non-tame? The ideas used in the proof suggest the appropriate generalization of the persistence diagram of such a function is a measure defined over the extended plane. Are there other interesting metrics, on functions and persistence diagrams, for which an inequality like the one in the Main Theorem holds? For example, if the distance between persistence diagrams is measured as in [4], what metrics between functions give such an inequality, if any? Other open questions are concerned with applications of the Main Theorem, which may be viewed as a precise statement about the structural sensitivity of functions to noise. This interpretation ought to be useful in the design of algorithms that cope well with incomplete or imprecise data. How can we exploit the new insights to get improved algorithms for problems like the medial axis [5], surface reconstruction [1], [10], and alike? How far can we relax the requirements on the input data and still guarantee correct construction?

Acknowledgments

We thank Frédéric Chazal for suggesting a simplification in the proof of the bottleneck stability. We also thank anonymous reviewers for their useful comments.

References

1. N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* **22** (1999), 481–504.
2. M. d'Amico, P. Frosini, and C. Landi. Optimal matching between reduced size functions. Tech. Report No. 35, DISMI, Università di Modena e Reggio Emilia, 2003.
3. F. L. Bookstein. *Morphometric Tools for Landmark Data*. Cambridge University Press, Cambridge, 1999.
4. G. Carlsson, A. Collins, L. Guibas, and A. Zomorodian. Persistence barcodes for shapes. In *Proc. 2nd Sympos. Geometry Process.*, pp. 127–138, 2004.
5. F. Chazal and A. Lieutier. The lambda medial axis. *Graphical Models* **67** (2005), 304–331.
6. F. Chazal and A. Lieutier. Weak feature size and persistent homology: computing homology of solids in $\mathbb{R}^n$ from noisy point samples. In *Proc. 21st Ann. Sympos. Comput. Geom.*, pp. 255–262, 2005.
7. D. Cohen-Steiner and H. Edelsbrunner. Inequalities for the curvature of curves and surfaces. In *Proc. 21st Ann. Sympos. Comput. Geom.*, pp. 272–277, 2005.
8. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*, second edition. MIT Press, Cambridge, Massachusetts, 2001.
9. V. de Silva and G. Carlsson. Topological estimation using witness complexes. In *Proc. Sympos. Point-Based Graphics*, pp. 157–166, 2004.
10. T. K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. In *Proc. 20th Ann. Sympos. Comput. Geom.*, pp. 330–339, 2004.
11. H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511–533.
12. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Comput. Graphics* **13** (1994), 43–72.
13. J. H. G. Fu. Tubular neighborhoods in Euclidean spaces. *Duke Math. J.* **52** (1995), 1025–1046.
14. M. Goresky and R. MacPherson. *Stratified Morse Theory.* Springer-Verlag, New York, 1988.
15. L. Lovász and M. D. Plummer. *Matching Theory.* North-Holland, Amsterdam, 1986.
16. J. Milnor. *Morse Theory.* Princeton University Press, New Jersey, 1963.
17. J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, California, 1984.
18. P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. Manuscript, Toyota Technological Institute, Chicago, Illinois, 2004.
19. V. Robins. Toward computing homology from finite approximations. *Topology Proceedings* **24** (1999), 503–532.
20. G. Sapiro. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, Cambridge, 2001.
21. E. Spanier. *Algebraic Topology.* Springer-Verlag, New York, 1966.
22. G. Strang. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, Wellesley, Massachusetts, 1993.
23. A. Zomorodian and G. Carlsson. Computing persistent homology. In *Proc. 20th Ann. Sympos. Comput. Geom.*, pp. 347–356, 2004.

Received July 19, 2005, *and in revised form September* 17, 2005. *Online publication December* 5, 2006.