

Extending Persistence Using Poincaré and Lefschetz Duality *

David Cohen-Steiner[†], Herbert Edelsbrunner[‡] and John Harer[§]

communicated by Konstantin Mischaikow

Abstract

Persistent homology has proven to be a useful tool in a variety of contexts, including the recognition and measurement of shape characteristics of surfaces in \mathbb{R}^3 . Persistence pairs homology classes that are born and die in a filtration of a topological space, but does not pair its actual homology classes. For the sublevelset filtration of a surface in \mathbb{R}^3 persistence has been extended to a pairing of essential classes using Reeb graphs. In this paper, we give an algebraic formulation that extends persistence to essential homology for any filtered space, present an algorithm to calculate it, and describe how it aids our ability to recognize shape features for codimension 1 submanifolds of Euclidean space. The extension derives from Poincaré duality but generalizes to non-manifold spaces. We prove stability for general triangulated spaces and duality as well as symmetry for triangulated manifolds.

AMS classification. 55N99: Algebraic topology, Homology and cohomology theories; 68W30: Computer science, Algorithms, symbolic computation and algebraic computation.

Keywords. Computational topology, homology, persistence, manifolds, duality, simplicial complexes, algorithms.

1 Introduction

Persistent homology was recently introduced to measure and possibly remove topological features in continuous data. This ability has broad applications in medical imaging, scientific visualization, and other fields of human study. In this paper, we extend persistence to essential homology classes.

*Research by all three authors is partially supported by DARPA under grant HR0011-05-1-0007. Research by the second author is also partially supported by NSF under grant CCR-00-86013.

[†]INRIA, 2004 Route des Lucioles, BP93, Sophia-Antipolis, France.

[‡]Department of Computer Science, Duke University, Durham, and Geomagic, Research Triangle Park, North Carolina, USA.

[§]Department of Mathematics and Center for Computational Science, Engineering and Medicine, Duke University, Durham, North Carolina, USA.

Motivation. Continuous functions have been used in a variety of contexts to recognize and extract features of embedded surfaces. Examples include local curvature approximation functions, used e.g. for protein docking [4] and global alignment [7], and the mean geodesic distance, used for shape matching and indexing [8]. Particularly relevant to this paper is the elevation function introduced in [1] and applied to coarse protein docking in [11]. To define the elevation at a point p of a surface S , we consider the directional height function given by projecting S onto its normal line through p , and pair up the critical points of that function using persistence. The elevation at p is then the absolute directional height difference between p and its paired point. To define the function for all points of S , [1] extends the persistence pairing to essential cycles that represent the homology of S using a rule about forks and loops in the corresponding Reeb graph [1]; see also Section 2. For all this to make sense, the pairing needs to be symmetric so that we get consistent values in opposite height directions.

A unique property of the elevation function is its independence of scale. More precisely, the scale that enters in the definition of the value at a point varies depending on the scale of the feature this point is chosen to represent. The elevation function thus lies somewhere between the primarily local curvature and curvature approximation functions and the primarily global mean distance function.

Results and prior work. This paper gives an algebraic justification for the choices made in [1], and it provides the topological results needed to generalize elevation to smoothly embedded d -manifolds in \mathbb{R}^{d+1} . The main technical results are

- (a) the extension of persistence to essential homology classes using Poincaré and Lefschetz duality;
- (b) a characterization of extended persistence in terms of optimally local homology bases;
- (c) an efficient algorithm to compute the pairs of critical values that characterize extended persistence;
- (d) a proof that these pairs of critical values are stable;

- (e) a proof that for manifolds these pairs satisfy a duality relation between dimensions r and $d - r$;
- (f) a proof that for manifolds these pairs are symmetric, that is, the same for a function and its negative.

The concept of persistent homology has originally been defined for modulo 2 homology in [5] and has later been generalized to fields in [12]. In both formulations, persistence does not measure essential homology classes, which is limiting in some applications. Our result (a) overcomes this limitation. Extended persistence has an intuitive interpretation in terms of homology bases whose classes lexicographically minimize thickness, providing the characterization (b). The hallmark of the original work on persistence is an efficient reduction algorithm, which we adapt to extended persistence giving us (c). The stability result (d) is proved by reducing it to the known stability of ordinary persistence [2]. This reduction is to the combinatorial formulation given in [3] and has the side-benefit of implying a linear-time algorithm for maintaining the pairing under continuous changes of the function. The duality result (e) is based on Lefschetz duality and breaks down for non-manifold spaces. Using duality, we prove symmetry (f), again for manifolds. Both results are new except for the 2-manifold case [1].

Outline. Section 2 gives the idea behind extended persistence. Section 3 presents background on filtrations and ordinary persistence. Section 4 extends ordinary persistence. Section 5 proves that extended persistence selects optimally local homology bases. Section 6 gives the algorithm for extended persistence and proves its stability. Section 7 proves duality and symmetry for manifolds. Section 8 concludes the paper.

2 Intuition

In this section, we describe persistence and its extension intuitively, for a Morse function on a manifold. We hope this prepares the reader for the formal treatment of the more general case and conveys a feeling for the type of applications that can benefit from this concept. Throughout the paper we will work with homology with $\mathbb{Z}/2\mathbb{Z}$ coefficients, so we will write $H_r(X) = H_r(X, \mathbb{Z}/2\mathbb{Z})$ and denote its rank by $\beta_r(X)$, calling it the r -th Betti number. We point out that this is not the normal usage in algebraic topology, where $\beta_r(X)$ denotes the rank of the homology of X with rational coefficients. The usual Betti numbers are generally smaller than ours since, by the universal coefficient theorem [9], $H_r(X, \mathbb{Z}/2\mathbb{Z})$ also includes 2-torsion.

Birth and death. Let \mathbb{M} be a manifold of dimension d and $f : \mathbb{M} \rightarrow \mathbb{R}$ a Morse function on the manifold. We can imagine that \mathbb{M} is smoothly embedded in \mathbb{R}^{d+1} and f maps every point to its height above some hyperplane. Recall that a Morse function has only non-degenerate critical points, all

of which have distinct critical values. We choose regular values $t_0 < t_1 < \dots < t_m$ bracketing the m critical values, and let $\mathbb{M}_k = f^{-1}(-\infty, t_k]$ be the sublevel set containing the first k critical points. Morse theory tells us that \mathbb{M}_k is homotopy equivalent to the result of attaching an r -dimensional cell to \mathbb{M}_{k-1} , where r is the index of the k -th critical point. Specifically, \mathbb{M}_k is homeomorphic to \mathbb{M}_{k-1} with an r -handle attached. This means that either $\beta_r(\mathbb{M}_k) = \beta_r(\mathbb{M}_{k-1}) + 1$ or $\beta_{r-1}(\mathbb{M}_k) = \beta_{r-1}(\mathbb{M}_{k-1}) - 1$. To distinguish the two cases, we call the critical point in the first case *positive* since it increases the sum of Betti numbers and in the second case *negative* since it decreases the sum of Betti numbers.

Persistence gives a pairing between positive critical points of index r and negative critical points of index $r + 1$. To be more specific, it introduces the idea of a homology class coming into existence at a particular moment and leaving again at some particular, later moment. To make sense of this idea, we use the maps between homology groups induced by the inclusions $\mathbb{M}_i \subseteq \mathbb{M}_j$ whenever $i \leq j$. We say a homology class $\alpha \in H_r(\mathbb{M}_k)$ is *born at* \mathbb{M}_k if it does not lie in the image of the map induced by $\mathbb{M}_{k-1} \subseteq \mathbb{M}_k$. Furthermore, if α is born at \mathbb{M}_k we say it *dies entering* \mathbb{M}_ℓ if the image of the map induced by $\mathbb{M}_{k-1} \subseteq \mathbb{M}_\ell$ does not contain the image of α but the image of the map induced by $\mathbb{M}_k \subseteq \mathbb{M}_\ell$ does. If α is born at \mathbb{M}_k and dies entering \mathbb{M}_ℓ then we pair the corresponding critical points, p_k and p_ℓ , and say their *persistence* is $\ell - k$ or $f(p_\ell) - f(p_k)$, depending on the application we have in mind. The latter is more common because it measures size in terms of the function that gives rise to the class in the first place. Homology classes that are born at \mathbb{M}_k and do not die are not paired by this method. We call these classes the *essential* homology classes of \mathbb{M} . They represent interesting features so it would be useful to extend the persistence pairing to include them. How to do this is the topic of this paper.

Torus example. To get a feeling for the information captured by the pairs of critical points, we consider the example illustrated in Figure 1. On the left, we see a torus (a sphere with a tunnel) embedded in \mathbb{R}^3 such that its height function has two minima, four saddles, and two maxima. On the right, we see the Reeb graph of the height function [10]. Each point of this graph represents a component of a level set. As we sweep the torus from bottom to top, each minimum starts a component, each saddle either splits a component into two or merges two components into one, and each maximum ends a component. The pairs of critical points are (b, B) for $r = 0$ and (d, D) for $r = 1$. Indeed, the minimum b gives birth to a component of the sublevel set which dies when B merges this component with the other one given birth to earlier by a . Similarly, the saddle d gives birth to a 1-dimensional homology class which dies at the hand of D . The remaining four critical points, a, c, C, A , are not paired as they give birth to the four essential homology classes of the torus. Our aim is to find an algebraic justification for pairing a with A and c




Figure 1: Left: smoothly embedded torus with select level sets bracketing the critical points of the height function. Right: Reeb graph of the height function.

with C . The former pair represents the component that eventually grows into the entire surface as well as the 2-cycle the surface forms when complete. The latter pair represents the tunnel, marking where it begins and where it ends if viewed during the sweep of the height function.

Elementary rules for pairing the extra critical points have been given in [1] and are easiest explained in terms of the nodes in the Reeb graph. Each component of the 2-manifold corresponds to a component of the Reeb graph and we pair the global minimum (the lowest node) with the global maximum (the highest node). Furthermore, each saddle corresponds to a degree-3 node which either forks up or down. A loop in the Reeb graph has the minimum at its lowest up-fork and the maximum at its highest down-fork, and we say the two forks *span* the loop. According to [1], each up-fork is to be paired with the lowest down-fork that spans a loop with it. It can be shown that every critical point belongs to exactly one pair and that reversing the direction in which we measure height gives the same pairing.

Cohomology and relative homology. In this paper, we extend persistence to essential homology using Poincaré duality, which provides an isomorphism between the dimension r homology group and the dimension $d - r$ cohomology group of a d -manifold, $H_r(\mathbb{M}) \cong H^{d-r}(\mathbb{M})$. The inclusions $\mathbb{M}_k \subseteq \mathbb{M}_{k+1}$ induce maps $H^{d-r}(\mathbb{M}_{k+1}) \rightarrow H^{d-r}(\mathbb{M}_k)$ so we get an extended sequence

$$\begin{aligned} 0 = H_r(\mathbb{M}_0) &\rightarrow \dots \rightarrow H_r(\mathbb{M}_m) \\ &\rightarrow H^{d-r}(\mathbb{M}_m) \rightarrow \dots \rightarrow H^{d-r}(\mathbb{M}_0) = 0, \end{aligned}$$

where $\mathbb{M}_0 = \emptyset$ and $\mathbb{M}_m = \mathbb{M}$. Classes get born and die also during the new, second half of the sequence. In particular, each class born at \mathbb{M}_k that lives all the way to the end of the first half now dies entering some \mathbb{M}_ℓ in the second half. The

only problem with this definition is that to make good sense of it requires that we understand persistence for cohomology. While this can certainly be done [6], we instead use Lefschetz duality to replace cohomology by relative homology, since it gives an isomorphism $H^{d-r}(\mathbb{M}_k) \cong H_r(\mathbb{M}_k, \partial\mathbb{M}_k)$ [9]. Letting $\mathbb{M}^{m-k} = f^{-1}[t_k, \infty)$ be the superlevel set containing the last $m - k$ critical points, we use excision to get an isomorphism $H_r(\mathbb{M}_k, \partial\mathbb{M}_k) \cong H_r(\mathbb{M}, \mathbb{M}^{m-k})$ [9] and another version of the above sequence,

$$\begin{aligned} 0 = H_r(\mathbb{M}_0) &\rightarrow \dots \rightarrow H_r(\mathbb{M}_m) \\ &\rightarrow H_r(\mathbb{M}, \mathbb{M}^0) \rightarrow \dots \rightarrow H_r(\mathbb{M}, \mathbb{M}^m) = 0, \end{aligned}$$

where $\mathbb{M}_0 = \mathbb{M}^0 = \emptyset$ and $\mathbb{M}_m = \mathbb{M}^m = \mathbb{M}$. This version is not only easier to analyze algorithmically but also generalizes naturally to simplicial complexes with a total ordering on their vertices. The next sections elaborate on this idea.

Working with this sequence, we can understand extended persistence in the following way. Suppose that passing from \mathbb{M}_{k-1} to \mathbb{M}_k gives birth to an essential, r -dimensional homology class. Descending through the superlevel sets, we look for the first $\mathbb{M}^{m-\ell}$ (the largest ℓ) that contains a class homologous in \mathbb{M} to this essential class. This class then dies entering $\mathbb{M}^{m-\ell}$ and we pair p_k with p_ℓ . For 2-manifolds, this is precisely the elementary rule about forks and loops in Reeb graphs given above. We define the *extended persistence* of this pair to be $|k - \ell|$ or $|f(p_k) - f(p_\ell)|$, depending on the application.

3 Persistence

In this section, we formally introduce persistence for a nested sequence of complexes. Two particularly useful such sequences mimic the sweep in the direction of increasing and decreasing function values.

Pair groups. Let K be a simplicial complex of finite dimension d . A *filtration* of K is a nested sequence of sub-complexes that starts with the empty complex and ends with the complete complex,

$$\emptyset = K_0 \subset K_1 \subset \dots \subset K_n = K.$$

Intuitively, persistence is a measure of how long a homology class lives in this filtration. Fix a dimension r and let $H_r^i = H_r(K_i)$ be the r -dimensional homology group of K_i . Letting $h_r^{i,j} : H_r^i \rightarrow H_r^j$ be the homomorphism induced by the inclusion of K_i in K_j , the image, $\text{im } h_r^{i,j}$, consists of the classes that live at least from H_r^i to H_r^j . For a particular class α we can measure its persistence by tracking where it is born and where it dies.

DEFINITION. A class $\alpha \in H_r^i - \text{im } h_r^{i-1,i}$ is *born* at K_i and it *dies entering* K_j if $h_r^{i,j-1}(\alpha) \notin \text{im } h_r^{i-1,j-1}$ but $h_r^{i,j}(\alpha) \in \text{im } h_r^{i-1,j}$. The *persistence* of α is the number of groups in which it lives, i.e. $j - i$.

Note that each class born at K_i forms a coset of $\text{im } h_r^{i-1,i}$ in H_r^i , so the set of new classes born corresponds to the quotient group $H_r^i / \text{im } h_r^{i-1,i}$. Furthermore, if α is born at K_i and dies entering K_j then there exists another such class $\tilde{\alpha}$ whose image in H_r^j vanishes. More formally, there exists a class $\tilde{\alpha} \in H_r^i - \text{im } h_r^{i-1,i}$ such that

- $\tilde{\alpha} - \alpha \in \text{im } h_r^{i-1,i}$,
- $h_r^{i,j-1}(\tilde{\alpha}) \notin \text{im } h_r^{i-1,j-1}$, and
- $h_r^{i,j}(\tilde{\alpha}) = 0$.

This suggests that we introduce the *pair group* of all classes born at K_i that vanish entering K_j :

$$P_r^{i,j} = \frac{\text{im } h_r^{i,j-1} \cap \ker h_r^{j-1,j}}{\text{im } h_r^{i-1,j-1} \cap \ker h_r^{j-1,j}}. \quad (1)$$

Indeed, the numerator collects all classes that exist in H_r^i and vanish as we pass from K_{j-1} to K_j . Some of these classes are born at K_i and others exist as images of classes in H_r^{i-1} . The denominator collects the latter and thus maps them to zero in the quotient group.

Lower and upper stars and links. We have introduced persistence for two rather different topological categories, Morse functions on manifolds and filtrations of simplicial complexes. We feel the need to relate the two and we prepare this unification by introducing local substructures around the vertices in a simplicial complex. The *star* (sometimes called the *open star*) of a vertex $u \in K$ is the set of simplices that contain u as a face, including u itself, and the *link* consists of all faces of simplices in the star that do not belong to the star:

$$\begin{aligned} \text{St } u &= \{ \sigma \in K \mid u \in \sigma \}, \\ \text{Lk } u &= \{ \tau \in K - \text{St } u \mid \tau \subseteq \sigma \in \text{St } u \}. \end{aligned}$$

The union of the star and the link is the *closed star*, $\overline{\text{St}} u = \text{St } u \cup \text{Lk } u$.

Suppose we have an injective, real-valued function f defined on the vertex set. Imagining that f measures height, we say x is *lower than* u and u is *higher than* x if $f(x) < f(u)$. The *lower star* of a vertex u consists of all simplices in the star for which u is the highest vertex, and the *lower link* consists of all simplices in the link that are faces of the simplices in the lower star:

$$\begin{aligned} \text{St}_- u &= \{ \sigma \in \text{St } u \mid x \in \sigma \Rightarrow f(x) \leq f(u) \}, \\ \text{Lk}_- u &= \{ \tau \in \text{Lk } u \mid \tau \subseteq \sigma \in \text{St}_- u \}. \end{aligned}$$

The union of the lower star and the lower link is the *closed lower star*, $\overline{\text{St}}_- u = \text{St}_- u \cup \text{Lk}_- u$. Every simplex has a unique highest vertex and thus belongs to the lower star of a unique vertex. This implies that the lower stars of the vertices partition K . We also define the *upper star*, $\text{St}^+ u$, the *upper link*, $\text{Lk}^+ u$, and the *closed upper star*, $\overline{\text{St}}^+ u = \text{St}^+ u \cup \text{Lk}^+ u$, in the analogous way. Similar to the lower stars, the upper stars of the vertices partition K .

Vertex ordering filtrations. Assuming an injective, real-valued function f on the vertex set, as before, we order the vertices as $f(u_1) < f(u_2) < \dots < f(u_n)$. Let K_i be the subcomplex of K spanned by the first i vertices. Hence, K_i is obtained from K_{i-1} by adding the lower star of u_i or, more precisely, by attaching the closed lower star along the lower link of u_i . Similarly, let L_{n-i} be the subcomplex of K spanned by the last $n-i$ vertices, so L_{n-i} is obtained from L_{n-i-1} by adding the upper star of u_{i+1} . The K_i form an ascending filtration and the L_j form a descending filtration,

$$\emptyset = K_0 \subset K_1 \subset \dots \subset K_n = K,$$

$$\emptyset = L_0 \subset L_1 \subset \dots \subset L_n = K,$$

where the words ‘ascending’ and ‘descending’ refer to the ordering of the vertices as they enter the filtrations. We will use the ascending filtration to compute persistence, and the two filtrations together to define extended persistence. To compute the persistence pairing, it is convenient to totally order the simplices and not just the lower and the upper stars. We do this by listing the lower stars in order and listing the simplices within each lower star by dimension, breaking the remaining ties arbitrarily. We write $\kappa_1, \kappa_2, \dots, \kappa_m$ for the resulting ascending sequence, where m is the number of simplices in K . We follow the analogous rule in turning the ordering of upper stars into a total order, and write $\lambda_1, \lambda_2, \dots, \lambda_m$ for the resulting descending sequence of simplices.

By adding simplices one at a time, we obtain alternative, finer filtrations of K which can be used to compute persistence. To insure that the answer is independent of the choice of order within the lower and upper stars, however, we use the indices or function values of the vertices to define persistence. Classes that are born and then die within a single lower or upper star will therefore have zero persistence.

Induced partition. The complexes K_i and L_{n-i} partition the vertices but they do not necessarily contain all simplices of K . Let J_i be the part of K contained in neither K_i nor L_{n-i} and note that J_i contains the simplices that connect vertices from the two sets, but has no vertices itself. In sit-




Figure 2: The set J_i consists of the shaded triangles and thick edges that connect the vertices of K_i below and vertices of L_{n-i} above the belt they form. The dotted edges indicate the stars in $\text{Sd } K$ of some of the vertices in L_{n-i} .

uations in which it is more convenient to work with a cover,

we may subdivide the simplices and add half of J_i to K_i and the other half to L_{n-i} . More formally, we construct the first barycentric subdivision, $\text{Sd } K$ of K , and let K_i^* be the union of closed stars in $\text{Sd } K$ of the vertices in K_i . Define L_{n-i}^* similarly. The following lemma is not difficult to prove.

PARTITION LEMMA. For each i , K_i is homotopy equivalent to K_i^* , L_{n-i} is homotopy equivalent to L_{n-i}^* , and J_i is homotopy equivalent to the intersection $K_i^* \cap L_{n-i}^*$.

This latter intersection is simultaneously the boundary of K_i^* and of L_{n-i}^* . Note that when K is a triangulation of a d -manifold, then K_i^* and L_{n-i}^* are d -manifolds with boundary $K_i^* \cap L_{n-i}^*$.

4 Extension

In this section, we review the definitions of Poincaré and Lefschetz duality and show how they can be used to define extended persistence. For manifolds, we prove a duality result which we then use to prove that both ordinary and extended persistence are symmetric.

Non-degenerate pairings. We find it convenient to express topological duality results in terms of pairings between homology groups over fields, which are vector spaces. Letting U and V be two finite-dimensional vector spaces, a bilinear and symmetric map $\langle \cdot, \cdot \rangle : U \times V \rightarrow \mathbb{Z}_2$ is a *non-degenerate pairing* if for every non-zero $u \in U$ there is a $v \in V$ with $\langle u, v \rangle = 1$ and, symmetrically, for every non-zero $v \in V$ there is a $u \in U$ with $\langle u, v \rangle = 1$. Importantly, a non-degenerate pairing induces a perfect matching between well-chosen bases. Since there are usually no canonical bases and no canonical perfect matchings between given bases, non-degenerate pairings provide a convenient language to talk about them without making choices.

PAIRING LEMMA. For any non-degenerate pairing of two finite-dimensional vector spaces there are bases for which the pairing induces a perfect matching.

PROOF. We construct the two bases that satisfy the claim inductively, starting with a non-zero $u_0 \in U$ and $v_0 \in V$ that satisfies $\langle u_0, v_0 \rangle = 1$. To continue the construction, consider the orthogonal subspaces,

$$\begin{aligned} U_0 &= \{u \in U \mid \langle u, v_0 \rangle = 0\}, \\ V_0 &= \{v \in V \mid \langle u_0, v \rangle = 0\}. \end{aligned}$$

Note that U is the union of U_0 and the set U_1 of vectors $u \in U$ with $\langle u, v_0 \rangle = 1$. This implies that U is isomorphic to the direct sum $U_0 \oplus \text{span}(u_0)$, where the latter is the line spanned by u_0 . Similarly, V is isomorphic to $V_0 \oplus \text{span}(v_0)$.

To prove that the restriction of $\langle \cdot, \cdot \rangle$ to U_0 and V_0 is again non-degenerate, we take a non-zero $u \in U_0$. If $\langle u, v \rangle = 0$ for every $v \in V_0$ then $\langle u, v + v_0 \rangle = 0$ by linearity. But this

exhausts all vectors in V which contradicts that the original pairing was non-degenerate. Symmetrically, for each non-zero $v \in V_0$ there exists $u \in U_0$ with $\langle u, v \rangle = 1$. By induction, there are bases of U_0 and V_0 for which $\langle \cdot, \cdot \rangle$ induces a perfect matching. We add u_0 and v_0 and get bases of U and V , again with an induced perfect matching. \square

Poincaré and Lefschetz duality. Let \mathbb{M} be a triangulated manifold of dimension d . Since we work with \mathbb{Z}_2 coefficients, we do not need to assume that \mathbb{M} is orientable. When \mathbb{M} has no boundary, Poincaré duality provides an isomorphism between the dimension r homology group and the dimension $d - r$ cohomology group of \mathbb{M} , for each r . An alternative formulation that does not use cohomology states that there is a non-degenerate pairing

$$\langle \cdot, \cdot \rangle : H_r(\mathbb{M}) \times H_{d-r}(\mathbb{M}) \rightarrow \mathbb{Z}_2$$

given by counting intersections of cycles in \mathbb{M} . What this means is the following. Suppose that $\alpha \in H_r(\mathbb{M})$ and $\beta \in H_{d-r}(\mathbb{M})$. Then we can find cycles a and b representing α and β in, say, the first barycentric subdivision of \mathbb{M} that intersect transversally. Since they have complementary dimensions, a and b intersect in a finite number of points. The pairing is then obtained by counting intersections, $\langle \alpha, \beta \rangle = \text{card}(a \cap b) \pmod{2}$. A standard result is that this pairing is well-defined, bilinear, symmetric, and non-degenerate.

When \mathbb{M} has non-empty boundary $\partial\mathbb{M}$, Poincaré duality is replaced by Lefschetz duality, which provides a non-degenerate pairing

$$\langle \cdot, \cdot \rangle : H_r(\mathbb{M}, \partial\mathbb{M}) \times H_{d-r}(\mathbb{M}) \rightarrow \mathbb{Z}_2$$

also defined by counting intersections. Note that the pairing without $\partial\mathbb{M}$ is still well-defined, but it is degenerate. A simple example is the annulus $\mathbb{M} = \mathbb{S}^1 \times [0, 1]$, drawn as the lower portion of a torus in Figure 3. The dimension of the annulus is $d = 2$ so for $r = 1$ we have $d - r = 1$. Both $H_1(\mathbb{M})$ and $H_1(\mathbb{M}, \partial\mathbb{M})$ are rank one. The self-intersection of the generator of $H_1(\mathbb{M})$ is zero, since the cycles $\mathbb{S}^1 \times 0$ and $\mathbb{S}^1 \times 1$ both represent the generator yet do not intersect. Thus the pairing is degenerate. A generator of $H_1(\mathbb{M}, \partial\mathbb{M})$ intersects every representative of the generator of $H_1(\mathbb{M})$ in one point, modulo 2. So the pairing between $H_1(\mathbb{M}, \partial\mathbb{M})$ and $H_1(\mathbb{M})$ is non-degenerate.

Extended persistence. Suppose now that K is a triangulation of \mathbb{M} and f is an injective, real-valued function on the vertex set. As in Section 3, we order the vertices as $f(u_1) < f(u_2) < \dots < f(u_n)$ and we construct the ascending and the descending filtrations by letting K_i be the subcomplex spanned by the first i vertices and L_{n-i} the subcomplex spanned by the last $n - i$ vertices. The sequence of




Figure 3: The circle around the annulus generates $H_1(\mathbb{M})$ and intersects the arc generating $H_1(\mathbb{M}, \partial\mathbb{M})$ in one point.

homology groups $H_r(K_i)$, for $0 \leq i \leq n$, can be used to define persistence of the homology classes that die before the sequence ends. To extend persistence to all classes, we continue the sequence using the Poincaré duality isomorphism $H_r(K) \rightarrow H^{d-r}(K)$:

$$\begin{aligned} 0 = H_r(K_0) &\rightarrow \dots \rightarrow H_r(K_n) \\ &\rightarrow H^{d-r}(K_n) \rightarrow \dots \rightarrow H^{d-r}(K_0) = 0. \end{aligned}$$

Using excision and Lefschetz duality, we can replace all cohomology groups by relative homology groups and get

$$\begin{aligned} 0 = H_r(K_0) &\rightarrow \dots \rightarrow H_r(K_n) \\ &\rightarrow H_r(K, L_0) \rightarrow \dots \rightarrow H_r(K, L_n) = 0. \end{aligned}$$

To express this more compactly, we identify K_n with (K_n, \emptyset) and write $K_{n+i} = (K, L_i)$, for $0 \leq i \leq n$. Furthermore we write $H_r^j = H_r(K_j)$, where this is of course ordinary homology for $0 \leq j \leq n$ and relative homology for $n < j \leq 2n$. Using this notation, we rewrite the extended sequence of homology groups as

$$0 = H_r^0 \rightarrow \dots \rightarrow H_r^n \rightarrow \dots \rightarrow H_r^{2n} = 0.$$

Since every class eventually dies, we simply adapt the definition above. Specifically, a class that is born at K_i or K_{n+i} and dies entering K_j or K_{n+j} , for $0 < i, j \leq n$, has *extended persistence* $|j - i|$ or $|f(u_j) - f(u_i)|$, depending on the application. Although Poincaré and Lefschetz duality do not extend to general simplicial complexes, the extended sequence of homology groups does.

Case analysis. It is useful to know what happens when we go from one homology group to the next in this extended sequence. To simplify the analysis, we consider the finer filtration in which we add one simplex at a time, so we index the homology groups from H_r^0 to H_r^m and further to H_r^{2m} , where m is the number of simplices in K , as usual. We begin with the ascending filtration and consider the kernels $K_r^i = \ker h_r^{i,m}$ and the cokernels $C_r^i = H_r(K)/\text{im } h_r^{i,m}$ of the maps into the homology group of the full complex. The kernel consists of all inessential cycles in K_i , the ones that are trivial in K . The cokernel represents all new essential cycles in K , the ones that have no preimage in K_i . The maps

induced by the inclusions $K_{i-1} \subseteq K_i \subseteq K$ in turn induce maps between the kernels and the cokernels. To simplify the discussion, we drop subscripts and superscripts and write $k_A : K_{r-1}^{i-1} \rightarrow K_{r-1}^i$ and $c_A : C_{r-1}^{i-1} \rightarrow C_r^i$ for these maps. In going from $i-1$ to i , we get K_i by adding κ_i to K_{i-1} , and we distinguish between three types of events.

Case A.1 The map k_A has a non-trivial kernel, $\ker k_A \neq 0$. This means there is a non-trivial class in H_{r-1}^{i-1} that goes to zero in H_{r-1}^i . This class dies, the dimension $r-1$ homology group decreases in rank by one, and κ_i is a negative r -simplex in the ascending filtration that gets paired with an earlier, positive $(r-1)$ -simplex.

Case A.2 The same map, k_A , has a non-trivial cokernel, $\text{coker } k_A \neq 0$. This means there is a new class in H_{r-1}^i that is inessential in K . This class is born at K_i , the dimension $r-1$ homology group increases in rank by one, and κ_i is a positive $(r-1)$ -simplex that will get paired with a negative r -simplex later.

Case A.3 The map c_A has a non-trivial kernel, $\ker c_A \neq 0$. This means there is a new class in H_r^i that is essential in K . This class is born at K_i , the dimension r homology group increases in rank by one, and κ_i is a positive r -simplex that will not get paired within the ascending filtration.

There is no fourth case because c_A is necessarily surjective and thus cannot have a non-trivial cokernel. Note that this same analysis applies also to the descending filtration of the L_j . To relate the change in the homology of the L_j to that of the pairs (K, L_j) , we consider the long exact sequence of the pair,

$$\dots \rightarrow H_r(L_j) \rightarrow H_r(K) \rightarrow H_r(K, L_j) \rightarrow H_{r-1}(L_j) \rightarrow \dots$$

Reusing the notation for kernels and cokernels, but now defined for the descending complexes, we write K_r^j for the kernel of $H_r(L_j) \rightarrow H_r(K)$ and C_r^j for the cokernel of the same map. With this notation, we get the short exact sequence

$$0 \rightarrow C_r^j \rightarrow H_r(K, L_j) \rightarrow K_{r-1}^j \rightarrow 0.$$

Since we are working with \mathbb{Z}_2 coefficients, this sequence splits to give an isomorphism, $H_r(K, L_j) \cong C_r^j \oplus K_{r-1}^j$. As before, the inclusions $L_{j-1} \subseteq L_j \subseteq K$ induce maps between the corresponding homology groups which, in turn, induce maps between kernels and cokernels. These maps make the diagram

$$\begin{array}{ccccccc} 0 & \rightarrow & C_r^{j-1} & \rightarrow & H_r(K, L_{j-1}) & \rightarrow & K_{r-1}^{j-1} \rightarrow 0 \\ & & \downarrow c_D & & \downarrow & & \downarrow k_D \\ 0 & \rightarrow & C_r^j & \rightarrow & H_r(K, L_j) & \rightarrow & K_{r-1}^j \rightarrow 0 \end{array}$$

commute. This diagram helps us relate the changes between absolute and relative homology groups. In going from (K, L_{j-1}) to (K, L_j) we get L_j by adding λ_j to L_{j-1} . We again distinguish between three cases, in strict analogy with the ascending filtration.

Case D.1 $\ker k_D \neq 0$. Similar to Case A.1, λ_j is a negative r -simplex in the descending filtration and the absolute dimension $r - 1$ homology group decreases in rank by one. The diagram implies the same for the dimension r relative homology group, $\text{rank } H_r(K, L_j) = \text{rank } H_r(K, L_{j-1}) - 1$.

Case D.2 $\text{coker } k_D \neq 0$. Similar to Case A.2, λ_j is a positive $(r - 1)$ -simplex in the descending filtration and the absolute dimension $r - 1$ homology group increases in rank by one. The diagram implies the same for the dimension r relative homology group, $\text{rank } H_r(K, L_j) = \text{rank } H_r(K, L_{j-1}) + 1$.

Case D.3 $\ker c_D \neq 0$, Similar to Case A.3, λ_j is a positive r -simplex in the descending filtration and the absolute dimension r homology group increases in rank by one, with an essential class. The diagram implies the opposite change for the relative homology group of the same dimension, $\text{rank } H_r(K, L_j) = \text{rank } H_r(K, L_{j-1}) - 1$.

As before there is no fourth case because c_D is surjective.

5 Locality

In practical situations, it can be useful to find the most local basis of the homology of a given space. There are many ways to measure locality and we choose thickness, which we define below. This will lead to a new interpretation of extended persistence as selecting a lexicographically thinnest homology basis.

Thickness of a homology class. It is convenient to assume a Morse function $f : \mathbb{M} \rightarrow \mathbb{R}$ on a manifold, although the arguments in this section also work for filtrations of general simplicial complexes. For any interval I , we define $\mathbb{M}^I = f^{-1}(I)$. Given a class $\alpha \in H_r(\mathbb{M})$, we let $\mathcal{I}(\alpha)$ be the set of closed intervals I such that α lies in the image of $H_r(\mathbb{M}^I)$ in $H_r(\mathbb{M})$ by inclusion. In other words, \mathbb{M}^I supports a representative cycle of α . We begin by proving that the intersection of two intervals in $\mathcal{I}(\alpha)$ is either empty or again belongs to $\mathcal{I}(\alpha)$. We distinguish two kinds of homology classes, calling α *horizontal* if $\mathcal{I}(\alpha)$ contains at least two non-empty disjoint intervals and *vertical* otherwise.

INTERSECTION LEMMA. For vertical α , $\mathcal{I}(\alpha)$ is closed under intersection.

PROOF. Let $I' = [w, y]$ and $J' = [x, z]$ be two intervals in $\mathcal{I}(\alpha)$. Since α is vertical, their intersection is non-empty. If they are nested then $I' \cap J'$ belongs to $\mathcal{I}(\alpha)$ for trivial reasons. If not, we assume $w \leq x \leq y \leq z$ and let $I = (-\infty, y]$ and $J = [x, \infty)$. The Mayer-Vietoris sequence for the two corresponding preimages is

$$\dots \rightarrow H_r(\mathbb{M}^{I \cap J}) \rightarrow H_r(\mathbb{M}^I) \oplus H_r(\mathbb{M}^J) \rightarrow H_r(\mathbb{M}) \rightarrow \dots$$

Since $I' \subseteq I$ and $J' \subseteq J$, the class $\alpha \in H_r(\mathbb{M})$ is the image of classes $\beta \in H_r(\mathbb{M}^I)$ and $\gamma \in H_r(\mathbb{M}^J)$. Because the second map in the Mayer-Vietoris sequence is induced by taking the difference of the two inclusions, (β, γ) lies in its kernel. By exactness, this implies that there is a class $\delta \in H_r(\mathbb{M}^{I \cap J})$ whose images in $H_r(\mathbb{M}^I)$ and $H_r(\mathbb{M}^J)$ are β and γ . It follows that the image of δ in $H_r(\mathbb{M})$ is α . Hence $I' \cap J' = I \cap J \in \mathcal{I}(\alpha)$, as claimed. \square

We call the common intersection of the intervals in $\mathcal{I}(\alpha)$ the *minimal support* of α , $I(\alpha) = \bigcap \mathcal{I}(\alpha)$. The *thickness* of α is either the length of this interval or zero, depending on the type of the class,

$$\text{th}(\alpha) = \begin{cases} |I(\alpha)| & \text{if } \alpha \text{ is vertical,} \\ 0 & \text{if } \alpha \text{ is horizontal.} \end{cases}$$

In the horizontal case, $\mathcal{I}(\alpha)$ contains two disjoint intervals and we can choose x so one is contained in $(-\infty, x]$ and the other in $[x, \infty)$. Applying the same argument as we used in the proof of the Intersection Lemma, we see that the degenerate interval $[x, x]$ is in $\mathcal{I}(\alpha)$ so that a representative cycle is supported in a single level set, thus the name horizontal. Note that this level set is not unique. Notice also that in the vertical case, we can have $\text{th}(\alpha) = 0$ in which case there is a unique level set $f^{-1}(x)$ supporting α and $I(\alpha) = [x, x]$.

Relation to extended persistence. Suppose a class α is born at y in the upward pass and dies entering x in the downward pass. Then $(-\infty, y]$ and $[x, \infty)$ both belong to $\mathcal{I}(\alpha)$. Furthermore, $x \leq y$ implies α is vertical and $x > y$ implies α is horizontal. To shed additional light on the relation between thickness and extended persistence, we prove a statement about sums of essential cycles. Call a vertical class $\alpha \in H_r(\mathbb{M})$ *decomposable* if it can be written as a sum of classes whose minimal supports are strict subsets of $I(\alpha)$. This includes the possibility that α is the sum of horizontal classes.

DECOMPOSITION LEMMA. Let $\alpha \in H_r(\mathbb{M})$ be a vertical class. The endpoints of $I(\alpha)$ are paired by extended persistence iff α is not decomposable.

PROOF. Let $I(\alpha) = [x, y]$ with $x \leq y$. Since the class is vertical, α is born at y in the upward pass and dies entering x in the downward pass. Suppose that y is paired with y' . Since α is supported in $[x, y]$, the image of α in $H_r(\mathbb{M}, \mathbb{M}^{[x, \infty)})$ is zero, which implies $x \leq y'$. The claim is therefore equivalent to $x = y'$ if α is non-decomposable and $x < y'$ if α is decomposable. In the remainder of this proof we simplify notation by writing $\mathbb{M}_y = \mathbb{M}^{(-\infty, y]}$ and $\mathbb{M}^x = \mathbb{M}^{[x, \infty)}$ for the sub- and superlevel sets.

To prove the first implication we show that $x < y'$ implies α is decomposable. As mentioned in Section 3, there is a class β that is born at y in the upward pass and dies entering y' in the downward pass. If β is horizontal it has a representative cycle in a level set above y' , and if it is vertical then

$I(\beta) \subseteq [y', y]$ is a non-empty proper subinterval of $I(\alpha)$. By assumption of f being Morse, there is a sufficiently small $\varepsilon > 0$ such that the only difference in Betti numbers at y and $y - \varepsilon$ is $\beta_r(\mathbb{M}_y) = \beta_r(\mathbb{M}_{y-\varepsilon}) + 1$. Hence, the relative group $H_r(\mathbb{M}_y, \mathbb{M}_{y-\varepsilon})$ has rank one, which implies that α and β have the same image there. This means the difference $\alpha - \beta$ goes to zero. It follows that $\alpha - \beta \in H_r(\mathbb{M})$ is the image of a class $\gamma \in H_r(\mathbb{M}_{y-\varepsilon})$. If γ is vertical then $I(\gamma) \subseteq [x, y - \varepsilon]$ is a proper subinterval of $I(\alpha)$. It follows that α is decomposable.

To prove the second implication, we show that α decomposable implies $x < y'$. We may write $\alpha = \beta + \gamma$, where β is supported in $\mathbb{M}_{y-\varepsilon}$ and γ is supported in $\mathbb{M}^{x+\varepsilon}$, for small enough $\varepsilon > 0$. Note that β is already alive at $y - \varepsilon$ and that the images of α and β in $H_r(\mathbb{M}, \mathbb{M}^{x+\varepsilon})$ are the same. By definition of pair group, α is therefore already dead at $x + \varepsilon$. Hence x and y are not paired. \square

Moreover, for any extended persistence pair (y, x) with $x \leq y$ there is a class with minimal support $[x, y]$. Indeed, there is a class $\alpha \in H_r(\mathbb{M})$ that is born at y and maps to zero for the first time in $H_r(\mathbb{M}, \mathbb{M}^x)$. Hence α is supported in \mathbb{M}_y as well as in \mathbb{M}^x . By the Intersection Lemma, $I(\alpha) \subseteq [x, y]$ and strict inclusion would lead to a contradiction.

Most local bases. Let $\eta = (\eta_i)_{i=1..n}$ and $\varphi = (\varphi_i)_{i=1..n}$ be two bases of $H_r(\mathbb{M})$ and assume that basis elements are sorted by non-decreasing thickness. We say η is *thinner* than φ if $th(\eta_j) < th(\varphi_j)$ for the smallest integer j for which $th(\eta_j) \neq th(\varphi_j)$. This is, of course, lexicographic ordering by thickness. Obviously, this order on bases admits at least one global optimum. It can be constructed greedily by taking the i -th element to have the smallest thickness among the classes that are linearly independent of the first $i - 1$ basis elements. We claim that extended persistence gives optimally local bases.

THIN BASIS THEOREM. In an optimally thin basis, the minimal supports of the basis elements with positive thickness are the intervals $[x, y]$ for which (y, x) is an extended persistence pair.

PROOF. Let $\eta = (\eta_i)_{i=1..n}$ be an optimally thin basis of $H_r(\mathbb{M})$, sorting basis elements in the order of non-decreasing thickness, as usual. Let η_j be the first basis element with positive thickness and let $x < y$ be the endpoints of its minimal support, $I(\eta_j) = [x, y]$. Then (y, x) is an extended persistence pair, else η_j would be a sum of classes with strictly smaller thickness, contradicting the optimality of η . Similarly, we prove that for $i \geq j$, η_i corresponds to the extended persistence pair with $(i - j + 1)$ -st smallest value of $y - x$. Finally, all extended persistence pairs (y, x) with $x < y$ correspond to basis elements.

Indeed, if such a pair (y, x) is not a basis element, we could form a new basis by adjoining it to the list of η_j with smaller persistence and extending to a full basis. This new

basis would be strictly thinner, contradicting our assumption that η is the thinnest. \square

6 Algorithm

In this section, we adapt the algorithm of [5] to extended persistence. We phrase the algorithm in the language of column operations on boundary matrices, prove its correctness, and show that the diagrams it produces are stable.

Matrix reduction. The definition of extended persistence suggests a multi-phase algorithm, sweeping the complex K first in ascending and then in descending order. By setting up the data structure appropriately, the seemingly different actions during the phases become the same and we can express the algorithm without making a distinction. This data structure is a $2m$ -by- $2m$ zero-one matrix

$$M = \begin{bmatrix} A & P \\ 0 & D \end{bmatrix},$$

where A represents the ascending filtration, D represents the descending filtration, and P stores the permutation that connects the two sequences of simplices. Recall that the ascending filtration is given by the sequence of simplices $\kappa_1, \kappa_2, \dots, \kappa_m$. The rows and columns of A correspond to the simplices in this order and we have $A[i, j] = 1$ iff κ_i is a co-dimension one face of κ_j . The descending filtration is given by the sequence of simplices $\lambda_1, \lambda_2, \dots, \lambda_m$ and we have $D[i, j] = 1$ iff λ_i is a co-dimension one face of λ_j . Finally, $P[i, j] = 1$ iff $\kappa_i = \lambda_j$.

The matrix M may be interpreted as the boundary matrix of a new complex, $K_\omega = \omega \cdot K$, obtained by coning K from a new vertex ω , but with ω removed. Indeed, A is a boundary matrix of K and so is D . The permutation matrix adds a single 1 above each column of D effectively increasing the dimension of the represented simplex by one. Specifically, the j -th column of D represents the simplex λ_j and the corresponding, $(m + j)$ -th column of M represents the cone $\omega_j = \omega \cdot \lambda_j$. By construction, every complex K_i in the ascending filtration corresponds to an upper-left square submatrix of M , and since it has m or fewer rows and columns, this is also a submatrix of A . The relative homology of the pair (K, L_j) in the extended sequence is the same as that of $(K \cup \omega \cdot L_j, \omega)$. Again by construction, the latter pair corresponds to an upper-left square submatrix of M , which now contains and extends beyond A . With this set-up, we can compute the pairing by reducing the matrix M using column operations. To describe what this means, let $\text{low}(j)$ be the maximum row index i for which $M[i, j] = 1$. If the column j is zero then $\text{low}(j)$ is undefined. We call M *reduced* if $\text{low}(j) \neq \text{low}(k)$ whenever $j \neq k$ are two non-zero columns. Following [3], the algorithm reduces M from left

to right, modifying each column by adding columns on its left, if necessary.

```

for  $j = 1$  to  $2m$  do
  while  $\exists k < j$  with  $\text{low}(k) = \text{low}(j)$  do
    add column  $k$  to column  $j$ 
  endwhile
endfor.

```

Each column operation moves the lowest 1 up until it is either not preceded by another lowest 1 in the same row or the entire column is zero, which happens after fewer than $2m$ steps. Each step takes time at most $O(m)$. Similar to Gaussian elimination, the algorithm takes a total of at most $O(m^2)$ steps and is therefore guaranteed to finish within time $O(m^3)$. In practice, one saves space and time using a sparse matrix representation of M similar to the ones described in [3] or [5].

Interpretation. In an effort to relate the algorithm to the algebra of the preceding sections, we interpret the columns generated by the algorithm in terms of chains, cycles, and boundaries. We distinguish three phases, respectively reducing A , D , and P . Phase I corresponds to $1 \leq j \leq m$ and Phases II and III correspond to $m < j \leq 2m$ but are intermingled in the algorithm. Phase I affects only matrix A and is precisely the algorithm in [5] executed on the ascending filtration. It distinguishes between positive simplices, which may get paired or remain unpaired, and negative simplices, which necessarily get paired.

- We recognize κ_j as a *positive simplex* if column j of A is zero, after reduction. The simplex κ_j is *paired* if row j of A contains a lowest 1. The corresponding column stores an inessential cycle and κ_j is its youngest simplex, in the sense of being introduced last in the ascending filtration. The simplex κ_j is *unpaired* if row j does not contain a lowest 1. In this case, the columns used to zero out column j form an essential cycle and κ_j is its youngest simplex.
- We recognize κ_j as a *negative simplex* if column j of A is non-zero, after reduction. The column corresponds to a chain and stores its boundary, a trivial cycle. Furthermore, κ_j is paired with κ_i , where $i = \text{low}(j)$, and κ_i is positive and the youngest simplex in that cycle.

Note that all cycles stored in A after reduction are boundaries, as indicated in Figure 4. Phase II consists of the actions within $m < j \leq 2m$ that affect matrix D . It is precisely the algorithm of [5] executed on the descending filtration. It thus distinguishes between paired and unpaired positive simplices and paired negative simplices as defined by the descending filtration. Similar to A , all cycles stored in D after reduction are boundaries. Phase III reduces P and thus completes the reduction of M . It modifies only columns of P above zero columns of D . Let $m + j$ be the index of such a column in




Figure 4: The matrix M after Phases I and II but before Phase III. Positive simplices give rise to zero columns and negative simplices to non-zero columns storing boundaries. Without reducing P , we get cycles and chains whose boundaries are stored in D below P .

M and note that its upper half, column j of P , stores a cycle, namely the cycle whose zero boundary is stored below it in D . To this column we add boundaries from A and other cycles from P . The reduction ends when the lowest 1 in the column is in a unique row, with index $i = \text{low}(m + j)$, or the column becomes zero. In the former case, the column stores an essential cycle and we pair κ_i with λ_j .

Correctness. We argue correctness in two steps, first making sure that all simplices are paired, and second that the pairs agree with the algebraic definition of extended persistence. For the first step, we add a new first row and a new first column representing ω to M . The reduction of the resulting matrix is precisely the algorithm in [5] executed on the filtration of K_ω defined by the sequence of simplices $\omega, \kappa_1, \dots, \kappa_m, \omega_1, \dots, \omega_m$. Since K_ω consists of a single connected component and has otherwise no non-trivial homology, ω is the only unpaired simplex, after reduction. This implies that the addition of the new first row and column does not affect the result and we get the same pairs with or without this addition.

The pairs we get from A are the ordinary persistence pairs for the ascending filtration, simply because Phase I of our algorithm is the same as the ordinary persistence algorithm given in [5]. Similarly, the pairs we get from D are the ordinary persistence pairs for the descending filtration. It remains to show that the pairs we get from P are the extended persistence pairs as defined by the algebra. If (κ_i, λ_j) is such a pair then column $m + j$ of M has a non-zero upper half (column j of P) and a zero lower half (column j of D), after reduction. It defines a cycle whose youngest simplex is κ_i , with $i = \text{low}(m + j)$. Since there is no column in A with lowest 1 in row i , the cycle is essential, and since no column operation from the left can change $\text{low}(m + j)$, it is not possible to push this essential cycle lower in K . In other words, K_i is the first complex in the ascending filtration that

contains the essential cycle. By construction, L_j is the first complex in the descending filtration that contains it. It follows that (κ_i, λ_j) is an extended persistence pair as defined by the algebra.

Stability. Following [2], we express the stability of extended persistence in terms of 2-dimensional diagrams, which we now define. Recall that a class in $\mathbb{P}_r^{i,j}$ is born at K_i and dies entering K_j . Encoding the birth and death using coordinates, this class is represented by the point (f_i, f_j) , where $f_i = f(u_i)$ if $i \leq n$ and $f_i = f(u_{2n-i+1})$ if $n < i$, and similarly for j . Collecting the points for all classes of dimension r and adding the points on the diagonal with infinite multiplicity, as in [2], we get the *dimension r persistence diagram* of f , which we denote as $\text{Dgm}_r(f)$. It is convenient to partition $\text{Dgm}_r(f)$ into the *ordinary sub-diagram*, $\text{Ord}_r(f)$, for $i \leq j \leq n$, the *extended sub-diagram*, $\text{Ext}_r(f)$, for $i \leq n < j$, and the *relative sub-diagram*, $\text{Rel}_r(f)$, for $n < i \leq j$. Besides the diagonal points, the ordinary sub-diagram has only points above the diagonal, the extended sub-diagram has points on both sides, and the relative sub-diagram has only points below the diagonal.

Let K be a simplicial complex, not necessarily a triangulation of a manifold, and let $f, g : K \rightarrow \mathbb{R}$ be obtained by linear extension of two real-valued functions on the vertex set. We compare the two functions using the L_∞ -norm of their difference, $\|f - g\|_\infty = \max_x |f(x) - g(x)|$. We compare their persistence diagrams using the bottleneck distance,

$$d_B(\text{Dgm}_r(f), \text{Dgm}_r(g)) = \inf_{\gamma} \sup_p \|p - \gamma(p)\|_\infty,$$

where $p \in \text{Dgm}_r(f)$ and $\gamma : \text{Dgm}_r(f) \rightarrow \text{Dgm}_r(g)$ is a bijection. This is the L_∞ -length of the longest edge in the best matching between the two multisets. Here we use points on the diagonal to complete the matching, if necessary, or to obtain a shorter longest edge, if possible.

STABILITY THEOREM. Letting f and g be piecewise linear functions on a common simplicial complex, we have $d_B(\text{Dgm}_r(f), \text{Dgm}_r(g)) \leq \|f - g\|_\infty$, for all r .

PROOF. We use the reduction algorithm applied to the cone complex $K_\omega = \omega \cdot K$ to generate the diagrams. The sequence of the simplices is $\omega, \kappa_1, \dots, \kappa_m, \omega_1, \dots, \omega_m$, as before. To generate the diagrams for f we assign real numbers to the simplices,

$$\begin{aligned} F(\omega) &= -\infty, \\ F(\kappa_j) &= \max\{f(u) \mid u \in \kappa_j\}, \\ F(\omega_j) &= \min\{f(u) \mid \omega \neq u \in \omega_j\}, \end{aligned}$$

for $1 \leq j \leq m$. All simplices of K in the lower star of a vertex u get assigned the value $f(u)$, and so do all

cones over simplices of K in the upper star of u . The algorithm matches up all simplices in pairs, except for ω . Each pair maps to a point in the persistence diagram whose dimension is that of the first simplex in the pair. We get $\text{Dgm}_r(f) = \text{Dgm}_r(F)$, for all dimensions r . Similarly, we define the assignment G for the function g and construct the diagrams $\text{Dgm}_r(g) = \text{Dgm}_r(G)$, for all r . As proved in [3], the bottleneck distance between $\text{Dgm}_r(F)$ and $\text{Dgm}_r(G)$ is bounded from above by the absolute difference between the assigned values, $d_B(\text{Dgm}_r(F), \text{Dgm}_r(G)) \leq \max_{\sigma \in K_\omega} |F(\sigma) - G(\sigma)|$. By construction, this maximum absolute difference is $\|f - g\|_\infty$, which completes the proof. \square

7 Structure

In this section, we prove that in the case of a d -manifold, both ordinary and extended persistence satisfy a duality result relating dimensions r and $d - r$. We then use this result to prove that extended persistence is symmetric.

Compatibility. Let K be a triangulation of a d -manifold with a total order of its vertices. As usual, K_i is spanned by the first i vertices and L_{n-i} by the last $n - i$ vertices. We recall that $H_r^i = H_r(K_i)$ and $H_{d-r}^{2n-i} = H_{d-r}(K, L_{n-i})$, for $0 \leq i \leq n$. Using the Partition Lemma and excision, we see that the latter, relative homology group is isomorphic to $H_{d-r}(K, L_{n-i}^*) \cong H_{d-r}(K_i^*, \partial K_i^*)$, where K_i^* and L_{n-i}^* are the complementary subcomplexes of the barycentric subdivision of K introduced in Section 3. Lefschetz duality therefore provides a non-degenerate pairing

$$\langle \cdot, \cdot \rangle : H_r^i \times H_{d-r}^{2n-i} \rightarrow \mathbb{Z}_2$$

for each $0 \leq i \leq n$ and each dimension r , given by counting intersections modulo 2. By symmetry, such a pairing exists for each $0 \leq i \leq 2n$. To prepare the proof of the duality result, we show that the pairings given for $i \leq j$ are compatible. By this we mean that if $\alpha \in H_r^i$ and $\beta \in H_{d-r}^{2n-j}$ then we get the same answer if we map α forward to H_r^j and compare it with β or we map β backward to H_{d-r}^{2n-i} and compare it with α : $\langle h_r^{i,j}(\alpha), \beta \rangle = \langle \alpha, h_{d-r}^{2n-j, 2n-i}(\beta) \rangle = \text{card}(a \cap b) \pmod{2}$, where a and b are representatives of α and β that intersect transversally. To see this, we need to check the cases illustrated in Figure 5. In the case $i \leq j \leq n$ on the left, a is a cycle in K_i , b is a relative cycle in (K, L_{n-j}) , and their intersection lies entirely in K_i . We can use a as a representative of the image of α in H_r^j and we can use b as a representative of the image of β in H_{d-r}^{2n-i} , so both pairings are computed from the same number of intersections and are therefore the same. The case $n < i \leq j$ on the right is similar. In the case $i \leq n < j$ in the middle, a is a cycle in K_i and b is a cycle in K_{2n-j} , with intersection in the smaller of the two complexes. Again we can use the same representations




Figure 5: From left to right: Cycles in closed and open sublevel sets of the torus for $i \leq j \leq n$, $i \leq n < j$, $n < i \leq j$. The cycle a is a representative of the homology class α of K_i as well as its image in K_j . The cycle b is a representative of the homology class β of K_{2n-j} as well as its image in K_{2n-i} .

for the images and the number of intersections we count is again the same.

Duality. As before, we assume K is a triangulation of a d -manifold and we use the notation for complexes, homology groups, and induced homomorphisms introduced earlier. With this notation, the definition of pair groups given in (1) extends to $0 \leq i < j \leq 2n$.

DUALITY THEOREM (VERSION 1). A real-valued function on a d -manifold has non-degenerate pairings on the pair groups,

$$P : \mathbb{P}_r^{i,j} \times \mathbb{P}_{d-r}^{2n-j+1, 2n-i+1} \rightarrow \mathbb{Z}_2,$$

for all $0 \leq i < j \leq 2n$ and all dimensions r .

PROOF. To simplify notation, we write $k = 2n - i + 1$ and $\ell = 2n - j + 1$ so that the second pair group in the non-degenerate pairing is $\mathbb{P}_{d-r}^{\ell,k}$. The following diagram shows all groups and maps needed to define the two pair groups referred to in the claim:

$$\begin{array}{ccccccc} \mathbb{H}_r^{i-1} & \rightarrow & \mathbb{H}_r^i & \rightarrow & \mathbb{H}_r^{j-1} & \rightarrow & \mathbb{H}_r^j \\ \times & & \times & & \times & & \times \\ \mathbb{H}_{d-r}^k & \leftarrow & \mathbb{H}_{d-r}^{k-1} & \leftarrow & \mathbb{H}_{d-r}^\ell & \leftarrow & \mathbb{H}_{d-r}^{\ell-1} \end{array}$$

We may assume that both pair groups are non-trivial, else there is nothing to prove. We define the pairing P in the natural way, letting

$$\begin{aligned} \gamma &\in \text{im } h_r^{i,j-1} \cap \ker h_r^{j-1,j} \subseteq \mathbb{H}_r^{j-1}, \\ \delta &\in \text{im } h_{d-r}^{\ell,k-1} \cap \ker h_{d-r}^{k-1,k} \subseteq \mathbb{H}_{d-r}^{k-1} \end{aligned}$$

represent the non-zero classes $[\gamma] \in \mathbb{P}_r^{i,j}$ and $[\delta] \in \mathbb{P}_{d-r}^{\ell,k}$. Then choose δ' in \mathbb{H}_{d-r}^ℓ whose image in \mathbb{H}_{d-r}^{k-1} is δ and set

$$P([\gamma], [\delta]) = \langle \gamma, \delta' \rangle.$$

By compatibility, the value does not depend on the choice of δ' and the pairing is therefore well-defined. Similarly, we

can see that the pairing is bilinear and symmetric. It remains to prove that it is also non-degenerate. Since $\langle \cdot, \cdot \rangle$ is non-degenerate for all vertically arranged pairs, we can choose β in \mathbb{H}_{d-r}^ℓ such that $\langle \gamma, \beta \rangle \neq 0$ and $\langle \alpha, \beta \rangle = 0$ for all α in the image of $h_r^{i-1,j-1}$. By compatibility, this latter condition implies $\langle \alpha'', \beta'' \rangle = 0$ for all α'' in \mathbb{H}_r^{i-1} , where β'' is the image of β in \mathbb{H}_{d-r}^k . Non-degeneracy thus leaves only the possibility that $\beta'' = 0$. Hence, the image β' of β in \mathbb{H}_{d-r}^{k-1} represents an element of $\mathbb{P}_{d-r}^{\ell,k}$ for which $P([\gamma], [\beta']) \neq 0$. This shows that P is non-degenerate and completes the proof. \square

Version 1 of the Duality Theorem implies relationships between the persistence diagrams, which we express by writing Dgm^T for the reflection of Dgm along the main diagonal. In other words, $(x, y) \in \text{Dgm}_r(f)$ iff $(y, x) \in \text{Dgm}_r^T(f)$. Similarly, we use a superscript T to indicate reflections of the three sub-diagrams. By duality, the ordinary dimension r persistent classes correspond to the dimension $d - r$ relative persistent classes or, more formally, $\text{Ord}_r(f) = \text{Rel}_{d-r}^T(f)$. Indeed, if $i \leq j \leq n$ then a class in the pair group $\mathbb{P}_r^{i,j}$ is born by adding the lower star of the vertex u_i to K_{i-1} , and it is killed by adding the lower star of u_j to K_{j-1} . Symmetrically, a class in $\mathbb{P}_{d-r}^{\ell,k}$ is born by adding the upper star of $u_{2n-\ell+1} = u_j$ to L_{n-j} , and it is killed by adding the upper star of $u_{2n-k+1} = u_i$ to L_{n-i} . Similarly, $\text{Ext}_r(f) = \text{Ext}_{d-r}^T(f)$ and $\text{Rel}_r(f) = \text{Ord}_{d-r}^T(f)$. We combine the three cases into one statement.

DUALITY THEOREM (VERSION 2). A real-valued function f on a d -manifold has persistence diagrams that are reflections of each other, $\text{Dgm}_r(f) = \text{Dgm}_{d-r}^T(f)$, for all dimensions r .

Symmetry. As before, K is a triangulation of a d -manifold and f is defined by a real-valued function on the vertex set. We claim that duality implies that persistence is symmetric in the sense that f and $-f$ give the same diagrams up to reflections and dimensions. However, this time we use the superscript R to indicate reflection across the minor diagonal, mapping a point (x, y) to $(-y, -x)$, and the superscript 0 to indicate reflection through the origin, mapping (x, y) to $(-x, -y)$.

SYMMETRY THEOREM. For a real-valued function f on a d -manifold, we have

$$\begin{aligned} \text{Ord}_r(f) &= \text{Ord}_{d-r-1}^R(-f), \\ \text{Ext}_r(f) &= \text{Ext}_{d-r}^0(-f), \\ \text{Rel}_r(f) &= \text{Rel}_{d-r=1}^R(-f), \end{aligned}$$

for all dimensions r .

PROOF. To see this symmetry, we note that the vertex ordering defined by $-f$ simply switches K_i and L_{n-i} . There are then three sets of equalities to consider.

For the first set of equalities, let (f_i, f_j) be a point in the ordinary sub-diagram, $i \leq j \leq n$. When a dimension r homology class is born at K_i and dies entering K_j , duality provides a dimension $d - r$ relative homology class that is born at $K_\ell = (K, L_{n-j+1})$ and dies entering $K_k = (K, L_{n-i+1})$. This class is represented by the point (f_j, f_i) in $\text{Rel}_{d-r}(f)$. The relative class born at K_ℓ is Case D.2 in the classification of the Section 4 so it is accompanied by the birth of a dimension $d - r - 1$ class in L_{n-j+1} . Similarly, the relative class that dies entering K_k is Case D.1, so the dimension $d - r - 1$ class born at L_{n-j+1} dies entering L_{n-i+1} . This class is therefore represented by the point $(-f_j, -f_i)$ in $\text{Ord}_{d-r-1}(-f)$, proving the first set of equalities.

For the second set of equalities, let (f_i, f_j) be a point in the extended sub-diagram, $i \leq n < j$. A dimension r homology class is born at K_i and dies entering $K_j = (K, L_{j-n})$. The Duality Theorem then provides a dimension $d - r$ relative homology class born at K_ℓ and dying entering $K_k = (K, L_{n-i+1})$. We note that the class born at K_ℓ is Case D.3, so it is accompanied by a dimension $d - r$ absolute homology class and the same is true for the class that dies entering K_k . This class is represented by the point $(-f_i, -f_j)$ in $\text{Ext}_{d-r}(-f)$, thus proving the second set of inequalities.

The argument for the third set of equalities is the same, and we have again a shift in dimension, same as in the first case. \square

Morse function example. Assuming f is a Morse function, each critical point belongs to two pairs, one for the ascending and one for the descending direction. The Duality Theorem implies that the two pairs are the same, so we have a perfect matching. The Symmetry Theorem implies that $-f$ gives the same perfect matching. We follow up with an illustration of duality and symmetry, letting \mathbb{M} be the 2-manifold sketched in Figure 6 and $f : \mathbb{M} \rightarrow \mathbb{R}$ measure height above a horizontal base plane. In the example, f is Morse and has three minima, nine saddles, and two maxima. Since \mathbb{M} is orientable, this implies its genus is three. Hence, there are eight essential homology classes, the component represented by the global minimum, a , six 1-cycles represented by the saddles u, v, w, U, V, W , and the 2-cycle represented by the global maximum, A . The pairs formed by extended persistence define points making up the extended sub-diagrams, (a, A) for dimension 0, $(u, U), (v, V), (w, W), (W, w), (V, v), (U, u)$ for dimension 1, and (A, a) for dimension 2. The remaining pairs correspond to points in the other sub-diagrams, $(b, B), (c, C)$ and (d, D) for the ordinary sub-diagrams of dimensions 0 and 1, and $(C, c), (B, b)$ and (D, d) for the relative sub-diagrams of dimensions 2 and 1. Figure 7 overlays the sub-diagrams showing all fourteen points as circle-shaped dots in one picture. By the Duality Theorem, this picture is symmetric with respect to the main diagonal. As we switch from f to $-f$, minima become maxima, saddles remain saddles, and maxima become minima. By the Symmetry Theorem, the pairs

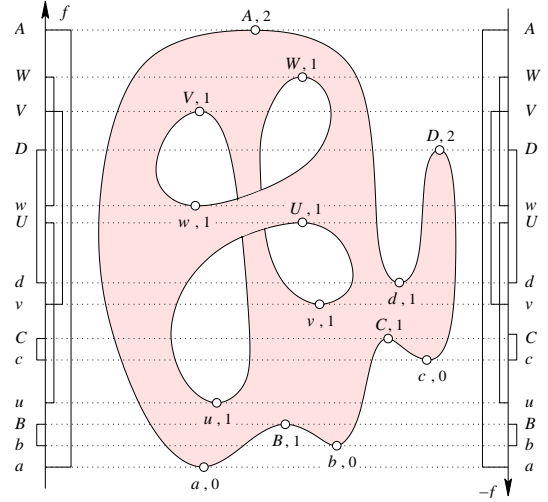


Figure 6: Two Morse functions on a 2-manifold, measuring height above and below a base plane.

remain the same except the coordinates become negative and switch their order. In other words, we get the diagrams of $-f$ by reflection along the minor diagonal. If we overlay the diagrams of f and $-f$ we get a picture that is symmetric with respect to both diagonals, as in Figure 7. Note, however, that the type of the sub-diagram and the dimension may change as we reflect a point across the major or the minor diagonal.

8 Discussion

The main contributions of this paper is the extension of ordinary persistence to essential homology classes. In this section, we briefly discuss the motivating application to elevation functions and list a few questions raised by the reported work.

Elevation. We recall that the elevation function of a surface embedded in \mathbb{R}^3 has been introduced and studied in [1]. The local maxima of this function have been applied to coarse protein docking in [11]. Here we briefly comment on how the results of this paper can be used to generalize elevation beyond \mathbb{R}^3 .

Let \mathbb{M} be a codimension 1 submanifold of \mathbb{R}^{d+1} , that is, \mathbb{M} is the smooth embedding of a d -manifold in \mathbb{R}^{d+1} . Each point $x \in \mathbb{M}$ has two unit normals, u_x and $-u_x$. Letting $f_u : \mathbb{M} \rightarrow \mathbb{R}$ be the height function in direction $u \in \mathbb{S}^d$, the point is critical iff u is equal to u_x or to $-u_x$. Assume $u = u_x$ and let $y \in \mathbb{M}$ be the critical point of f_u that is paired with x by (extended) persistence. By the Symmetry Theorem, x and y are also paired for $u = -u_x$. The function

$$\text{Elevation} : \mathbb{M} \rightarrow \mathbb{R}$$

defined by $\text{Elevation}(x) = |f_u(x) - f_u(y)|$ therefore makes sense. Similar to the case $d = 2$, the function is smooth almost everywhere but there are measure-zero violations even




Figure 7: Overlay of the persistence diagrams of f and $-f$ of all dimensions. We draw circles for f and squares for $-f$, color-coding ordinary, extended, relative sub-diagrams white, gray, black, and marking the dots by the dimension of their diagrams. The moments of birth and death of points above the main diagonal are marked by the corresponding critical points. We have $f(a) = -f(A)$ and therefore two double points on the minor diagonal.

for continuity. It is possible to do surgery on \mathbb{M} to make Elevation continuous so local maxima can be defined. For every finite dimension d , there are finitely many generic cases, namely maxima determined by $2 \leq k + \ell \leq d + 2$ points on \mathbb{M} . Counting all possibilities that satisfy $1 \leq k \leq \ell$ we get $(d^2 + 4d + 4)/4$ cases for even d and $(d^2 + 4d + 3)/4$ cases for odd d . For surfaces embedded in \mathbb{R}^3 , elevation is one of few known functions essentially different from variations of curvature that capture shape information. It would be interesting to find applications in which the shape of codimension 1 submanifolds in Euclidean space of dimension beyond 3 plays a role.

Open questions. The algorithm for computing extended persistence applies to triangulated manifolds as well as to general simplicial complexes, but the Duality and Symmetry Theorems do not. As a consequence, the information we get for general simplicial complexes is more difficult to interpret. Already for the height function of the figure-8 space shown in Figure 8, we have pairs defining a graph that is more complicated than a matching.

- Is there a non-manifold counterexample to the weaker version of the Symmetry Theorem that claims we get the same collection of pairs for f and for $-f$?
- Can we quantify the extent to which a diagram violates the Duality and Symmetry Theorems and this way mea-




Figure 8: The homology of the sublevel set changes at three values. Extended persistence forms three pairs, (A, C) , (C, B) , (B, A) . We get the same three pairs, each in opposite order, if we measure height from top to bottom.

sure how close a complex is to being a manifold?

The Stability Theorem can be strengthened to apply to the three sub-diagrams individually. In particular, we have $d_B(\text{Ext}_r(f), \text{Ext}_r(g)) \leq \|f - g\|_\infty$, for each dimension r .

- Can we extend this result to functions f and g defined on different spaces and use the bottleneck distance to measure how similar or different these spaces are?

The extended sub-diagram needs to be interpreted differently from the ordinary and the relative sub-diagrams. For instance, we can have points on the diagonal representing non-negligible topological features. An example is the projective plane, $\mathbb{R}P^2$, with $\beta_0 = \beta_1 = \beta_2 = 1$. The essential dimension 1 homology class is born when we pass the saddle at height x in the upward pass and it dies again when we pass the same saddle in the downward pass. It follows that (x, x) is the only point in the dimension 1 extended sub-diagram. Another such example is provided by $\mathbb{C}P^2$, an orientable 4-manifold. We may therefore reconsider the definition of bottleneck distance as it applies to extended sub-diagrams, for example by not adding extra points on the diagonal.

Acknowledgments

The authors thank Dmitry Morozov for insightful discussions on the topic of this paper.

References

- [1] P. K. AGARWAL, H. EDELSBRUNNER, J. HARER AND Y. WANG. Extreme elevation on a 2-manifold. *Discrete Comput. Geom.* **36** (2006), 553–572.
- [2] D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), 103–120.
- [3] D. COHEN-STEINER, H. EDELSBRUNNER AND D. MOROZOV. Vines and vineyards by updating persistence in linear time. In “Proc. 22nd Ann. Sympos. Comput. Geom., 2006”, 119–126.
- [4] M. L. CONNOLLY. Shape complementarity at the hemo-globin albl subunit interface. *Biopolymers* **25** (1986), 1229–1247.

- [5] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511–533.
- [6] R. GONZÁLES-DÍAZ AND P. REAL. On the cohomology of 3D digital images. *Discrete Appl. Math.* **147** (2005), 245–263.
- [7] N. GELFAND, N. J. MITRA, L. J. GUIBAS AND H. POTTMANN. Robust global registration. In “Proc. Sympos. Geom. Process., 2005”, 197–206.
- [8] M. HILAGA, Y. SHINAGAWA, T. KOMURA AND T. L. KUNII. Topology matching for fully automatic similarity estimation of 3D shapes. In “Computer Graphics Proc., SIGGRAPH 2001”, 203–212.
- [9] J. R. MUNKRES. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, California, 1984.
- [10] G. REEB. Sur les points singuliers d’une forme de Pfaff complètement intégrable ou d’une fonction numérique. *Comptes Rendus de L’Académie des Séances* **222** (1946), 847–849.
- [11] Y. WANG, P. K. AGARWAL, P. BROWN, H. EDELSBRUNNER AND J. RUDOLPH. Coarse and reliable geometric alignment for protein docking. In “Proc. Pacific Sympos. Biocomput., 2005”, 65–75.
- [12] A. ZOMORODIAN AND G. CARLSSON. Computing persistent homology. *Discrete Comput. Geom.* **33** (2005), 249–274.