

Persistent Homology for Kernels, Images, and Cokernels ^{*}

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

Abstract

Motivated by the measurement of local homology and of functions on noisy domains, we extend the notion of persistent homology to sequences of kernels, images, and cokernels of maps induced by inclusions in a filtration of pairs of spaces. Specifically, we note that persistence in this context is well defined, we prove that the persistence diagrams are stable, and we explain how to compute them.

1 Introduction

Natural phenomena are often modeled in terms of spaces and functions on these spaces. We argue that it is almost always more appropriate to use a multi-scale hierarchy instead of a single space. One reason is the prevalent multi-scale organization we find in nature, another is that the data we gather about nature is necessarily incomplete and requires interpolation. The multi-scale aspect helps bridging the gap between the data about natural phenomena and the idealized mathematical concepts we use for exploration. In particular, we consider homology groups, which are algebraic structures that define and count holes in a topological space [12]. Their multi-scale extensions are persistent homology groups introduced in [10, 14]. Similar to homology which not only counts but also defines, persistent homology not only measures but also creates the hierarchy.

In all previous settings, the hierarchy is defined by a nested sequence of spaces, $\mathbb{X}_0 \subseteq \mathbb{X}_1 \subseteq \dots \subseteq \mathbb{X}_m$, and persistent homology arises from considering the corresponding sequence of homology groups, $H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \dots \rightarrow$

$H(\mathbb{X}_m)$, connected from left to right by homomorphic maps induced by inclusion. Persistence tracks when a homology class is born and when it dies. This can also be done for an arbitrary sequence of vector spaces connected by homomorphic maps. Motivation for studying such more general sequences is derived from recent investigations in computational topology. First, Bendich et al. describe a multi-scale assessment of local homology for the purpose of reconstructing a stratified space from a point sample [2]. We will see how sequences of kernels can be used to refine their construction. Second, we use sequences of images to introduce a notion of persistence that filters out noise induced by imprecise specifications of domains. This contrasts standard persistence which can handle imprecise function values but not imprecise domains. As an application, we will approximate the persistence diagram of a function knowing only its values at a finite set of points. The main contributions of this paper are two-fold:

- an algorithm that computes the persistence diagrams of sequences of kernels, images, and cokernels in time at most cubic in the size of the simplicial complexes representing the data;
- applications of the algebraic and algorithmic results to measuring local homology and to approximating persistence diagrams of noisy functions on noisy domains.

Outline. The remainder of this paper is structured as follows. Section 2 introduces the algebra of persistent homology including its extension to sequences of kernels, images and cokernels. Section 3 explains the algorithms for computing the corresponding persistence diagrams for a nested sequence of pairs of spaces and Section 4 proves their correctness. Section 5 presents the two applications of our algebraic and algorithmic results. Section 6 concludes the paper.

2 Algebra

Beginning with a review of persistent homology, this section extends this concept to sequences of kernels, images, and

^{*}This research is partially supported by the Defense Advanced Research Projects Agency (DARPA) under grants HR0011-05-1-0007 and HR0011-05-1-0057 and by CNRS under grant PICS-3416.

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

[‡]Departments of Computer Science and of Mathematics, Duke University, Durham, Berlin Mathematical School, Berlin, Germany, and Gemagic, Research Triangle Park, North Carolina, USA.

[§]Departments of Mathematics and of Computer Science, Duke University, Durham, North Carolina, USA.

[¶]Department of Computer Science, Duke University, Durham, North Carolina, USA.

cokernels. It also proves that the persistence diagrams of these extensions are stable.

Persistent homology. This concept is a recent addition to classical homology theory and was originally introduced for ordered simplicial complexes [10]. We follow the exposition in [6] in which we have a topological space \mathbb{X} and a continuous function $f : \mathbb{X} \rightarrow \mathbb{R}$. The *sublevel set* defined by $a \in \mathbb{R}$ consists of all points with function value at most the threshold, $\mathbb{X}_a = f^{-1}(-\infty, a]$. We use the algebraic language of homology theory to characterize how \mathbb{X}_a is connected, see e.g. [12]. Adding chains with modulo-2 arithmetic, we write $H_p(\mathbb{X}_a)$ for the *dimension p homology group* over $\mathbb{Z}/2\mathbb{Z}$ of \mathbb{X}_a and $H(\mathbb{X}_a) = (\dots, H_p(\mathbb{X}_a), H_{p+1}(\mathbb{X}_a), \dots)$ for the infinite series obtained by collecting the groups for all dimensions. Of course, only the groups for p between 0 and the dimension of \mathbb{X}_a are possibly non-trivial. To simplify language, we will often ignore the difference between a single homology group and the entire series. Given $a \leq b$, the inclusion between the sublevel sets, $\mathbb{X}_a \subseteq \mathbb{X}_b$, induces a homomorphism, $\mathbf{f}_{a,b} : H(\mathbb{X}_a) \rightarrow H(\mathbb{X}_b)$. For $a = b$ this is an isomorphism and for $a < b$ it may or may not be an isomorphism. A value $a \in \mathbb{R}$ is a *homological critical value* of f if there is no sufficiently small $\varepsilon > 0$ for which $\mathbf{f}_{a-\varepsilon, a+\varepsilon}$ is an isomorphism. We assume that f is *tame*, that is, it has only finitely many critical values and every sublevel set has only finite rank homology groups.

Let $a_1 < a_2 < \dots < a_m$ be the critical values of f and consider an interleaved sequence $s_{i-1} < a_i < s_i$ for all i . This gives a sequence of spaces, $\mathbb{X}_0 \subseteq \mathbb{X}_1 \subseteq \dots \subseteq \mathbb{X}_m = \mathbb{X}$, where we simplify notation by writing $\mathbb{X}_i = \mathbb{X}_{s_i}$, and a corresponding sequence of homology groups connected by homomorphisms,

$$H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \dots \rightarrow H(\mathbb{X}_m).$$

Persistence concerns itself with the history of individual homology classes within this sequence. Specifically, a class γ in $H(\mathbb{X}_i)$ is *born* at a_i if it is not in the image of $\mathbf{f}_{i-1,i} = \mathbf{f}_{s_{i-1}, s_i}$. More precisely, an entire coset is born at a_i . Furthermore, if γ is born at a_i we say it *dies entering* a_j if $\mathbf{f}_{i,j-1}(\gamma)$ is not contained in the image of $\mathbf{f}_{i-1,j-1}$ but $\mathbf{f}_{i,j}(\gamma)$ is contained in the image of $\mathbf{f}_{i-1,j}$. The images of the maps $\mathbf{f}_{i,j}$ are referred to as *persistent homology groups* since they consist of all homology groups born at or before a_i that live beyond a_j . Nothing about the definition of birth and death is specific to homology groups. In other words, persistence makes perfect sense for any sequence of vector spaces connected by homomorphisms.

It is convenient to represent the fact that γ is born at a_i and dies entering a_j by drawing the point (a_i, a_j) in the two-dimensional plane. By collecting the points for all p -dimensional classes we get the *dimension p persistence diagram* which we denote as $\text{Dgm}_p(p)$. Since birth necessarily happens before death all points lie above the diagonal. It is also possible that a class γ is born at a_i but does not die since it represents a class of $\mathbb{X}_m = \mathbb{X}$. In this case, we draw

γ as the point (a_i, ∞) in the diagram. For technical reasons that will become clear later, we consider all points on the diagonal to be part of the persistence diagram. Similar to homology groups we get a diagram for each dimension and we write $\text{Dgm}(f)$ for the infinite series of diagrams. Same as for homology groups and for the maps between them we will simplify language by ignoring the difference between a single diagram and an entire series.

Kernels, images, and cokernels. For the extension of persistence to kernels, images, and cokernels we consider two functions, $f : \mathbb{X} \rightarrow \mathbb{R}$ and a majorizing function $g : \mathbb{Y} \rightarrow \mathbb{R}$ defined on a subspace $\mathbb{Y} \subseteq \mathbb{X}$, that is, $f(y) \leq g(y)$ for all $y \in \mathbb{Y} \subseteq \mathbb{X}$. Assuming both functions are tame, we order the collection of critical values of f and g and interleave them with a sequence of real values s_i . The corresponding sequences of sublevel sets give rise to two parallel sequences of homology groups,

$$\begin{array}{ccccccc} H(\mathbb{X}_0) & \rightarrow & H(\mathbb{X}_1) & \rightarrow & \dots & \rightarrow & H(\mathbb{X}_m) \\ & & \uparrow j_0 & & \uparrow j_1 & & \dots & & \uparrow j_m \\ H(\mathbb{Y}_0) & \rightarrow & H(\mathbb{Y}_1) & \rightarrow & \dots & \rightarrow & H(\mathbb{Y}_m), \end{array}$$

where $\mathbb{X}_i = f^{-1}(-\infty, s_i]$ and $\mathbb{Y}_i = g^{-1}(-\infty, s_i]$. The two sequences are connected by homomorphisms $j_i : H(\mathbb{Y}_i) \rightarrow H(\mathbb{X}_i)$ induced by the inclusions $\mathbb{Y}_i \subseteq \mathbb{X}_i$. We call this the *two function setting*, in contrast to the more special *one function setting* in which g is the restriction of f to \mathbb{Y} . More about the relationship between the two settings later. We are interested in the kernels, images, and cokernels of the connecting homomorphisms,

$$\begin{aligned} \ker j_i &= \{\gamma \in H(\mathbb{Y}_i) \mid j_i(\gamma) = 0 \in H(\mathbb{X}_i)\}; \\ \text{im } j_i &= \{j_i(\gamma) \in H(\mathbb{X}_i) \mid \gamma \in H(\mathbb{Y}_i)\}; \\ \text{cok } j_i &= H(\mathbb{X}_i)/\text{im } j_i. \end{aligned}$$

We note that the “coimage” of the map j_i , in symbols $H(\mathbb{Y}_i)/\ker j_i$, is isomorphic to the image of this map, and therefore does not deserve any special attention. The square defined by the four homology groups $H(\mathbb{Y}_i), H(\mathbb{Y}_{i+1}), H(\mathbb{X}_i)$, and $H(\mathbb{X}_{i+1})$ commutes. It follows that the inclusion $\mathbb{Y}_i \subseteq \mathbb{Y}_{i+1}$ induces a homomorphism $\ker j_i \rightarrow \ker j_{i+1}$. Similarly, the inclusion $\mathbb{X}_i \subseteq \mathbb{X}_{i+1}$ induces a homomorphism $\text{im } j_i \rightarrow \text{im } j_{i+1}$ and another homomorphism $\text{cok } j_i \rightarrow \text{cok } j_{i+1}$. We thus get sequences of kernels, images, and cokernels,

$$\begin{aligned} \text{Ker}(g \rightarrow f) &: \ker j_0 \rightarrow \ker j_1 \rightarrow \dots \rightarrow \ker j_m; \\ \text{Im}(g \rightarrow f) &: \text{im } j_0 \rightarrow \text{im } j_1 \rightarrow \dots \rightarrow \text{im } j_m; \\ \text{Cok}(g \rightarrow f) &: \text{cok } j_0 \rightarrow \text{cok } j_1 \rightarrow \dots \rightarrow \text{cok } j_m, \end{aligned}$$

all connected from left to right by homomorphisms. Homology classes are born and die in these sequences same as in the sequences of homology groups. We can therefore define *persistent kernels*, *persistent images*, and *persistent cokernels*

as well as construct the corresponding persistence diagrams, which we denote as $\text{Dgm}(\ker g \rightarrow f)$, $\text{Dgm}(\text{im } g \rightarrow f)$, and $\text{Dgm}(\text{cok } g \rightarrow f)$.

Birth-death combinations. We consider the generic case in which changes happen one at a time. An event thus corresponds to a birth, a death, or no change in the kernel, in the image, and in the cokernel, giving rise to 27 different combinations. But the ranks of these groups are not independent, that is,

$$\begin{aligned} \text{rank } \ker j_i + \text{rank } \text{im } j_i &= \text{rank } H(\mathbb{Y}_i); \\ \text{rank } \text{im } j_i + \text{rank } \text{cok } j_i &= \text{rank } H(\mathbb{X}_i), \end{aligned}$$

for all i . We can therefore relate the births and deaths in the three sequences using the births and deaths in the sequences of homology groups of the \mathbb{Y}_i and of the \mathbb{X}_i . The first equation eliminates two of the nine combinations for kernels and images. Another combination is eliminated by $\ker j_i$ being a subgroup of $H(\mathbb{Y}_i)$, hence a death in the kernel implies a death in the homology group. Case A occurs for

Case	$\ker j_i$	$\text{im } j_i$	$\text{cok } j_i$	$H(\mathbb{Y}_i)$	$H(\mathbb{X}_i)$
A	birth	death	—	—	death
B	—	birth	—	birth	birth
C	—	—	birth	—	birth
D	—	—	death	—	death
E	—	death	—	death	death
F	death	—	birth	death	birth
P	birth	—	—	birth	—
Q	—	birth	death	birth	—

Table 1: The eight cases in the two function setting relating kernels and images, and images and cokernels. Except for Cases P and Q they also occur in the one function setting.

example when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1} = \mathbb{Y}_i$ is a circle and \mathbb{X}_i is obtained by adding a spanning disk. Case B occurs when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1}$ is a point and $\mathbb{Y}_i = \mathbb{X}_i$ is obtained by adding an arc that completes the point to a circle. Case C occurs when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1} = \mathbb{Y}_i$ is a point and \mathbb{X}_i is again obtained by adding an arc that forms a circle. We also retain ranks in Case D which occurs when \mathbb{X}_{i-1} is a circle, $\mathbb{Y}_{i-1} = \mathbb{Y}_i$ is a point on this circle, and \mathbb{X}_i is obtained by adding a spanning disk. Case E occurs when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1}$ is a circle and $\mathbb{Y}_i = \mathbb{X}_i$ is obtained by adding a spanning disk. Case F occurs when \mathbb{X}_{i-1} is a disk, \mathbb{Y}_{i-1} is its boundary circle, and we get \mathbb{Y}_i and \mathbb{X}_i by adding another spanning disk to both spaces. Finally, Case P occurs when $\mathbb{X}_{i-1} = \mathbb{X}_i$ is a disk, \mathbb{Y}_{i-1} is a point of its boundary circle, and \mathbb{Y}_i is obtained by adding the rest of the circle. This last case happens in the two function setting but not in the one function setting because it requires points that are added to the sublevel set of g strictly after they are added to the sublevel set of f .

Similarly, the second equation eliminates two of the nine combinations for images and cokernels. Another combination is eliminated by $\text{im } j_i$ being a subgroup of $H(\mathbb{X}_i)$, hence a death in the image implies a death in the homology group.

Cases A to F and P have been described above. Case Q occurs when $\mathbb{X}_{i-1} = \mathbb{Y}_i = \mathbb{X}_i$ is a circle and \mathbb{Y}_{i-1} is a point on that circle. Cases P and Q do not happen in the one function setting in which every change has a non-zero effect on the rank of the homology group of \mathbb{X}_i .

Mapping cylinder. We reduce the two function setting to the one function setting using a construction that will be exploited by the algorithm described in Section 3. Specifically, the *mapping cylinder* of the pair $\mathbb{Y} \subseteq \mathbb{X}$ is the space $\mathbb{X}' = \mathbb{X} \cup (\mathbb{Y} \times [0, 1])$ obtained by gluing $\mathbb{Y} \subseteq \mathbb{X}$ to $\mathbb{Y} \times \{0\} \subseteq \mathbb{Y} \times [0, 1]$. The function $f' : \mathbb{X}' \rightarrow \mathbb{R}$ agrees with f on \mathbb{X} and with g on $\mathbb{Y}' = \mathbb{Y} \times \{1\}$, linearly interpolating in between, that is, $f'(x) = f(x)$ for every $x \in \mathbb{X}$ and $f'(y, t) = (1 - t)f(y) + tg(y)$ for every $y \in \mathbb{Y}$ and every $t \in [0, 1]$. The pair of functions f' and $g' = f'|_{\mathbb{Y}'}$ induces homomorphisms $j'_i : H(\mathbb{Y}'_i) \rightarrow H(\mathbb{X}'_i)$. The corresponding sequences of kernels, images, and cokernels are

$$\begin{aligned} \text{Ker}(g' \rightarrow f') &: \ker j'_0 \rightarrow \ker j'_1 \rightarrow \dots \rightarrow \ker j'_m; \\ \text{Im}(g' \rightarrow f') &: \text{im } j'_0 \rightarrow \text{im } j'_1 \rightarrow \dots \rightarrow \text{im } j'_m; \\ \text{Cok}(g' \rightarrow f') &: \text{cok } j'_0 \rightarrow \text{cok } j'_1 \rightarrow \dots \rightarrow \text{cok } j'_m. \end{aligned}$$

We claim that they contain the same information as the sequences $\text{Ker}(g \rightarrow f)$, $\text{Im}(g \rightarrow f)$, and $\text{Cok}(g \rightarrow f)$.

MAPPING CYLINDER LEMMA. The pairs of functions f, g and f', g' define the same persistence diagrams for kernels, images, and cokernels: $\text{Dgm}(\text{grp } g \rightarrow f) = \text{Dgm}(\text{grp } g' \rightarrow f')$ for $\text{grp} \in \{\ker, \text{im}, \text{cok}\}$.

Details are omitted in this abridged conference version of the paper.

Stability. An important property of the persistence diagrams is their stability originally proved in [6]. More precisely, the bottleneck distance between the diagrams of two functions f and f'' is bounded from above by the difference between the functions, $d_B(\text{Dgm}(f), \text{Dgm}(f'')) \leq \|f - f''\|_\infty$. Here d_B is the maximum of $\|u - \gamma(u)\|_\infty$, where u is a point in the diagram of f and γ is a dimension-preserving bijection between the diagrams of f and of f'' . Recall that the points on the diagonal belong to the diagrams and can therefore be used in the effort to find a matching γ that minimizes the length of the longest edge. The proof of stability given in [6] can be adapted to the setting in this paper. Specifically, we consider the maps $j_a : H(\mathbb{Y}_a) \rightarrow H(\mathbb{X}_a)$ and $j''_{a+\varepsilon} : H(\mathbb{Y}''_{a+\varepsilon}) \rightarrow H(\mathbb{X}''_{a+\varepsilon})$, where ε is the larger of the two differences between functions, $\|f - f''\|_\infty$ and $\|g - g''\|_\infty$, and $\mathbb{Y}_a, \mathbb{X}_a, \mathbb{Y}''_{a+\varepsilon}, \mathbb{X}''_{a+\varepsilon}$ are the sublevel sets of g, f, g'', f'' for thresholds a and $a + \varepsilon$. To adapt the proof we need that the maps induced by the inclusions $\mathbb{Y}_a \subseteq \mathbb{Y}''_{a+\varepsilon}$ and $\mathbb{X}_a \subseteq \mathbb{X}''_{a+\varepsilon}$ send the kernel of j_a into the kernel of $j''_{a+\varepsilon}$. But this follows from the commutativity of the diagram involving groups $H(\mathbb{Y}_a), H(\mathbb{Y}''_{a+\varepsilon}), H(\mathbb{X}_a)$, and $H(\mathbb{X}''_{a+\varepsilon})$ with maps induced by the inclusions of respective spaces. Similarly, we need that the inclusions $\mathbb{Y}''_a \subseteq \mathbb{Y}_{a+\varepsilon}$ and $\mathbb{X}''_a \subseteq \mathbb{X}_{a+\varepsilon}$ send $\ker j''_a$ into $\ker j_{a+\varepsilon}$ which follows by

symmetry. With this property the original proof of stability goes through and we refer to [6] for details. The arguments for the images and the cokernels are the same and we state the results.

STABILITY THEOREM. Let $f, f'' : \mathbb{X} \rightarrow \mathbb{R}$ and $g, g'' : \mathbb{Y} \rightarrow \mathbb{R}$ with $f(y) \leq g(y)$ and $f''(y) \leq g''(y)$ for every $y \in \mathbb{Y} \subseteq \mathbb{X}$ and $\varepsilon = \max\{\|f - f''\|_\infty, \|g - g''\|_\infty\}$. Then the bottleneck distance between the persistence diagrams is bounded from above by the difference between the functions:

$$d_B(\text{Dgm}(\text{grp } g \rightarrow f), \text{Dgm}(\text{grp } g'' \rightarrow f'')) \leq \varepsilon,$$

for $\text{grp} \in \{\ker, \text{im}, \text{cok}\}$, provided f, g, f'' , and g'' are continuous and tame and there is a triangulation of \mathbb{X} in which \mathbb{Y} arises as a subcomplex.

3 Algorithms

In this section, we describe the algorithms for computing the persistence diagrams of the sequences of kernels, images, and cokernels. At their core is the reduction of a matrix as introduced in [8] which we describe first.

Matrix reduction. To get a finite representation of the data we now substitute a function on a simplicial complex for the continuous function on a topological space. In particular, we assume a (finite) simplicial complex K and an injective function $f : K \rightarrow \mathbb{R}$ that maps each simplex to a real number. The only additional requirement is $f(\sigma) < f(\tau)$ whenever σ is a face of τ . Equivalently, $K_a = f^{-1}(-\infty, a]$ is a subcomplex of K for every $a \in \mathbb{R}$. We index the simplices of K such that $f(\sigma_1) < f(\sigma_2) < \dots < f(\sigma_m)$ and let K_i be the complex consisting of the first i simplices in the sequence. To compute persistence for the sequence of complexes K_i we let D be the m -by- m incidence matrix defined by $D[\ell, i] = 1$ if σ_ℓ is a co-dimension one face of σ_i and $D[\ell, i] = 0$ otherwise. We *reduce* D using left-to-right modulo-2 column additions until the lowest one of every non-zero column is in a unique row. Initializing R and V to the incidence and the identity matrices and letting $low(i)$ be the row index of the lowest one in column i of R , or 0 if the entire column is zero, we can formalize the algorithm as follows.

```

R = D; V = I;
for i = 1 to m do
  while  $\exists k < i$  with  $low(k) = low(i) \neq 0$  do
    add column  $k$  to column  $i$  in  $R$  as well as in  $V$ .

```

Equivalently, the reduced matrix is obtained by multiplying the incidence matrix from the right with an upper-triangular matrix, $R = DV$, such that the map from the non-zero columns of R to the row indices of their lowest ones is injective. As proved in [8], R is not unique but the map is. By construction, the rows of R and V correspond to individual simplices, same as the rows of D , but the columns of R and V correspond to chains. Specifically, column i of

R stores the boundary of the chain stored in column i of V . We call σ_i *positive* if its addition to K_{i-1} gives birth to a homology class. Equivalently, column i of V stores a cycle and column i of R is zero. Symmetrically, we call σ_i *negative* if its addition to K_{i-1} gives death to a homology class. Equivalently, column i of V stores a chain that is not a cycle and column i of R is non-zero. The significance of the lowest one in this column of R is that the negative σ_i is paired with the positive σ_ℓ , with $\ell = low(i)$, which gives birth to the class that σ_i kills.

The persistence diagrams of f can be obtained from the reduced matrix. Specifically, each lowest one, $\ell = low(i)$, corresponds to a pair of simplices, σ_ℓ, σ_i , and we draw the point $(f(\sigma_\ell), f(\sigma_i))$ in the diagram whose dimension is the same as that of σ_ℓ .

Partial and reordered matrices. We prepare the computation of the persistence diagrams by reducing five matrices. By the Mapping Cylinder Lemma, we can restrict ourselves to the one function setting. We therefore assume two simplicial complexes, $L \subseteq K$, let $f : K \rightarrow \mathbb{R}$ be an injective function whose sublevel sets are subcomplexes of K , and let g be the restriction of f to L . We write D_f for the incidence matrix of K whose rows and columns are ordered by f . Similarly, we write D_g for the incidence matrix of L whose rows and columns are ordered by g .

Step 1 Reduce the two incidence matrices to get $R_f = D_f V_f$ and $R_g = D_g V_g$.

Step 2 Reorder the rows of D_f leaving the columns untouched to get a new matrix D_{im} . Specifically, its rows correspond to the simplices in L , ordered by g , followed by the simplices in $K - L$, ordered by f . Reduce the new matrix to get $R_{\text{im}} = D_{\text{im}} V_{\text{im}}$; see Figure 1.




Figure 1: Matrices computed in the reduction of the incidence matrix of K with reordered rows. The matrix V_{im} is upper-triangular with all ones in the diagonal.

Step 3 Delete some of the columns from V_{im} and reorder the rows to get a new matrix D_{ker} . Specifically, keep the columns that represent cycles and remove all others. Furthermore, reorder the rows so they correspond to the simplices in L , ordered by g , followed by the simplices in $K - L$, ordered by f . Finally, reduce the new matrix to get $R_{\text{ker}} = D_{\text{ker}} V_{\text{ker}}$.

Step 4 Starting again with D_f , replace some of the columns to get a new matrix D_{cok} . Specifically, substitute columns in V_g that represent cycles for the corresponding columns in D_f , adding zeros to compensate for the simplices in $K - L$, which are missing in V_g . Reduce the new matrix to get $R_{\text{cok}} = D_{\text{cok}} V_{\text{cok}}$.

We note that reducing D_f is redundant because the type information it furnishes is also available from R_{im} . We still use R_f because this clarifies which information is used where.

Births, deaths, and pairs. We use the reduced matrices to compute the persistence diagrams of the sequences of kernels, images, and cokernels. Specifically, R_f and R_g (and in one case R_{im}) decide which simplices give birth and which give death and R_{ker} , R_{im} , R_{cok} determine how the births match up with the deaths. We begin with the sequence of kernels and recall the relevant Cases A and F in Table 1.

Algorithm for kernels:

Birth. A simplex σ gives birth in $\text{Ker}(g \rightarrow f)$ iff $\sigma \in K - L$, σ is negative in R_f , and the lowest one in its column in R_{im} corresponds to a simplex in L .

Death. A simplex τ gives death in $\text{Ker}(g \rightarrow f)$ iff $\tau \in L$, τ is negative in R_g , and τ is positive in R_f . In this case, the lowest one in the column of τ in R_{ker} corresponds to a simplex $\sigma \in K - L$ that gives birth in $\text{Ker}(g \rightarrow f)$. Then (σ, τ) is a pair.

A dimension p homology class is given birth to in the kernel by a $(p+1)$ -simplex and it dies at the hand of another $(p+1)$ -simplex. The dimension p persistence diagram thus consists of all points $(f(\sigma), f(\tau))$ encoding pairs of $(p+1)$ -simplices identified in the Death case as well as all points $(f(\sigma), \infty)$ encoding unpaired $(p+1)$ -simplices identified in the Birth case. We continue with the sequence of images and recall the relevant Cases A, B, E in Table 1.

Algorithm for images:

Birth. A simplex σ gives birth in $\text{Im}(g \rightarrow f)$ iff $\sigma \in L$ and σ is positive in R_g .

Death. A simplex τ gives death in $\text{Im}(g \rightarrow f)$ iff τ is negative in R_f and the lowest one in its column in R_{im} corresponds to a simplex $\sigma \in L$. Then (σ, τ) is a pair.

Note that the Death case splits into Case A with $\tau \in K - L$ and Case E with $\tau \in L$. The dimension p persistence diagram consists of all points $(f(\sigma), f(\tau))$ encoding pairs of p - and $(p+1)$ -simplices identified in the Death case as well as points $(f(\sigma), \infty)$ encoding unpaired p -simplices identified in the Birth case. We continue with the sequence of cokernels and recall the relevant Cases C, F and D.

Algorithm for cokernels:

Birth. A simplex σ gives birth in $\text{Cok}(g \rightarrow f)$ iff σ is positive in R_f and it is either in $K - L$ or negative in R_g .

Death. A simplex τ gives death in $\text{Cok}(g \rightarrow f)$ iff τ is negative in R_f and the lowest one in its column in R_{im} corresponds to a simplex in $K - L$. In this case, the lowest one in the column of τ in R_{cok} corresponds to a simplex σ that gives birth in $\text{Cok}(g \rightarrow f)$. Then (σ, τ) is a pair.

The dimension p persistence diagram consists of all points $(f(\sigma), f(\tau))$ encoding pairs of p - and $(p+1)$ -simplices identified in the Death case as well as points $(f(\sigma), \infty)$ encoding unpaired p -simplices identified in the Birth case. The running time of the three algorithms is $O(m^3)$, same as the original persistence algorithm given in [10]. Furthermore, it is possible to extend the algorithm in [8] so that it maintains the reduced matrices in time $O(m)$ per transposition of contiguous simplices in the ordered sequences; details can be found in [11, Section 6.5] or Appendix A of the full version of this paper. This is the method of choice for computing the vineyard of a pair of 1-parameter families of functions f and g , as they arise in applications considered in Section 5.

4 Correctness

We prove the correctness of the algorithms inductively, by considering one simplex at a time. For each index i , we consider the *actual* births, deaths, and pairs that occur in the sequences up to j_i , and the *computed* births, deaths, and pairs reported by the algorithm working on the simplices up to σ_i . To prove that the corresponding sets are the same at the end, for $i = m$, we show they are the same throughout, for all i . We do this in two steps, first proving that the algorithms are necessary and second that they are sufficient. In other words, we first prove that the computed information is correct and second that it is complete.

Preparation. We begin with a few preliminary observations. Recall that Table 1 lists the possible combinations of births and deaths under the simplifying assumption that each group has at most one change happening at any one time. This is indeed the situation if we add individual simplices to a growing complex. We can therefore use the table in the correctness proof. Since we only consider the one function setting on a domain that is a simplicial complex, we emphasize the relevant information in Table 2. We get K_i by adding σ_i to K_{i-1} . If $\sigma_i \in L$ then $L_i = L_{i-1} \cup \{\sigma_i\}$ else $L_i = L_{i-1}$. In Cases B, E, F, the addition of σ_i changes the homology of L_{i-1} , which can only happen if $\sigma_i \in L$. In the remaining three cases, the addition of σ_i changes the homology of K_{i-1} but not that of L_{i-1} , which can only happen if $\sigma_i \in K - L$. Note also that the change in the homology of K_{i-1} is unambiguous in all cases, that is, σ_i is positive in Cases B, C, F, and negative in Cases A, D, E. We note that each death is paired with a unique birth but

Case	ker j_i	im j_i	cok j_i	H(L_i)	H(K_i)	σ_i
A	birth	death	—	—	death	$K - L$
B	—	birth	—	birth	birth	L
C	—	—	birth	—	birth	$K - L$
D	—	—	death	—	death	$K - L$
E	—	death	—	death	death	L
F	death	—	birth	death	birth	L

Table 2: In the one function setting there are six cases in which the addition of σ_i changes the kernel, the image, or the cokernel.

some births remain unpaired until the very end. It is convenient to rephrase the pairing condition in a form that is most directly useful in the argument below. We say a cycle *appears* in $\text{grp } j_l$ if the class it represents is born at that group, where $\text{grp} \in \{\text{ker}, \text{im}, \text{cok}\}$ as usual. We note that the cycle might exist in the complex before it appears in the group.

DEATH LEMMA. Let $l < i$ be indices and z a cycle that appears first in $\text{grp } j_l$ and is zero in $\text{grp } j_i$. If there is no index $i' < i$ for which there is a cycle that first appears in $\text{grp } j_l$ and is zero in $\text{grp } j_{i'}$ then the class represented by z is born at $\text{grp } j_l$ and dies entering $\text{grp } j_i$.

Next we consider the incidence matrices used to compute the persistence diagrams of the sequences of kernels, images, and cokernels. To simplify language, we let $M[i]$ be the column of matrix M that corresponds to σ_i or, alternatively, the set of simplices whose corresponding rows have a one in this column. We will refer to it as *column i* of M and note that in some cases it is not the i -th column from the left, for example when $M = R_{\text{ker}}$.

OBSERVATION. Recall that $R_f, R_g, R_{\text{ker}}, R_{\text{im}}, R_{\text{cok}}$ are the reduced incidence matrices computed by the algorithms in Section 3.

- (i) If $R_g[i] = 0$ then $R_f[i] = 0$.
- (ii) $R_f[i] = 0$ iff $R_{\text{im}}[i] = 0$.
- (iii) If $\sigma_i \in L$ and $R_f[i] \neq 0$ then the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex in L .
- (iv) The columns of R_{ker} are all non-zero.
- (v) If $\sigma_i \in K - L$ then the lowest one in $R_{\text{ker}}[i]$ corresponds to σ_i .
- (vi) If $\sigma_i \in L$, $R_g[i] \neq 0$, and $R_f[i] = 0$ then the lowest one in $R_{\text{ker}}[i]$ corresponds to a simplex in $K - L$.

We omit the proof of this observation from the conference version of the paper.

Inductive step. We assume inductively that the actual and the computed sets of births, deaths, and pairs are the same up to index $i - 1$. This is clearly true for $i - 1 = 0$, when all sets are empty.

Necessity. Using this as the induction hypothesis, we first show that the computed sets of births, deaths, and pairs are subsets of the corresponding actual sets up to index i .

Images. The algorithm for the images reports a birth for the new simplex iff $\sigma_i \in L$ and $R_g[i] = 0$. In this case, $V_g[i]$ stores a cycle that represents a new class in the homology group of L_i as well as in $\text{im } j_i$. Hence, σ_i gives rise to an actual birth in the sequence of images.

The algorithm reports a death iff $R_f[i] \neq 0$ and the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex $\sigma_l \in L$. The reduced column is a sum of boundaries in K_i ,

$$R_{\text{im}}[i] = \sum_{\sigma_k \in V_{\text{im}}[i]} D_{\text{im}}[k].$$

This cycle first appears in $\text{im } j_l$ and it is zero in K_i and therefore also in $\text{im } j_i$. Furthermore, there is no $i' < i$ for which there is a chain that first appears in $\text{im } j_l$ and is zero in $\text{im } j_{i'}$. Otherwise, $(\sigma_l, \sigma_{i'})$ would be a pair and the lowest one in $R_{\text{im}}[i']$ would correspond to σ_l , by inductive hypothesis. But then R_{im} could be reduced further, a contradiction. By the Death Lemma, σ_i gives rise to an actual death and (σ_l, σ_i) is an actual pair in the sequence of images.

Cokernels. The algorithm for the cokernels reports a birth iff $R_f[i] = 0$ and either $\sigma_i \in K - L$ or else $\sigma_i \in L$ and $R_g[i] \neq 0$. In this case, we have indeed a new class in the cokernel, namely the one represented by $V_f[i]$.

The algorithm reports a death iff $R_f[i] \neq 0$ and the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex in $K - L$. By Observation (iii), this implies $\sigma_i \in K - L$. Letting σ_l correspond to the lowest one in $R_{\text{cok}}[i]$, the algorithm reports (σ_l, σ_i) as a pair. In this case, the reduced column is a sum of boundaries in K_i and cycles in L_i ,

$$R_{\text{cok}}[i] = \sum_{\sigma_k \in V_{\text{cok}}[i]} D_{\text{cok}}[k].$$

This cycle appears first in $\text{cok } j_l$ and it is zero in K_i and therefore also in $\text{cok } j_i$. Furthermore, there is no index $i' < i$ for which there is a chain that first appears in $\text{cok } j_l$ and is zero in $\text{cok } j_{i'}$. As before, we use induction and the fact that R_{cok} is reduced to prove this claim. By the Death Lemma, σ_i gives rise to an actual death and (σ_l, σ_i) is an actual pair in the sequence of cokernels.

Kernels. The algorithm for the kernels reports a birth iff $\sigma_i \in K - L$, $R_f[i] \neq 0$, and the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex in L . In this case, $R_{\text{im}}[i]$ is a cycle in L and $V_{\text{im}}[i]$ is a chain whose boundary is this cycle. Furthermore, i is the smallest index for which such a chain exists, else we could use induction to show that R_{im} can be further reduced. Since σ_i belongs to $V_{\text{im}}[i]$, this chain does not belong to L . Hence, $R_{\text{im}}[i]$ represents a class in the kernel and σ_i gives rise to an actual birth.

The algorithm reports a death iff $\sigma_i \in L$, $R_g[i] \neq 0$, and $R_f[i] = 0$. By Observation (vi), the lowest one in $R_{\ker}[i]$ corresponds to a simplex $\sigma_l \in K - L$. The algorithm reports (σ_l, σ_i) as a pair. To prove that there is an actual death, we recall that the columns in D_{\ker} are cycles in K . We write each cycle as a sum of two chains, one in L and the other in its complement, $D_{\ker}[k] = \sum \lambda_\ell + \sum \kappa_\ell$, where the λ_ℓ belong to L and the κ_ℓ belong to $K - L$. The two chains share their boundary, which we denote as $z_k = \partial \sum \lambda_\ell = \partial \sum \kappa_\ell$. Clearly, z_k is a cycle in L . Because it bounds the sum of the κ_ℓ , the cycle belongs to the kernel, and because it bounds the sum of the λ_ℓ , the cycle is zero in L and therefore also zero in the kernel. Consider the cycle

$$z = \sum_{\sigma_k \in V_{\ker}[i]} z_k.$$

We claim that the class it represents in the kernel is born at $\ker j_l$. Indeed, if it were born earlier there would be $l' < l$ and a chain $c \in K_{l'}$ whose boundary is z . But then σ_l would be positive and by Observation (v) it would be the lowest one of its own column and not that of column i . Now, $V_{\ker}[i]$ provides the sum we need to finish the proof using the Death Lemma. As before, we use the induction hypothesis and the fact that R_{\ker} is reduced to conclude that there is no index $i' < i$ for which a cycle appears in $\ker j_l$ and is zero in $\ker j_{i'}$. Therefore, σ_i gives rise to an actual death and (σ_l, σ_i) is an actual pair in the sequence of kernels.

Sufficiency. We second show that the algorithms are complete, that is, the actual births, deaths, and pairs are subsets of the corresponding computed sets. Since there is a bijection between the deaths and the pairs, it suffices to prove the containments for the births and the deaths. We use Table 2 to do this by exhaustive case analysis.

Case 1 $\sigma_i \in K - L$.

Case 1.1 $R_f[i] = 0$. This is Case C in Table 2. The only change is a birth in the cokernel and this is correctly reported by the algorithms.

Case 1.2 $R_f[i] \neq 0$. Let σ_l correspond to the lowest one in $R_{\text{im}}[i]$.

Case 1.2.1 $\sigma_l \in K - L$. From the above analysis we know that this corresponds to a death in the cokernel. This is Case D in Table 2. There are no other changes and this is correctly reported by the algorithms.

Case 1.2.2 $\sigma_l \in L$. From the above analysis we know that this corresponds to a birth in the kernel and a death in the image. This is Case A in Table 2. There are no other changes and this is correctly reported by the algorithms.

Case 2 $\sigma_i \in L$.

Case 2.1 $R_f[i] = 0$.

Case 2.1.1 $R_g[i] \neq 0$. This is Case F in Table 2. There is a death in the kernel, a birth in the cokernel, and no

change in the image, and this is correctly reported by the algorithms.

Case 2.1.2 $R_g[i] = 0$. This is Case B in Table 2. The only change is a birth in the image and this is correctly reported by the algorithms.

Case 2.2 $R_f[i] \neq 0$. By Observation (i), this implies $R_g[i] \neq 0$. This is Case E in Table 2. The only change is a death in the image which is correctly reported by the algorithms.

We conclude that the actual births and deaths are subsets of the computed births and deaths, and similarly that the actual pairs are a subset of the computed pairs.

We now have the containment relations in both directions which implies that corresponding sets of computed and actual births, deaths, and pairs are in fact the same. This concludes the proof that the algorithms in Section 3 correctly compute the persistence diagrams of the sequences of kernels, images, and cokernels.

5 Applications

In this section, we use kernels to measure local homology and images to approximate persistence diagrams of noisy functions specified on noisy domains. There are additional applications that are sufficiently straightforward that we can leave the details to the interested reader. One is the denoising of alpha-beta witness complexes as introduced in [1]; see also [9]. By considering maps from one complex to another, more tolerantly constructed complex we can preserve persistent features without accidentally introducing new ones. By varying the scale parameter, α , we thus get a persistence diagram that is less noisy than the diagram of the sequence of complexes for fixed tolerance parameter β . Another application is the computation of rank invariants for a doubly-filtered space $\mathbb{X}_{i,k}$ as considered in [3]. Here we get a simple algorithm by encoding the ranks of the images of $H_p(\mathbb{X}_{i,k}) \rightarrow H_p(\mathbb{X}_{i',k'})$, for fixed $i < i'$ and variable $k < k'$, in a single persistence diagram of the images of the maps $H_p(\mathbb{X}_{i,k}) \rightarrow H_p(\mathbb{X}_{i',k})$. We can compute all such diagrams in time quartic in the number of simplices in the triangulation of \mathbb{X} .

5.1 Local Homology

We begin with the application of kernels to measuring the local homology of a space in \mathbb{R}^n at a point not necessarily in the space. Following earlier work, we assume that the space is not known other than indirectly through a finite set of points sampled in or near the space.

Measuring local homology. Bendich et al. study the reconstruction of a stratified space \mathbb{S} from a finite point sample U in \mathbb{R}^n [2]. Specifically, they use persistence to define a multi-scale version of the local homology of \mathbb{S} at a point $z \in \mathbb{R}^n$. Let \mathbb{S}_α be the sets of points at Euclidean

distance at most α from \mathbb{S} , \mathbb{S}^α the set of points at distance at least α from \mathbb{S} , and B_r the closed ball of radius r centered at z . They express the homology within a fixed distance r of z in terms of the persistence diagram of the sequence

$$(5.1) \quad \begin{aligned} 0 &\rightarrow H(\mathbb{S}_\alpha \cap B_r) \rightarrow \dots \rightarrow H(B_r) \\ &\rightarrow H(B_r, \mathbb{S}^\alpha \cap B_r) \rightarrow \dots \rightarrow 0, \end{aligned}$$

where α first goes up, from 0 to ∞ , and then down, from ∞ to 0. The first half of the sequence captures the development of the absolute homology of \mathbb{S}_α within B_r , and it can be shown that the second half captures the development of the relative homology of \mathbb{S}_α within the pair $(B_r, \partial B_r)$. We note that cycles that lie entirely inside the ball are captured twice, once in each half. Finally, Bendich et al. vary r from 0 to ∞ and this way obtain a vineyard that expresses the local homology of \mathbb{S} at z under the 2-parameter variation of α and r . They also prove relationships between this vineyard and the similarly defined vineyard of a finite set of points $U \subseteq \mathbb{R}^n$ sampled near \mathbb{S} .

In this paper, we propose to substitute a sequence of kernels for (5.1). Specifically, let $\mathbb{X} = B_r$, $\mathbb{Y} = \partial B_r$, and let $f : \mathbb{X} \rightarrow \mathbb{R}$, $g : \mathbb{Y} \rightarrow \mathbb{R}$ map each point to its Euclidean distance from \mathbb{S} . For each value α we write $\mathbb{X}_\alpha = f^{-1}(-\infty, \alpha]$, $\mathbb{Y}_\alpha = g^{-1}(-\infty, \alpha]$ and let $j_\alpha : H(\mathbb{Y}_\alpha) \rightarrow H(\mathbb{X}_\alpha)$ be the map induced by the inclusion $\mathbb{Y}_\alpha \subseteq \mathbb{X}_\alpha$. Assuming f and g are both tame we have a finite set of critical values and thus a finite sequence of kernels,

$$(5.2) \quad \text{Ker}(g \rightarrow f) : \ker j_0 \rightarrow \ker j_1 \rightarrow \dots \rightarrow \ker j_m,$$

which traces the evolution of the relative homology classes in (5.1) that have a non-zero boundary in $\mathbb{Y} = \partial B_r$. The Stability Theorem in Section 2 implies that varying r from 0 to ∞ gives a vineyard. It tracks a homology class as long as the boundary of the ball with radius r intersects all its representatives. The interval of radii expresses relevant size information, namely how far away from z the class starts and ends. It is thus no longer necessary to include absolute homology classes in the measurement and the vineyard simplifies into a form that more closely reflects the shape of the space in the neighborhood of the point z ; see Figure 2. Relationships between the vineyard of \mathbb{S} and that of a finite point sample $U \subseteq \mathbb{R}^n$ similar to the Local Homology Inference and the Inverse LHI Theorems in [2] can be proved using the same methods. Details are omitted.

Computing local homology. Following the prior work, we use the Delaunay triangulations of U restricted to the ball B_r and to the ball without the interior of the power cell of the point z , $Z_0(r) = B_r - \text{int } Z(r)$; see [2, Section 6] for details. We use distance functions and Delaunay triangulations for clear and solid theoretical footing. For practical high-dimensional implementations, like [2], these results need to be extended to Vietoris-Rips [13] and Witness complexes [9]. Let $K_\alpha = \text{Del}(U|_{U_\alpha \cap B_r})$ and $L_\alpha = \text{Del}(U|_{U_\alpha \cap Z_0(r)})$ be the Delaunay triangulations that are




Figure 2: The vineyard of the chain of loops defined using the sequence (5.1) on the left and the sequence (5.2) of kernels on the right. To compare the two diagrams we see that the kernel sequence reflects the left half of the left vineyard across the diagonal plane and corrects for the removed absolute homology groups.

further restricted to the set of points at distance at most α from U . Let furthermore

$$\begin{aligned} i_\alpha &: H(U_\alpha \cap Z_0(r)) \rightarrow H(U_\alpha \cap B_r); \\ l_\alpha &: H(L_\alpha) \rightarrow H(K_\alpha) \end{aligned}$$

be the maps between the homology groups induced by inclusion. To justify the use of the restricted Delaunay triangulations, we need to show that the persistence diagrams of the kernels of these maps are the same. This requires that the following diagram on the left is well-defined, commutative, and its vertical maps are isomorphisms whenever $\alpha \leq \alpha'$:

$$\begin{array}{ccc|cc} \ker j_\alpha & \rightarrow & \ker j_{\alpha'} & \begin{array}{c} H(U_\alpha \cap \partial B_r) \\ \downarrow \\ H(U_\alpha \cap Z_0(r)) \\ \uparrow \\ H(L_\alpha) \end{array} & \begin{array}{c} \xrightarrow{j_\alpha} \\ \xrightarrow{i_\alpha} \\ \xrightarrow{l_\alpha} \end{array} & \begin{array}{c} H(U_\alpha \cap B_r) \\ \downarrow \\ H(U_\alpha \cap B_r) \\ \uparrow \\ H(K_\alpha) \end{array} \\ \downarrow i_\alpha & & \downarrow i_{\alpha'} & & & \\ \ker i_\alpha & \rightarrow & \ker i_{\alpha'} & & & \\ \uparrow h_\alpha & & \uparrow h_{\alpha'} & & & \\ \ker l_\alpha & \rightarrow & \ker l_{\alpha'} & & & \end{array}$$

Consider the diagram on the right whose maps are all induced by inclusions except for the lower vertical maps which are motivated by the Nerve Subdivision Lemma [2, Section 8]. The top square commutes. The bottom square also commutes since the horizontal maps are induced by inclusion and the vertical map for L_α is the restriction of the one for K_α . Therefore the kernel diagram is well-defined. Since the kernels are subgroups of the domains of their defining maps, the kernel diagram is a restriction of a diagram considered in [2, Section 8]. The analysis there implies that it commutes and its vertical maps are isomorphisms, as required.

The Stability Theorem of Section 2 implies that the kernel persistence diagrams change continuously with the radius of the restricted ball B_r . We construct the implied

vineyard by maintaining the ordering of the simplices and the reduced matrices. However, unlike with relative homology in [2] we cannot use excision to maintain different orderings of a static complex. We need to be able to handle insertion of simplices into the ordering when the power cell of the point z expands with r to include new simplices. Fortunately, the new simplices are paired amongst each other so that these updates can be done in linear time per insertion. Once a simplex is inserted, its position in the ordering can be described by a continuous function; see Appendices A and B of [2]. It therefore suffices to maintain the decompositions in the four steps of the algorithm under transpositions of contiguous simplices. Details on how to perform these operations are given in [11, Section 6.5].

5.2 Noisy Domains

Persistent homology has proven to be well-suited for dealing with noisy functions. Indeed, the stability of persistence diagrams implies that the topological features of an unknown ideal function $\tilde{f} : \mathbb{X} \rightarrow \mathbb{R}$ can be approximately recovered knowing only a noisy approximation f of \tilde{f} . We claim that the persistence for images can be used to furthermore filter out the topological noise induced by the domain itself.

Stability. We assume an unknown ideal domain given as the zero sublevel set of the unknown ideal function $\tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}$, that is, $\tilde{\mathbb{X}} = \tilde{h}^{-1}(-\infty, 0]$. On this domain we consider another unknown ideal function but because we will vary the domain we assume it is defined on the entire ambient space, $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}$. Can we estimate the persistence diagram of the restriction $\tilde{f}|_{\tilde{\mathbb{X}}} : \tilde{\mathbb{X}} \rightarrow \mathbb{R}$ knowing only noisy approximations h, f of \tilde{h}, \tilde{f} ? We give an affirmative answer under mild requirements on the functions. To describe these requirements we use superscripts for sublevel sets of h and \tilde{h} and subscripts for sublevel sets of f and \tilde{f} , that is, $\mathbb{X}^u = h^{-1}(-\infty, u]$, $\mathbb{X}_a = f^{-1}(-\infty, a]$, $\mathbb{X}_a^u = \mathbb{X}_a \cap \mathbb{X}^u$ and similarly for $\tilde{\mathbb{X}}^u, \tilde{\mathbb{X}}_a, \tilde{\mathbb{X}}_a^u$. Writing $\varepsilon = \|h - \tilde{h}\|_\infty$ we require that \tilde{h} is smooth and the norm of its gradient is bounded away from 0 where this is relevant, that is, $\|\nabla \tilde{h}\| \geq \mu > 0$ on $\tilde{\mathbb{X}}^{2\varepsilon} - \tilde{\mathbb{X}}^{-2\varepsilon}$. Furthermore, we write $\delta = \|f - \tilde{f}\|_\infty$ and require that f is Lipschitz with constant κ , that is, $|f(x) - f(y)| \leq \kappa\|x - y\|$ for all $x, y \in \mathbb{R}^n$.

We note that the requirement on \tilde{h} implies a homotopy between the identity on $\tilde{\mathbb{X}}^{2\varepsilon}$ and a retraction ϱ from $\tilde{\mathbb{X}}^{2\varepsilon}$ to $\tilde{\mathbb{X}}^0$. To construct the homotopy we consider the integral lines of the vector field $-\nabla \tilde{h}$ starting at points $x \in \tilde{\mathbb{X}}^{2\varepsilon}$. Let $\varrho(x)$ be the first point on the curve starting at x that satisfies $\tilde{h}(\varrho(x)) = \tilde{h}(x) - 2\varepsilon$ or $\tilde{h}(\varrho(x)) \leq -2\varepsilon$. In words, ϱ moves the fringe outside the boundary of $\tilde{\mathbb{X}}$ to the fringe inside that boundary and it moves the fringe inside that boundary to the boundary of $\tilde{\mathbb{X}}^{-2\varepsilon}$. Since the gradient has norm no smaller than $\mu > 0$ along the integral line, the retraction $\varrho : \tilde{\mathbb{X}}^{2\varepsilon} \rightarrow \tilde{\mathbb{X}}^{2\varepsilon}$ is well defined. By construction, there is a homotopy between the identity on $\tilde{\mathbb{X}}^{2\varepsilon}$ and ϱ that moves points by at most $2\varepsilon/\mu$. We are now ready to state our result.

STABILITY THEOREM FOR NOISY DOMAINS. Let $h, \tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\varepsilon = \|h - \tilde{h}\|_\infty$, \tilde{h} smooth, and the norm of the gradient satisfying $\|\nabla \tilde{h}\| \geq \mu > 0$ on $\tilde{\mathbb{X}}^{2\varepsilon} - \tilde{\mathbb{X}}^{-2\varepsilon}$. Furthermore, let $f, \tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\delta = \|f - \tilde{f}\|_\infty$ and f Lipschitz with constant κ . Then

$$d_B(\text{Dgm}(\tilde{f}|_{\tilde{\mathbb{X}}}), \text{Dgm}(\text{im } f|_{\mathbb{X}^{-\varepsilon}} \rightarrow f|_{\mathbb{X}^\varepsilon})) \leq 2\kappa\varepsilon/\mu + \delta$$

provided the restrictions of \tilde{f} to $\tilde{\mathbb{X}}$ and of f to $\tilde{\mathbb{X}}, \mathbb{X}^\varepsilon, \mathbb{X}^{-\varepsilon}$ are continuous and tame and there exists a triangulation of \mathbb{X}^ε in which $\mathbb{X}^{-\varepsilon}$ and $\tilde{\mathbb{X}}$ arise as subcomplexes.

Proof. By the Stability Theorem for ordinary persistence we have $d_B(\text{Dgm}(\tilde{f}|_{\tilde{\mathbb{X}}}), \text{Dgm}(f|_{\tilde{\mathbb{X}}})) \leq \delta$. It remains to show that the bottleneck distance between $\text{Dgm}(f|_{\tilde{\mathbb{X}}})$ and $\text{Dgm}(\text{im } f|_{\mathbb{X}^{-\varepsilon}} \rightarrow f|_{\mathbb{X}^\varepsilon})$ is bounded from above by $c = 2\kappa\varepsilon/\mu$.

Writing F_a for $H(\tilde{\mathbb{X}}_a^0)$, the diagrams $\text{Dgm}(f|_{\tilde{\mathbb{X}}})$ is obtained from the sequence formed by the maps $F_a \rightarrow F_b$ induced by the inclusion $\tilde{\mathbb{X}}_a^0 \subseteq \tilde{\mathbb{X}}_b^0$ for all $a \leq b$. Similarly, writing J_a for the image of the map $H(\mathbb{X}_{a-\varepsilon}^-) \rightarrow H(\mathbb{X}_a^\varepsilon)$, the diagram $\text{Dgm}(\text{im } f|_{\mathbb{X}^{-\varepsilon}} \rightarrow f|_{\mathbb{X}^\varepsilon})$ is obtained from the sequence of maps $J_a \rightarrow J_b$ again induced by inclusion and for all $a \leq b$. To adapt the proof of stability given in [6], we need to connect these two sequences by maps $\phi_a : F_{a-c} \rightarrow J_a$ and $\psi_a : J_{a-c} \rightarrow F_a$, for all $a \in \mathbb{R}$, in such a way that the diagram formed by the two sequences together with the new maps commutes. We construct the new maps using the homotopy between the identity and the retraction $\varrho : \tilde{\mathbb{X}}^{2\varepsilon} \rightarrow \tilde{\mathbb{X}}^{2\varepsilon}$. As mentioned earlier, the homotopy moves a point by at most $2\varepsilon/\mu$ and since f is Lipschitz with constant κ , ϱ maps $\tilde{\mathbb{X}}_{a-c}^0$ to $\tilde{\mathbb{X}}_a^{-2\varepsilon}$ which is included in $\mathbb{X}_{a-\varepsilon}^-$. The induced homomorphism from $H(\tilde{\mathbb{X}}_{a-c}^0)$ to $H(\mathbb{X}_{a-\varepsilon}^-)$ composed with the induced homomorphism from $H(\mathbb{X}_{a-\varepsilon}^-)$ to $H(\mathbb{X}_a^\varepsilon)$ gives the map ϕ_a . ϕ_a connects the two sequences with a shift of $c = 2\kappa\varepsilon/\mu$. By a similar process, we construct the map ψ_a connecting the two sequences of vector spaces in the other direction and again with a shift of c . Because ϱ is homotopic to the identity, the diagram formed by the two sequences and the maps ϕ_a and ψ_a commutes. The remainder of the proof can be adapted directly from [6].

Diagram approximation. An interesting case of the above theorem arises when we consider a finite set of points, U , sampling an unknown shape, $\mathbb{S} \subseteq \mathbb{R}^n$. Let $\tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}$ be the distance function of \mathbb{S} , that is, $\tilde{h}(x) = \inf_{y \in \mathbb{S}} \|x - y\|$. Similarly, let $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be the distance function of U and set ε to the Hausdorff distance between \mathbb{S} and U . For technical reasons we may have to replace \tilde{h} by a smooth approximation, for example obtained by convolution with an infinitesimally narrow Gaussian so that the assumptions in the theorem are satisfied. In this setting, the requirement that the norm of the gradient of \tilde{h} is bounded from below by μ is equivalent to the μ -reach of \mathbb{S} exceeding 4ε . Here we recall that the μ -reach as recently introduced in [4] is a notion of feature size that permits the treatment of non-smooth

objects. Under this assumption, the Stability Theorem for Noisy Domains implies that it is possible to estimate the persistence diagram of a function \tilde{f} restricted to the 2ε -offset of \mathbb{S} knowing only a Lipschitz function f that approximates \tilde{f} and the point set U that samples \mathbb{S} .

To approximate the persistence diagram of \tilde{f} restricted to \mathbb{S} itself, we exploit the existence of an isotopy ι from $\mathbb{S}^{2\varepsilon}$ to an arbitrarily small offset \mathbb{S}^η of the shape such that the points move by less than $2\varepsilon/\mu$ during the deformation. The construction of ι is similar to the construction of the homotopy between the identity and ϱ described above. The isotopy implies that $\text{Dgm}(f|_{\mathbb{S}^{2\varepsilon}})$ equals $\text{Dgm}(f \circ \iota^{-1}|_{\mathbb{S}^\eta})$ which in turn is c -close to $\text{Dgm}(f|_{\mathbb{S}^\eta})$ by stability. Hence, the latter diagram can also be estimated from f and U with an accuracy of $2c$. Perhaps surprisingly, $\text{Dgm}(f|_{\mathbb{S}^\eta})$ may not converge to $\text{Dgm}(f|_{\mathbb{S}})$ as η goes to 0; see [5] for examples of shapes \mathbb{S} that lack this convergence property. However, convergence holds for sufficiently regular spaces S , such as smooth submanifolds or geometrically realized simplicial complexes. In these cases we can estimate $\text{Dgm}(f|_{\mathbb{S}})$ with precision $2c$.

In the more realistic case in which f is only known at a finite set of points, U , a valid approach replaces f by the function \tilde{f} that is constant on the Voronoi cells of the points and coincides with f on U . While \tilde{f} is not continuous, it is almost everywhere continuous in a way that does not disrupt the proof of stability. Furthermore, f and \tilde{f} differ by at most $4\kappa\varepsilon$ on $\tilde{h}^{-1}(-\infty, 4\varepsilon]$. By the Stability Theorem for Noisy Domains, the persistence diagrams of the images defined by f and \tilde{f} are close. The diagram for \tilde{f} can be computed using the alpha shape filtration of U . We thus get a practical algorithm for estimating the persistence diagram of functions given only at a finite set of points.

6 Discussion

In this paper, we consider persistent homology for sequences of kernels, images, and cokernels defined by a pair of topological spaces, $\mathbb{Y} \subseteq \mathbb{X}$, and two functions, $f : \mathbb{X} \rightarrow \mathbb{R}$ and $g : \mathbb{Y} \rightarrow \mathbb{R}$, with $f(y) \leq g(y)$ for every $y \in \mathbb{Y}$. Since g majorizes the restriction of f to \mathbb{Y} , its sublevel sets are contained in those of f , $\mathbb{Y}_a = g^{-1}(-\infty, a] \subseteq \mathbb{X}_a = f^{-1}(-\infty, a]$, and we have homomorphisms $j_a : H(\mathbb{Y}_a) \rightarrow H(\mathbb{X}_a)$ induced by the inclusions. To see that persistent homology is well defined we just need to note that the diagrams involving groups $H(\mathbb{Y}_a)$, $H(\mathbb{Y}_b)$, $H(\mathbb{X}_a)$, $H(\mathbb{X}_b)$ induced by inclusion between respective spaces commute for any $a \leq b$ and thus induce homomorphic maps $\ker j_a \rightarrow \ker j_b$, $\text{im } j_a \rightarrow \text{im } j_b$, and $\text{cok } j_a \rightarrow \text{cok } j_b$.

It is worth noting that the mapping cylinder construction described in Section 2 can be used to extend the framework from inclusion $\mathbb{Y} \subseteq \mathbb{X}$ to an arbitrary continuous map $j : \mathbb{Y} \rightarrow \mathbb{X}$. If we have two functions $f : \mathbb{X} \rightarrow \mathbb{R}$, and $g : \mathbb{Y} \rightarrow \mathbb{R}$, such that $j(\mathbb{Y}_a) \subseteq \mathbb{X}_a$ with the sublevel set of each space taken with respect to its own map,

then the maps $j_a : \mathbb{Y}_a \rightarrow \mathbb{X}_a$ between the sublevel sets induce homomorphisms on homology groups just the same. Constructing a mapping cylinder $\mathbb{X}' = \mathbb{X} \cup \mathbb{Y} \times [0, 1]$ by identifying $(y, 0) \in \mathbb{Y} \times \{0\}$ with $j(y) \in \mathbb{X}$, it is easy to verify that the sequences of kernels, images, and cokernels induced by inclusion $\mathbb{Y}_a = \mathbb{Y}_a \times \{1\} \subseteq \mathbb{X}'_a$ give the same persistence pairing as the three sequences induced by the continuous maps j_a defined above.

One can also extend the framework described in this paper from spaces $\mathbb{Y}_a \subseteq \mathbb{X}_a$ to pairs of spaces $(\mathbb{Y}_a, \mathbb{Y}_{a_0}) \subseteq (\mathbb{X}_a, \mathbb{X}_{a_0})$ using the cone construction exploited for the computation of extended persistence [7]. Indeed, observing that the relative homology groups $H(\mathbb{Y}_a, \mathbb{Y}_{a_0})$ are isomorphic to the homology groups of \mathbb{Y}_a with a cone on \mathbb{Y}_{a_0} rel the cone point, i.e. $H(\mathbb{Y}_a, \mathbb{Y}_{a_0}) \simeq H(\mathbb{Y}_a \cup CY_{a_0}, \omega)$, we can compute the persistence of kernels, images, and cokernels induced by the inclusions of pairs of spaces by using the algorithms of this paper on corresponding cones.

References

- [1] D. ATTALI, H. EDELSBRUNNER, J. HARER AND Y. MILEYKO. Alpha-beta witness complexes. In "Proc. 11th Workshop Alg. Data Struct., 2007", Springer-Verlag, Lecture Notes in Computer Science **4619**, 386–397.
- [2] P. BENDICH, D. COHEN-STEINER, H. EDELSBRUNNER, J. HARER AND D. MOROZOV. Inferring local homology from sampled stratified spaces. In "Proc. 48th Ann. Sympos. Found. Comput. Sci., 2007".
- [3] G. CARLSSON AND A. ZOMORODIAN. The theory of multidimensional persistence. In "Proc. 22nd Ann. Sympos. Comput. Geom., 2007", 184–193.
- [4] F. CHAZAL, D. COHEN-STEINER AND A. LIEUTIER. A sampling theory for compact sets in Euclidean space. In "Proc. 22nd Ann. Sympos. Comput. Geom., 2006", 319–326.
- [5] F. CHAZAL AND A. LIEUTIER. Stability and computation of topological invariants of solids in \mathbb{R}^n . *Discrete Comput. Geom.* **37** (2007), 601–617.
- [6] D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), 103–120.
- [7] D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.*, to appear.
- [8] 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.
- [9] V. DE SILVA AND G. CARLSSON. Topological estimation using witness complexes. In "Proc. Sympos. Point-Based Graphics, 2004", 157–166.
- [10] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511–533.
- [11] D. MOROZOV. *Homological Illusions of Persistence and Stability*. Ph.D. Thesis, Duke University, 2008.
- [12] J. R. MUNKRES. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, California, 1984.
- [13] L. VIETORIS. Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen. *Math. Ann.* **97** (1927), 454–472.
- [14] A. ZOMORODIAN AND G. CARLSSON. Computing persistent homology. *Discrete Comput. Geom.* **33** (2005), 249–274.