Topology and local optima in computer vision

Erik Carlsson¹ and John Carlsson²

¹Department of Mathematics, UC Davis ²Department of Industrial engineering, University of Southern California

June 18, 2021

Abstract

We present an application of persistent homology to the image correspondence problem, also known as image registration, which is used to produce 3D reconstruction of scenery from two or more cameras. We present a novel filtered complex in the sense of persistent homology, and show that nontrivial homology groups in its persistence diagrams correspond to recognizable anomalies in images pairs, such as repeated patterns, which contribute to nonconvexity of the relevant cost function. We present examples with actual image pairs, and prove a basic result that the corresponding homology classes are invariant under certain continuous deformations.

1 Introduction

Suppose we are given a pair of pictures of the same scene from different angles, represented by grayscale images called the source and target respectively. We will denote them as functions $I_j : \mathcal{D}_j \to \mathbb{R}$ for j = 1, 2 for the source and target respectively, where both domains are the unit square $\mathcal{D}_1 = \mathcal{D}_2 = [0, 1] \times [0, 1]$, and the value of $I_j(p) = I_j(x, y)$ is the intensity. The goal of image correspondence is to find a suitable transformation $T : \mathcal{D}_1 \to \mathcal{D}_2$ with the property that E(T) is as small as possible, where E is some cost function measuring the difference between $I_1(x, y)$ and the the composition $I_2(T(x, y))$. In typical applications, the correspondence T might be considered suitable if it is continuous, differentiable, of bounded variation, or some other niceness criteria. If the source and target come from cameras at the same height (e.g. a person's left and right eyes), it is natural to require that T satisfy the "epipolar" condition that corresponding pairs $p, \varphi(p)$, have the same y-coordinate. A basic example of a cost function is the L^2 -measure

$$E(\varphi) = \int_{\mathcal{D}_1} \|I_1(p) - I_2(\varphi(p))\|^2 \, dp,$$
(1)

which might also contain a term penalizing high variation in φ . Another possibility if I_1, I_2 are continuous is to replace the integral in (1) with the supremum over \mathcal{D}_1 .

A common problem with the optimization approach is that $E(\varphi)$ may have local minimizers. Some are simply due to noise in the images, which can be resolved by a smoothing of the objective function, for instance by blurring the images. More serious ones occur when the source image matches with more than one region in the target, or when there is occlusion by foreground objects, so that regions in the source do not have a unique match in the target. One way to deal with this is to only study key points of the two images, as is the case for the highly successful SIFT and SURF methods [1, 6].

In this paper, rather than smooth away local basins, we propose a novel construction which uses persistent homology to classify them, and identify regions in which they are present. We present a new filtered simplicial complex associated to a pair of images, with the property that its persistent homology groups encode robust categories of correspondences, which respect continuous deformation in a particular sense defined in Section 5. In Section 4, we present practical examples in which confounding properties of image pairs such as repeated patterns can be identified with long bars in the persistent diagram of the corresponding complex.

There are several ways such classifiers could be used to address the original correspondence problem. First, if persistent homology detects the presence of one or more correspondences in a particular domain, one could divise algorithms for producing one or more correspondences that represent a given homology class, which is part of the general problem of producing cycle representatives of homology classes when they exist. Such representatives may then serve as highly informed initial guesses that avoid obstructions, due to confounding features of the image detected by topology. Conversely, if there are no persistent homology classes, that represents a certificate that there are no viable correspondences in a particular domain, as described in Corollary 1 below. The solver would therefore be well served to move on to another region of the source images, rely on other cameras, or move on to a different formulation of the correspondence problem. This leads to the possibility of searching for subregions of the source image of maximal size which satisfy the property that topologically, they are expected to contain a unique correspondence.

1.1 Acknowledgments

Erik Carlsson was supported by NSF DMS-1802371 during part of this project. John Carlsson was supported by DARPA grant N660011824029, and NSF grant CMMI-1634193.

2 Definitions

In this section we recall some relevant concepts from persistent homology, and formulate the correspondence problem.

2.1 Persistent homology

In this paper, a simplicial complex X will mean an abstract simplicial complex with no particular imbeding in space. In other words, X is a collection of subsets Δ of some index set S such that if $\Delta \in X$, and $\Delta' \subset \Delta$, then $\Delta' \in X$. The subsets $\Delta' \subset \Delta$ are called the faces of Δ . Its geometric realization will be denoted |X|. For any *i*, let X_i denote the *i*-dimensional simplices. Fix a field F, and denote the set of *i*-chains, *i*-cycles, *i*-boundaries, and the boundary operator over F by $C_i(X), Z_i(X), B_i(X)$, and ∂ , respectively. In applications, it is most efficient to let F be a finite field. If $A \subset X$ is a subcomplex, then we have the long exact sequence in (reduced) relative homology,

$$\dots \to H_i(A) \to H_i(X) \to H_i(X, A) \xrightarrow{\delta} H_{i-1}(A) \to \dots$$
 (2)

where $H_i(X, A)$ is the relative homology group, and δ is the connecting homomorphism.

If X is induced from a triangulation of a closed region with boundary in \mathbb{R}^n , then its boundary is the image of a subcomplex $\partial X \subset X$. The orientation on \mathbb{R}^n determines a well-defined fundamental class

$$[X] = \sum_{\Delta \in X_n} \pm e_\Delta \in H_n(X, \partial X),$$

where there is a minus sign if the ordering of the indices on Δ is the reverse of the orientation on \mathbb{R}^n .

A filtration function on X is a function $f : X \to \mathbb{R} \cup \{\infty\}$ such that whenever Δ' is a face of $\Delta \in X$, we have $f(\Delta') \leq f(\Delta)$. For any a, the set

$$X^{a} = \left\{ \Delta \in X : f(\Delta) \le a \right\},\$$

is a subcomplex, with an inclusion map $\iota^{a,b} : X^a \to X^b$ for $a \leq b$. A complex together with a filtration function is called a filtered complex. We will denote $C_i^a(X) = C_i(X^a)$, and similarly for the cycles, boundaries, and homology groups. Then $\iota^{a,b}$ induces an inclusion map $i_*^{a,b} : C_i^a(K) \hookrightarrow C_i^b(X)$ that commutes with the boundary operator, which in turn induces a map $H_i^a(X) \to H_i^b(X)$, that need not be injective or surjective. For $a \leq b$, let $H_i^{a,b}(X)$ denote the persistent homology group, which is the image of $H_i^a(X) \in H_i^b(X)$.

For any numbers $a \leq b$, we have a nonnegative integer

$$\rho_K(a,b) = \dim H_i^{a,b}(K) = \operatorname{rk}\left(H_i^a(K) \to H_i^b(K)\right).$$

The ranks are encoded in the barcode diagram [2, 3, 9], which is the unordered collection of intervals in $\mathbb{R}_+ \cup \{\infty\}$, with the property that

$$\rho_K(a,b) = \# \{ [c,d] \in barcode(K) : [a,b] \subset [c,d] \}.$$
(3)

It is constructed by assuming that a, b take a discrete set values in $\mathbb{Z} \cdot \epsilon \subset \mathbb{Q}$. Then considering all the homology groups at once as a graded module

$$\mathbb{F}[x] \to M = \bigoplus_{a} H^{a}(X), \quad x \mapsto \iota = \sum_{a} \iota_{*}^{a,a+\epsilon}$$

The barcode is then determined by decomposing M as a module over a principal ideal domain. For an explanation of how these barcodes that can be generated, we refer to the JavaPlex tutorial [8].

2.2 Image correspondences

Denote a source and target image I_1, I_2 respectively by functions from the unit square

$$I_j: \mathcal{D}_j \to C, \quad \mathcal{D}_1 = \mathcal{D}_2 = [0, 1] \times [0, 1],$$

into some set of possible colors C, for instance $[0,1]^3$ for color images or [0,1] for grayscale. We suppose that both images are pictures of the same scene, taken from different predetermined locations and angles. For each $p \in \mathcal{D}_1$ belonging to the domain of the source image, we choose a parametrization of the epipolar line of the form

$$q_p: [0,1] \to \mathbb{R}^2, \quad q_{(x,y)}(t) = (x + ct + d, y)$$

for some numbers c, d determined by the relative placement of two cameras, and lower and upper bounds on the x coordinate of corresponding points. In other words, the line parametrizes the points in the target which could be in correspondence with points in the source, assuming the cameras are parallel to the ground, of the same height, and pointed in the same direction. Our examples are all of this form, which is called being *rectified*, but in general the parametrization could be more complicated.

Now select a triangulation of \mathcal{D}_1 represented by an inclusion of a pure 2-dimensional simplicial complex $p: |K| \to \mathcal{D}_1$. Suppose we are also given a continuous distance function $d(\mathfrak{T}, \mathfrak{T}')$ for every pair of triangles $\mathfrak{T}, \mathfrak{T}' \in \mathcal{D}_1$, measuring the distance between the restriction of I_1 to \mathfrak{T} , and the restriction of I_2 to \mathfrak{T}' . The basic function we will use in the examples section is to affinely map \mathfrak{T} and \mathfrak{T}' to the unit right triangle whose vertices are (0,0), (1,0), and (0,1), and evaluate the L^2 metric between them. That is,

$$d([p_0, p_0 + u_1, p_0 + u_2], [q_0, q_0 + v_1, q_0 + v_2]) = \int_{s=0}^{1} \int_{t=0}^{1-s} d_C \left(I_1(p_0 + su_1 + tu_2), I_2(q_0 + sv_1 + tv_2) \right) ds dt$$
(4)

where d_C is just the L^2 -metric on the color space $C = [0, 1]^3$. Another extension we will use involves a penalty term when \mathcal{T}' is highly warped, meaning it is very far from being equilateral, or has the opposite orientation as \mathcal{T} .

Define a simplicial correspondence to be an element of the set

$$\Gamma(K) = \{\varphi : K_0 \to I\}$$

of functions on the vertex set, representing a piecewise linear function on $|K| = \mathcal{D}_1$. The cost function $E : \Gamma(K) \to \mathbb{R}_+$ is given by the worst-case triangle

$$E(\varphi) = \max_{[i_0, i_1, i_2] \in K} d\left([p_0, p_1, p_2], [q_{p_0}(\varphi(i_0)), q_{p_1}(\varphi(i_1)), q_{p_2}(\varphi(i_2))] \right)$$
(5)

for $p_a = p(i_a)$.

3 Main construction

In this section we present the main object of the paper, define the complex described in the introduction in the case of two-dimensional images, and present the general pipeline to be followed in the examples of Section 4.

3.1 Illustration in the continuous case

We describe the idea in the case of one dimensional images. Suppose I_1, I_2 are functions on the interval $\mathcal{D}_1 = \mathcal{D}_2 = [0, 1]$. Consider a correspondence to be a continuous increasing function $T : \mathcal{D}_1 \to \mathcal{D}_2$, with the property that

$$d(x, T(x)) = ||I_1(x) - I_2(T(x))||^2 \le a$$

for some upper bound a > 0 on the dissimilarity between I_1 and I_2 . For the functions in Figure 1, and a > 0, there will be infinitely many correspondences. However, one can see that there are essentially three groups of them up to continuous deformation, corresponding to the upward and rightward moving paths from the lower boundary of Figure 1c to the upper one, avoiding the shaded regions which have high dissimilarity (the fact that the three paths only move up and to the right is equivalent to the requirement that T has to be increasing).

Let $\mathfrak{X} = \mathfrak{D}_1 \times \mathfrak{D}_2$ be the set of (x, y) pairs in Figure 1c. We have the persistent homology group for the sublevel set filtration

$$H_i^{a,b}(\mathfrak{X}) = \operatorname{im}\left(H_i\left(\mathfrak{X}^a\right) \to H_i\left(\mathfrak{X}^b\right)\right) \tag{6}$$

for any a < b, where $\mathfrak{X}^a = d^{-1}(-\infty, a] \subset \mathfrak{X}$. The relevant barcodes for these groups are shown in Figure 2. Let $\pi : \mathfrak{X} \to \mathcal{D}_1$ be the projection $\pi(x, y) = x$, and let

 $\mathcal{A} = \pi^{-1}(\partial \mathcal{D}_1) = \pi^{-1}(\{0,1\}) = \{0,1\} \times [0,1]$



Figure 1: A source and target function, shown in 1a and 1b, where the source is on the vertical axis and the target is on the horizontal axis. A correspondence is an increasing continuous function T satisfying $d(x, T(x)) \leq a$. For a relatively small, there are essentially three for these two signals, corresponding to three up and right moving paths from the bottom to the

top of 1c avoiding the shaded regions, for which d(x, y) > a.



Figure 2: The barcodes for the zeroth and first persistent homology groups of the function $d: [0,1] \times [0,1] \to \mathbb{R}$ from Figure 1c.

be the inverse image of the boundary $\partial \mathcal{D}_1$. We now similarly have persistent homology for the restriction of d to \mathcal{A} , and also the relative homology groups $H_i^{a,b}(\mathfrak{X},\mathcal{A})$.

A correspondence then induces a continuous section $\tilde{T}: \mathcal{D}_1 \to \mathfrak{X}$ satisfying

$$\pi(\tilde{T}(x)) = x, \quad d(x, \tilde{T}(x)) \le a_{x}$$

for all x. It induces a persistent homology class

$$[T]_{a,b} = \tilde{T}_*([\mathcal{D}_1]) \in \operatorname{im}\left(H_1(\mathcal{X}^a, \mathcal{A}^a) \to H_1^{a,b}(\mathcal{X}, \mathcal{A})\right)$$
(7)

for any b > a, where

$$[\mathcal{D}_1] \in H_1(\mathcal{D}_1, \partial \mathcal{D}_1) \cong H_1(S^1)$$

is a generator of the first homology of $\partial \mathcal{D}_1$, which is equivalent to the circle. If $T \sim T'$ are homotopy equivalent in $(\mathfrak{X}^b, \mathcal{A}^b)$, meaning T can be deformed into T' without ever leaving \mathfrak{X}^b , keeping the endpoints in \mathcal{A}^b , they will induce the same class $[T]_{a,b} = [T']_{a,b}$. We also have the image

$$[\partial T]_{a,b} = \delta\left([T]_{a,b}\right) \in H_0^{a,b}(\mathcal{A}).$$
(8)

under the connecting homomorphism from the long exact sequence (2). The elements $[T]_{a,b}$ and $[\partial T]_{a,b}$ are the one-dimensional version of the classifiers referred to in the introduction.

In the case of two-dimensional image correspondence, we have few differences. First, the domain of an image is now $\mathcal{D}_1 = \mathcal{D}_2 = [0,1] \times [0,1]$ as in Section 2.2. We would then expect the space \mathcal{X} from Section 3.1 to have dimension 4. But by the epipolar condition that corresponding points have the same *y*-coordinate, we instead take

$$\mathfrak{X} = \{ (p,q) \in \mathfrak{D}_1 \times \mathfrak{D}_2 : p_2 = q_2 \},\$$

which is three-dimensional. The preimage of the boundary $\partial \mathcal{D}_1$ is now a cylinder, $\mathcal{A} = \pi^{-1}(\partial \mathcal{D}_1) = S^1 \times [0, 1]$, where $\pi(p, q) = p$, and S^1 is identified with the boundary of the square. The dimensions of the persistent homology groups are moved up by one dimension,

$$[T]_{a,b} \in H_2(\mathfrak{X}, \mathcal{A}), \quad [\partial T]_{a,b} \in H_1(\mathcal{A}).$$

We have found that the second class $[\partial T]_{a,b}$ is favorable in practice, for one thing because homology in one-dimension is a smaller computation. In examples of correspondence, we will be interested in the persistence diagram of the filtered vector space containing this class, which is

$$\operatorname{im}\left(H_2(\mathfrak{X},\mathcal{A})\to H_1(\mathcal{A})\right) = \operatorname{ker}\left(H_1(\mathcal{A})\to H_1(\mathfrak{X})\right).$$
(9)

Its elements may be thought of as classes of correspondences on the boundary $\partial \mathcal{D}_1$, which may be extended to the interior, and therefore map to zero in $H_1(\mathfrak{X})$. The setup of each of our examples in Section 4 is to present a correspondence problem, exhibit the persistence diagram associated to the filtered vector space on the right side of (9), and show that features of the diagram, i.e. long bars, correspond to pertinent features of the image pair.

3.2 Definition of the filtered complex

The preceding section shows how persistent homology classes can in principal be used to classify image correspondences up to continuous deformation, if the relative persistent homology groups $H_i^{a,b}(\mathcal{X},\mathcal{A})$ can be computed in practice. In this section, we construct a simplicial complex X representing \mathcal{X} , together with a filtration function $f: X \to \mathbb{R}_{\geq 0}$ in place of d. We also define a simplicial complex K associated to a triangulation of the base \mathcal{D}_1 , and a complex $A \subset X$ representing $\mathcal{A} \subset \mathcal{X}$, which is also filtered by f.

Choosing a simplicial complex to represent a space is an interesting problem in general. The motivation for the construction of this paper is that it is fibered as a complex of the base $X \to K$ representing $\pi : \mathfrak{X} \to \mathcal{D}_1$, so that a correspondence T determines a section $\varphi : K \to X$. It is therefore possible to define

$$[\varphi]_{a,b} \in H_2^{a,b}(X,A), \quad [\partial\varphi]_{a,b} \in H_1^{a,b}(A)$$

analogous to (7) and (8). A biproduct of doing this is that X does not come from a triangulation of \mathfrak{X} , nor is it embeddable in \mathbb{R}^3 .

To do this, begin with a triangulation $p: |K| \to \mathcal{D}_1$, and let $\partial K \subset K$ denote the boundary, as in Section 2.2. The underlying complex of X is given as follows.

1. Choose some natural numbers N_i for every vertex $i \in K_0$, and define the set of 0-simplices as

$$X_0 = \{(i, j) : i \in K_0, \ 1 \le j \le N_i\}.$$

- 2. Include every 2-simplex $\Delta = [(i_0, j_0), (i_1, j_1), (i_2, j_2)]$ for which $[i_0, i_1, i_2]$ is a 2-simplex in K_2 , and call these the horizontal faces of X.
- 3. Add every 3-simplex $[(i_0, j_0), (i_1, j_1), (i_2, j_2), (i_3, j_3)]$ satisfying
 - (a) There are only three distinct elements in $\{i_0, i_1, i_2, i_3\}$, and they are the vertices of a 2-simplex in K.
 - (b) If $i_a = i_b$ for $a \neq b$, then $j_a = j_b \pm 1$.

In other words, we have added a 3-simplex whenever it includes two horizontal faces differing only in one coordinate by a j-value of one.

4. Include all the faces of every simplex added thus far, making X a legitimate simplicial complex. The 2-simplices that have been added as a result will only contain two distinct *i*-values, and will be called vertical faces.

The effect of adding the 3-simplices in item 3 is to "fill in" the space between correspondences. Notice that there is an obvious surjective map $\pi : X \to K$ of complexes

$$\pi([(i_0, j_0), \dots, (i_k, j_k)]) = [i_0, \dots, i_k],$$

that forgets the *j*-values. Let $A = \pi^{-1}(\partial K)$ be the subcomplex of X whose *i*-values lie in the boundary $\partial K \subset K$.

We next define a filtration on this complex. For every horizontal 2simplex, define

$$f([(i_0, j_0), (i_1, j_1), (i_2, j_2)]) = \min_{j_a - 1 \le s_a \le j_a} d\left([p_0, p_1, p_2], [q_{p_0}(s_0), q_{p_1}(s_1), q_{p_2}(s_2)]\right),$$
(10)

where d is choice of distance measure from Section 2.2. On every 3-simplex, we define the value of f to be the minimum of the two horizontal 2-simplices

that are its faces. For every remaining simplex, the weight is inductively defined as the maximum of all simplices for which it is a face. We also obtain a filtration on A by restriction.

As described in Section 3.1, the persistent homology groups we are interested in are the persistence digrams for the kernel of the map induced by the inclusion

$$\ker\left(H_1(A) \to H_1(X)\right). \tag{11}$$

In order to generate a persistence diagram in JavaPlex we will make use of a workaround introduced in [4], which was used to study the image (11) in the case where A is the ideal Klein bottle, and X is a filtered complex representing a dataset of natural images. In this setup, we select an parameter $t_0 > 0$, and define X_{t_0} to have the same underlying complex as X, but where the persistence value of all interior simplices $\Delta \in X - A$ (which includes all horizontal simplices) begin at t_0 . This encodes the map $A \to X$ into single complex, by having the persistence values in the interior X - A begin at t_0 , by simply shifting the persistence values. We have found this approach to be sufficient for our purposes, though in future applications we expect to study the kernel in (9) directly, for instance using persistence for kernels and images in Dionysus [7].

We now describe how to interpret the persistence diagram of X_{t_0} in terms of image correspondences. For each image pair, the persistence diagram will show the following types of bar:

- 1. Multiple short bars: these may be disregarded as noise.
- 2. Multiple long bars, which begin to the right of the chosen offset parameter t_0 : these represent partial solutions in some subregion in the interior, that do not extend to the boundary values.
- 3. Long bars, with left endpoint slightly greater than zero, and right endpoint slightly greater than t_0 : they represent elements of the kernel (11), which come from correspondences of the entire source image. In other words, they represent true solutions to the correspondence problem.
- 4. Even longer bars whose right endpoint is significantly greater than t_0 : They represent correspondences near the periphery of the source image, but do not extend to the entire diagram, and so are not in the kernel (11). In other words, they are partial solutions which solve the

correspondence problem near the boundary of the image, but which do not match on some regions in the interior.

3.3 Pipeline

Here we give an explicit description of the complex X_{t_0} that determines classes of correspondence between two images. We assume that we are given two images I_1 and I_2 , as indicated in Figure 3a. The construction is as follows:

- 1. Let L be a collection of evenly-distributed landmark points in \mathcal{D}_1 , such as a hexagonal lattice (Figure 3b).
- 2. Associated with each landmark point $p \in L$, we have a collection of possible images of that landmark point Q(p) in the domain of I_2 , which are the possible q-values from the last section. The set Q could be infinite or finite, and is determined by some prior knowledge about the camera placement or other initial pre-processing. In our example, we have two cameras that are horizontally aligned, so we restrict Q(p) to a horizontal interval, which we discretize to obtain a finite complex, as shown in Figure 3c. The set Q partially characterizes the "niceness" of the correspondence map that we seek, by restricting the plausible locations that we think that p could land under a "nice" mapping T. The 0-simplices of the complex X are the union $\bigcup_{p \in L} Q(p)$.
- 3. Build a Delaunay triangulation of L. For each triangle $\mathcal{T} = (x, y, z)$ in the triangulation, and for every triangle \mathcal{T}' of the form (u, v, w), with $u \in Q(x), v \in Q(y)$, and $w \in Q(z)$, we do the following:

If the shape of \mathcal{T}' is "similar" to that of \mathcal{T} , then add a 2-simplex to X (by "similar", we mean for example that the lengths of the perimeters of \mathcal{T} and \mathcal{T}' are not too different, and that their orientations are the same). Its persistence value is $d(\mathcal{T}, \mathcal{T}')$, where $d(\cdot, \cdot)$ is the distance measure in Equation (4), or any similar counterpart, such as earth mover's distance. We have done this instead of the minimum taken in equation (10) only in the interest of speed, and because in our examples the values do not vary much in that domain, making this an acceptable approximation. In more sensitive applications we would expect to produce each value by solving an



(a) Source and target images



(b) Landmark points in the source



(c) A landmark point p in the source and the set Q(p) in the target.



(d) A few triangles \mathcal{T}' associated with the magnified triangle \mathcal{T} in the source. The source triangle \mathcal{T} is an element of the Delaunay triangulation of the landmark points.

Figure 3: The construction of our proposed complex. The source and target images are slightly skewed from one another, to suggest different camera angles of the same scene. The middle of the five triangles in the target image of 3d will have the lowest value of $d(\mathcal{T}, \mathcal{T}')$.

actual optimization problem in a highly localized domain. This obviously adds component 1- and 0-simplices as well, whose persistence values are not defined yet.

This process is shown in Figure 3d.

- 4. Set the persistence values of the lower-dimensional simplices in the natural way: each 1-simplex has a persistence value equal to the minimum persistence value among all 2-simplices containing it, and each 0-simplex has a persistence value equal to the minimum persistence value among all 1-simplices containing it.
- 5. Choose a value $t_0 > 0$ which is somewhere between the expected left and right endpoints of the important bars in the persistence diagram. For each simplex Δ (of any dimension), if Δ contains an interior vertex, add t_0 to its persistence value. Call the complex with the new persistence values X_{t_0} .

4 Practical Examples

In this section present three example applications, following the description of the pipeline from Section 3.3. Each example contains an image pair and a description of the correspondence problem, technical information such as the choice of the dissimilarity function and the value of the t_0 parameter, and the persistence diagram for the complex X_{t_0} described at the end of Section 3.2.

4.1 Example: identical black discs

Our first example is primarily a conceptual warmup example which illustrates some of the interesting features in Example 4.3. We choose our source and target to be the two identical black opaque circles shown in Figure 4.

The source image is triangulated by a complex K with 100 equilateral triangles. Our distance function is the L^2 -distance from (4). We remove all 2-simplices of X which are either orientation reversed, or are sheared to a width of more than double that of the based triangle, by setting the persistence score to infinity. The t_0 parameter is chosen to be 4.6×10^4 . In place of the desired min in equation (5), we take a rough approximation of the





(a) Source image (with Delaunay triangulation)

(b) Target image

Figure 4: Two identical source and target images

min in (10) by only sampling a single value for speed purposes as described in Section 3.3.

The persistence diagram is shown in Figure 5. The important information is that there are three long bars, shown bolder in the picture, beginning before t_0 . Two of these continue well beyond t_0 , while the lower one stops almost immediately after it. The bar that stops near t_0 represents correspondences which correctly correspond points in the entire image. The two longer bars correspond to two types of correspondence of the boundary which cannot be extended to the entire picture, represented by those which carry the boundary of the source entirely in the white space to the left of the black disc in the target, and those that are entirely on the right.

4.2 Example: dot mesh

The second correspondence problem is to find a mapping φ that preserves vertical coordinates between the point clouds shown in Figures 6a and 6b. Although not remotely apparent to the naked eye, there are actually two correspondences between the two, as illustrated in Figures 6c, 6d and 6e. The two triangulations show that there are essentially two different ways to map the source points into the second.

Our distance measure is defined by

$$E(\varphi) = \max_{p \in A} \min_{q \in B} ||p - q||.$$
(12)

For speed, we have simply replaced the maximum by the values at a sample



Figure 5: The JavaPlex barcode diagram for the images of black discs in Figure 4.



Figure 6: Two different correspondences of a mesh filled with a point cloud into the same target.

we did for the circle example. The 2-simplices in X whose width is more than double that of the associated triangles in the domain of the source image \mathcal{D}_1 are effectively removed by setting their persistence score to infinity, and also those triangles in which the orientation has been reversed. The t_0 parameter is set to 180. Figure 7 shows the complexes X^a for a few choices of a. Notice that the interior triangles begin appearing in Figure 7c, once the t_0 parameter is passed.

The persistence diagram shown in Figure 8 shows two long bars, which detect the two possible types of correspondence shows in Figures 6d and 6e.

4.3 Example: Pennies on a tablecloth

Next consider two real images of a penny sitting on a tablecloth with an interesting pattern shown in Figure 9. We began by quantizing color values using a filter, and then used a Wasserstein distance function. We eliminated highly skewed triangles using the same criteria from Examples 4.1 and 4.2.



Figure 7: Some diagrams showing the complex $X^{< a}$ for a few values of a for the dot mesh shown in Figure 6.



Figure 8: The JavaPlex barcode diagram for the images in Figure 6.



(a) Source image (b) Target image

Figure 9: Two pictures of a penny on a tablecloth



Figure 10: The JavaPlex barcode diagram for the images in Figure 9

The parameter is $t_0 = 2.6 \times 10^4$.

The persistence diagram is shown in Figure 10. We have removed all barcodes of length < 30 for visual purposes. The first Betti number output appears similar to the first example, which is to be expected. However, beyond the extra noise due to the fact that have used a picture from real life, this image has additional features due to the repeating pattern of the tablecloth:

1. We now have interesting Betti zero features (connected components of X) corresponding to the different translations of the tablecloth pattern. They represent the fact that the boundary of the source image can be matched to three different positions on the target.

2. The Betti one features are the same as in Example 4.1, but for a different reason. This time every correspondence of the boundary encircles the penny, but there are three different classes of them due to the pattern. If the left and right borders of the frame in the source are moved in even slightly, so that the border is in the black stripe region at the same height as the source penny, the two additional Betti one barcodes will disappear. This is because the penny in the target will raise the persistence value even for boundary correspondences.

5 A deformation invariance result

We now define the classes associated to a correspondence φ described in Section 3.1, and prove that they are invariant under certain smooth deformations in Theorem 1. The existence of a theorem such as this one is the reason the discussion in each example in Section 4 is valid, as they are implicitly describing classes of correspondences which are equivalent up to homotopy.

Let us call φ an *a*-correspondence if $E(\varphi) < a$, denote the set of these elements by $\Gamma(K, a)$. Let us also define $X^{<a}$ in the same way as X^a but with strict inequality, and the same for $A^{<a}$. Assume the dissimilarity function *d* is continuous as a function of the six vertices of a pair of triangles, so that we may regard $\Gamma(K, a)$ as an open subset $U \subset \mathbb{R}^N$, where *N* is the number of landmark points, i.e. zero simplices in K_0 .

We will say that φ, φ' are *b*-equivalent, and write $\varphi \sim_b \varphi'$ if there exists a continuous function $h : [0, 1] \to \Gamma(K, b)$ satisfying $h(0) = \varphi, h(1) = \varphi'$. for all $s \in I$. In other words, they are in the same connected component of $\Gamma(K, b) \subset \mathbb{R}^N$, since path connected and connected are the same for open subsets. We will be interested in studying the set $\Gamma(K, a)/\sim_b$ of *a*-correspondences up to *b*-equivalence for $a \leq b$, the set being empty otherwise. For instance, $\Gamma(K, a)/\sim_a$ is just the discrete set of connected components of the open set U, whereas choosing b > a results in collapsing more of these components. If $\varphi \in \Gamma(K, a)$, the equivalence class of φ will be written $\tilde{\varphi}_{a,b} \in \Gamma(K, a)/\sim_b$.

For any correspondence $\varphi \in \Gamma(K, a)$, we have an injective map of complexes $s^{\varphi} : K \to X$ by

$$s^{\varphi}([i_0, ..., i_k]) = [(i_0, j_0), ..., (i_k, j_k)], \quad j_a = \lfloor N_{i_a}\varphi(i_a) \rfloor.$$
(13)

It is a section of π , meaning $\pi s^{\varphi} = \text{Id}$, which in particular implies that $s^{\varphi}(\partial K) \subset A$. By (10), we can see that if $\varphi \in \Gamma(K, a)$, then $s^{\varphi}(K) \subset X^{\leq a}$.

Definition 1. For any *a*-correspondence $\varphi \in \Gamma(K, a)$, let

$$[\varphi]_a = s^{\varphi}_* ([K]) \in H_2(X^{< a}, A^{< a}), \quad [\partial \varphi]_a = \delta ([\varphi]_a) \in H_1(A^{< a}),$$

where δ is the connecting homomorphism from (2). Let $[\varphi]_{a,b}, [\partial \varphi]_{a,b}$ denote their images in $H_2^{a,b}(X, A), H_1^{a,b}(A)$, where in this section we use the persistent homology groups using strict inequality, i.e. replacing X^a, A^a with $X^{\leq a}, A^{\leq a}$.

Theorem 1. If $\varphi, \varphi' \in \Gamma(K, a)$, and $\varphi \sim_b \varphi'$, then $[\varphi]_{a,b} = [\varphi']_{a,b}$, and $[\partial \varphi]_{a,b} = [\partial \varphi']_{a,b}$. In particular, we have well-defined maps

$$\Gamma(K,a)/\sim_b \to H_2^{a,b}(X,A) \to H_1^{a,b}(A),$$

sending $\tilde{\varphi}_{a,b}$ to $[\varphi]_{a,b}$, and $[\partial \varphi]_{a,b}$, respectively. Furthermore, for any φ , both of these classes are nonzero.

Corollary 1. If either $H_2^{a,b}(X,A)$, or $H_1^{a,b}(A)$ are trivial, then $\Gamma(K,a)$ is empty.

We need a simple lemma:

Lemma 2. If p, q are in the same connected component of an open subset $U \subset \mathbb{R}^N$, there is a continuous function $h : [0, 1] \to U$ satisfying

- a) It connects the points h(0) = p, h(1) = q.
- b) For every $1 \le i \le N$, the set of points $t \in [0, 1]$ with coordinate function $h_i(t) \in \mathbb{Z}$ is finite.
- c) No point in the path has two integer-valued coordinates simulatenously, i.e. $h_i(t)$ and $h_j(t)$ are not both integers for $i \neq j$.

Proof. Since connected implies path connected for open sets, we may suppose there is a continuous function satisfying part a), and by Whitney approximation, we may assume that it is smooth. Part b) can easily be obtained using standard transversality results, noting that the set of points where the *i*th coordinate is integral is a codimension one submanifold of U, and a submanifold of [0, 1] of dimension zero is finite. See for instance [5]. Part c) is then easy to obtain by a simple perturbation argument.

We now move onto the proof of Theorem 1.

Proof. Let $U = \Gamma(K, a)$ be as above, so that φ, φ' define points p, q in the same connected component of U. Suppose we have a function h as in Lemma 2. Then applying the definition of j_k in (13) to h(s) gives functions $j_k : [0,1] \to \{1, ..., N_k\}$ such that $j_k(s)$ has jumps from one natural number to an adjacent one at certain values of s. By the lemma, we may suppose that each $j_k(s)$ has jumps at finitely many values of s, and that $j_k(s)$ and $j_l(s)$ never have simultaneous jumps for $k \neq l$.

We now have a sequence $\varphi = \varphi_0, ..., \varphi_n = \varphi' \in \Gamma(K, b)$ by taking $\varphi_i = h(s_i)$ where $s_0 = 0$, $s_n = 1$, and s_i is any value between the between the *i*th and (i + 1)st jump for $1 \leq i \leq n - 1$. It now suffices to show that $[\varphi_i]_{a,b} - [\varphi_{i+1}]_{a,b} = 0$ for all $i \in \{0, ..., n - 1\}$. Given any such *i*, let *k* be the coordinate that jumps between *i* and i + 1. Then we have

$$s_*^{\varphi_i}([K]) - s_*^{\varphi_{i+1}}([K]) = \pm \partial \left[\operatorname{St}[(k, j_k(s_i)), (k, j_k(s_{i+1}))] \right]$$

where $[\operatorname{St} \Delta]$ is the class of the star of Δ , in other words the sum of all 3simplices which have Δ as a face, which in the equation is a 1-simplex. By construction, each of these 3-simplices are contained in X, since the second coordinates differ by 1. They are also in $X^{\leq b}$, as we may take φ_i as an upper bound for the the min in (10).

Finally, we check that neither class vanishes. Since $\pi i^{a,b} s^{\varphi} : K \to K$ is the identity map, we have

$$\pi_*([\partial\varphi]_{a,b}) = \pi_*(\delta([\varphi]_{a,b})) = \delta([K]) = [\partial K] \in H_1(\partial K),$$

which is nonzero, so that $[\partial \varphi]_{a,b} \neq 0$. Since $[\partial \varphi]_{a,b} = \delta([\varphi]_{a,b})$, we must have $[\varphi]_{a,b} \neq 0$ as well.

6 Conflict of interest statement

On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

[1] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. *ECCV*, 2006.

- Gunnar Carlsson and Afra Zomorodian. Computing persistent homology. Discrete Comput. Geom, 33:249–274, 2004.
- [3] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification, 2000.
- [4] Carlsson G., T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. Int J Comput Vis, 76:1–12, 2008.
- [5] V. Guillemin and A. Pollack. *Differential Topology*. AMS Chelsea Publishing. AMS Chelsea Pub., 2010.
- [6] David G. Lowe. Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision, 60:91–110, 2004.
- [7] D. Morozov. Dionysus documentation. http://www.mrzv.org/ software/dionysus/.
- [8] Andrew Tausz, Mikael Vejdemo-Johansson, and Henry Adams. JavaPlex: A research software package for persistent (co)homology. In Han Hong and Chee Yap, editors, *Proceedings of ICMS 2014*, Lecture Notes in Computer Science 8592, pages 129–136, 2014. Software available at http://appliedtopology.github.io/javaplex/.
- [9] Afra Joze Zomorodian. Computing and comprehending topology: Persistence and hierarchical morse complexes, 2001. Ph.D. thesis, University of Illinois at Urbana-Champaign.