

Persistent Cohomology and Circular Coordinates

Vin de Silva^{*}
 Department of Mathematics
 Pomona College
 Claremont, California
 vin.desilva@pomona.edu

Mikael Vejdemo-Johansson[†]
 Department of Mathematics
 Stanford University
 Stanford, California
 mik@math.stanford.edu

ABSTRACT

Nonlinear dimensionality reduction (NLDR) algorithms such as Isomap, LLE and Laplacian Eigenmaps address the problem of representing high-dimensional nonlinear data in terms of low-dimensional coordinates which represent the intrinsic structure of the data. This paradigm incorporates the assumption that real-valued coordinates provide a rich enough class of functions to represent the data faithfully and efficiently. On the other hand, there are simple structures which challenge this assumption: the circle, for example, is one-dimensional but its faithful representation requires two real coordinates. In this work, we present a strategy for constructing circle-valued functions on a statistical data set. We develop a machinery of persistent cohomology to identify candidates for significant circle-structures in the data, and we use harmonic smoothing and integration to obtain the circle-valued coordinate functions themselves. We suggest that this enriched class of coordinate functions permits a precise NLDR analysis of a broader range of realistic data sets.

Categories and Subject Descriptors

G.3 [Probability and statistics]: multivariate statistics;
 I.5.1 [Pattern recognition]: models—*geometric*

General Terms

algorithms, theory

Keywords

dimensionality reduction, computational topology, persistent homology, persistent cohomology

^{*}VdS has been partially supported by DARPA, through grants HR0011-07-1-0002 and HR0011-05-1-0007.

[†]MVJ has been partially supported by the Office of Naval Research, through grant N00014-08-1-0931.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'09, June 8–10, 2009, Aarhus, Denmark.

Copyright 2009 ACM 978-1-60558-501-7/09/06 ...\$5.00.

1. INTRODUCTION

Nonlinear dimensionality reduction (NLDR) algorithms address the following problem: given a high-dimensional collection of data points $X \subset \mathbb{R}^N$, find a low-dimensional embedding $\phi : X \rightarrow \mathbb{R}^n$ (for some $n \ll N$) which faithfully preserves the ‘intrinsic’ structure of the data. For instance, if the data have been obtained by sampling from some unknown manifold $M \subset \mathbb{R}^N$ — perhaps the parameter space of some physical system — then ϕ might correspond to an n -dimensional coordinate system on M . If M is completely and non-redundantly parametrized by these n coordinates, then the NLDR is regarded as having succeeded completely.

Principal components analysis, or linear regression, is the simplest form of dimensionality reduction; the embedding function ϕ is taken to be a linear projection. This is closely related to (and sometimes identified with) classical multidimensional scaling [2].

When there are no satisfactory linear projections, it becomes necessary to use NLDR. Prominent algorithms for NLDR include Locally Linear Embedding [14], Isomap [16], Laplacian Eigenmaps [1], Hessian Eigenmaps [5], and many more.

These techniques share an implicit assumption that the unknown manifold M is well-described by a finite set of coordinate functions $\phi_1, \phi_2, \dots, \phi_n : M \rightarrow \mathbb{R}$. Explicitly, some of the correctness theorems in these studies depend on the hypothesis that M has the topological structure of a convex domain in some $\mathbb{R}^n$. This hypothesis guarantees that good coordinates exist, and shifts the burden of proof onto showing that the algorithm recovers these coordinates.

In this paper we ask what happens when this assumption fails. The simplest space which challenges the assumption is the circle, which is one-dimensional but requires two real coordinates for a faithful embedding. Other simple examples include the annulus, the torus, the figure eight, the 2-sphere, the last three of which present topological obstructions to being embedded in the Euclidean space of their natural dimension. We propose that an appropriate response to the problem is to enlarge the class of coordinate functions to include circle-valued coordinates $\theta : M \rightarrow S^1$. In a physical setting, circular coordinates occur naturally as angular and phase variables. Spaces like the annulus and the torus are well described by a combination of real and circular coordinates. (The 2-sphere is not so lucky, and must await its day.)

The goal of this paper is to describe a natural procedure for constructing circular coordinates on a nonlinear data set using techniques from classical algebraic topology and its

21st-century grandchild, persistent topology. We direct the reader to [9] as a general reference for algebraic topology, and to [17] for a streamlined account of persistent homology.

1.1 Related work

There have been other attempts to address the problem of finding good coordinate representations of simple non-Euclidean data spaces. One approach [13] is to use modified versions of multidimensional scaling specifically devised to find the best embedding of a data set into the cylinder, the sphere and so on. The target space has to be chosen in advance. Another class of approaches [10, 4] involves cutting the data manifold along arcs and curves until it has trivial topology. The resulting configuration can then be embedded in Euclidean space in the usual way. In our approach, the number of circular coordinates is not fixed in advance, but is determined experimentally after a persistent homology calculation. Moreover, there is no cutting involved; the coordinate functions respect the original topology of the data.

1.2 Overview

The principle behind our algorithm is the following equation from homotopy theory, valid for topological spaces X with the homotopy type of a cell complex (which covers everything we normally encounter):

$$[X, S^1] = H^1(X; \mathbb{Z}) \quad (1)$$

The left-hand side denotes the set of equivalence classes of continuous maps from X to the circle S^1 ; two maps are equivalent if they are homotopic (meaning that one map can be deformed continuously into the other); the right-hand side denotes the 1-dimensional cohomology of X , taken with integer coefficients. In other language: S^1 is the classifying space for H^1 , or equivalently S^1 is the Eilenberg–MacLane space $K(\mathbb{Z}, 1)$. See section 4.3 of [9].

If X is a contractible space (such as a convex subset of $\mathbb{R}^n$), then $H^1(X; \mathbb{Z}) = 0$ and Equation (1) tells us not to bother looking for circular functions: all such functions are homotopic to a constant function. On the other hand, if X has nontrivial topology then there may well exist a nonzero cohomology class $[\alpha] \in H^1(X; \mathbb{Z})$; we can then build a continuous function $X \rightarrow S^1$ which in some sense reveals $[\alpha]$.

Our strategy divides into the following steps.

1. Represent the given discrete data set as a simplicial complex or filtered simplicial complex.
2. Use persistent cohomology to identify a ‘significant’ cohomology class in the data. For technical reasons, we carry this out with coefficients in the field $\mathbb{F}_p$ of integers modulo p , for some prime p . This gives us $[\alpha_p] \in H^1(X; \mathbb{F}_p)$.
3. Lift $[\alpha_p]$ to a cohomology class with integer coefficients: $[\alpha] \in H^1(X; \mathbb{Z})$.
4. Smoothing: replace the integer cocycle α by a harmonic cocycle in the same cohomology class: $\bar{\alpha} \in C^1(X; \mathbb{R})$.
5. Integrate the harmonic cocycle $\bar{\alpha}$ to a circle-valued function $\theta : X \rightarrow S^1$.

The paper is organized as follows. In Section 2.1, we derive what we need of equation (1). Steps (1–5) of the algorithm are addressed in Sections 2.2–2.6, respectively. In Section 3 we report some experimental results.

2. ALGORITHM DETAILS

2.1 Cohomology and circular functions

Let X be a finite simplicial complex. Let X^0, X^1, X^2 denote the sets of vertices, edges and triangles of X , respectively. We suppose that the vertices are totally ordered (in an arbitrary way). If $a < b$ then the edge between vertices a, b is always written ab and not ba . Similarly, if $a < b < c$ then the triangle with vertices a, b, c is always written abc .

Cohomology can be defined as follows. Let $\mathbb{A}$ be a commutative ring (for example $\mathbb{A} = \mathbb{Z}, \mathbb{F}_p, \mathbb{R}$). We define 0-cochains, 1-cochains, and 2-cochains as follows:

$$\begin{aligned} C^0 &= C^0(X; \mathbb{A}) &= \{ \text{functions } f : X^0 \rightarrow \mathbb{A} \} \\ C^1 &= C^1(X; \mathbb{A}) &= \{ \text{functions } \alpha : X^1 \rightarrow \mathbb{A} \} \\ C^2 &= C^2(X; \mathbb{A}) &= \{ \text{functions } A : X^2 \rightarrow \mathbb{A} \} \end{aligned}$$

These are modules over $\mathbb{A}$. We now define coboundary maps $d_0 : C^0 \rightarrow C^1$ and $d_1 : C^1 \rightarrow C^2$.

$$\begin{aligned} (d_0 f)(ab) &= f(b) - f(a) \\ (d_1 \alpha)(abc) &= \alpha(bc) - \alpha(ac) + \alpha(ab) \end{aligned}$$

Let $\alpha \in C^1$. If $d_1 \alpha = 0$ we say that α is a *cocycle*. If $d_0 f = \alpha$ admits a solution $f \in C^0$ we say that α is a *coboundary*. The solution f , if it exists, can be thought of as the discrete integral of α . It is unique up to adding constants on each connected component of X .

It is easily verified that $d_1 d_0 f = 0$ for any $f \in C^0$. Thus, coboundaries are always cocycles, or equivalently $\text{Im}(d_0) \subseteq \text{Ker}(d_1)$. We can measure the difference between coboundaries and cocycles by defining the 1-cohomology of X to be the quotient module

$$H^1(X; \mathbb{A}) = \text{Ker}(d_1) / \text{Im}(d_0).$$

We say that two cocycles α, β are *cohomologous* if $\alpha - \beta$ is a coboundary.

We now consider integer coefficients. The following proposition fulfils part of the promise of equation (1), by producing circle-valued functions from integer cocycles. It will be helpful to think of S^1 as the quotient group $\mathbb{R}/\mathbb{Z}$.

PROPOSITION 1. *Let $\alpha \in C^1(X; \mathbb{Z})$ be a cocycle. Then there exists a continuous function $\theta : X \rightarrow \mathbb{R}/\mathbb{Z}$ which maps each vertex to 0, and each edge ab around the entire circle with winding number $\alpha(ab)$.*

PROOF. We can define θ inductively on the vertices, edges, triangles, ... of X . The vertices and edges follow the prescription in the statement of the proposition. To extend θ to the triangles, it is necessary that the winding number of θ along the boundary of each triangle abc is zero. And indeed this is $\alpha(bc) - \alpha(ac) + \alpha(ab) = d_1 \alpha(abc) = 0$. Since the higher homotopy groups of S^1 are all zero ([9], section 4.3), θ can then be extended to the higher cells of X without obstruction. $\square$

The construction in Proposition 1 is unsatisfactory in the sense that all vertices are mapped to the same point. All variation in the circle parameter takes place in the interior of the edges (and higher cells). This is rather unsmooth. For more leeway, we consider real coefficients.

PROPOSITION 2. *Let $\bar{\alpha} \in C^1(X; \mathbb{R})$ be a cocycle. Suppose we can find $\alpha \in C^1(X; \mathbb{Z})$ and $f \in C^0(X; \mathbb{R})$ such that $\bar{\alpha} =$*

$\alpha + d_0 f$. Then there exists a continuous function $\theta : X \rightarrow \mathbb{R}/\mathbb{Z}$ which maps each edge ab linearly to an interval of length $\bar{\alpha}(ab)$, measured with sign.

In other words, we can construct a circle-valued function out of any real cocycle $\bar{\alpha}$ whose cohomology class $[\bar{\alpha}]$ lies in the image of the natural homomorphism $H^1(X; \mathbb{Z}) \rightarrow H^1(X; \mathbb{R})$.

PROOF. Define θ on the vertices of X by setting $\theta(a)$ to be $f(a) \bmod \mathbb{Z}$. For each edge ab , we have

$$\begin{aligned} \theta(b) - \theta(a) &= f(b) - f(a) \\ &= d_0 f(ab) \\ &= \bar{\alpha}(ab) - \alpha(ab) \end{aligned}$$

which is congruent to $\bar{\alpha}(ab) \bmod \mathbb{Z}$, since $\alpha(ab)$ is an integer.

It follows that θ can be taken to map ab linearly onto an interval of signed length $\bar{\alpha}(ab)$. Since $\bar{\alpha}$ is a cocycle, θ can be extended to the triangles as before; then to the higher cells. $\square$

Proposition 2 suggests the following tactic: from an integer cocycle α we construct a cohomologous real cocycle $\bar{\alpha} = \alpha + d_0 f$, and then define $\theta = f \bmod \mathbb{Z}$ on the vertices of X . If we can construct $\bar{\alpha}$ so that the edge-lengths $|\bar{\alpha}(ab)|$ are small, then the behaviour of θ will be apparent from its restriction to the vertices. See Section 2.5.

2.2 Point-cloud data to simplicial complex

We now begin describing the workflow in detail. The input is a point-cloud data set: in other words, a finite set $S \subset \mathbb{R}^N$ or more generally a finite metric space. The first step is to convert S into a simplicial complex and to identify a stable-looking integer cohomology class. This will occupy the next three subsections.

The first lesson of point-cloud topology [7] is that point-clouds are best represented by 1-parameter nested families of simplicial complexes. There are several candidate constructions: the Vietoris–Rips complex $X^\epsilon = \text{Rips}(S, \epsilon)$ has vertex set S and includes a k -simplex whenever all $k+1$ vertices lie pairwise within distance ϵ of each other. The witness complex $X^\epsilon = \text{Witness}(L, S, \epsilon)$ uses a smaller vertex set $L \subset S$ and includes a k -simplex when the $k+1$ vertices lie close to other points of S , in a certain precise sense (see [3, 8]). In both cases, $X^\epsilon \subseteq X^{\epsilon'}$ whenever $\epsilon \leq \epsilon'$. Either of these constructions will serve our purposes, but the witness complex has the computational advantage of being considerably smaller.

We determine X^ϵ only up to its 2-skeleton, since we are interested in H^1 .

2.3 Persistent cohomology

Having constructed a 1-parameter family $\{X^\epsilon\}$, we apply the principle of persistence to identify cocycles that are stable across a large range for ϵ . Suppose that $\epsilon_1, \epsilon_2, \dots, \epsilon_m$ are the critical values where the complex X^ϵ gains new cells. The family can be represented as a diagram

$$X^{\epsilon_1} \longrightarrow X^{\epsilon_2} \longrightarrow \dots \longrightarrow X^{\epsilon_m}$$

of simplicial complexes and inclusion maps. For any coefficient field $\mathbb{F}$, the cohomology functor $H^1(-; \mathbb{F})$ converts this diagram into a diagram of vector spaces and linear maps over $\mathbb{F}$; the arrows are reversed:

$$H^1(X^{\epsilon_1}; \mathbb{F}) \longleftarrow H^1(X^{\epsilon_2}; \mathbb{F}) \longleftarrow \dots \longleftarrow H^1(X^{\epsilon_m}; \mathbb{F})$$

According to the theory of persistence [6, 17], such a diagram decomposes as a direct sum of 1-dimensional terms indexed by half-open intervals of the form $[\epsilon_i, \epsilon_j)$. Each such term corresponds to a cochain $\alpha \in C^i(X^\epsilon)$ that satisfies the cocycle condition for $\epsilon < \epsilon_j$ and becomes a coboundary for $\epsilon < \epsilon_i$. The collection of intervals can be displayed graphically as a *persistence diagram*, by representing each interval $[\epsilon_i, \epsilon_j)$ as a point (ϵ_i, ϵ_j) in the Cartesian plane above the main diagonal. We think of long intervals as representing trustworthy (i.e. stable) topological information.

CHOICE OF COEFFICIENTS. The persistence decomposition theorem applies to diagrams of vector spaces over a field. When we work over the ring of integers $\mathbb{Z}$, however, the result is known to fail: there need not be an interval decomposition. This is unfortunate, since we require integer cocycles to construct circle maps. To finesse this problem, we pick an arbitrary prime number p (such as $p = 47$) and carry out our persistence calculations over the finite field $\mathbb{F} = \mathbb{F}_p$. The resulting $\mathbb{F}_p$ cocycle must then be converted to integer coefficients: we address this in Section 2.4.

In principle we can use the ideas in [17] to calculate the persistent cohomology intervals and then select a long interval $[\epsilon_i, \epsilon_j)$ and a specific $\delta \in [\epsilon_i, \epsilon_j)$. We then let $X = X^\delta$ and take α to be the cocycle in $C^1(X; \mathbb{F})$ corresponding to the interval.

Explicitly, persistent cocycles can be calculated in the following way. We thank Dmitriy Morozov for this algorithm. Suppose that the simplices in the filtered complex are totally ordered, and labelled $\sigma_1, \sigma_2, \dots, \sigma_m$ so that σ_i arrives at time ϵ_i . For $k = 0, 1, \dots, m$ we maintain the following information:

- a set of indices $I_k \subseteq \{1, 2, \dots, k\}$ associated with ‘live’ cocycles;
- a list of cocycles $(\alpha_i : i \in I_k)$ in $C^*(X^{\epsilon_k}; \mathbb{F})$.

The cocycle α_i involves only σ_i and those simplices of the same dimension that appear later in the filtration sequence (thus only σ_j with $j \geq i$).

Initially $I_0 = \emptyset$ and the list of cocycles is empty.

To update from $k-1$ to k , we compute the coboundaries of the cocycles $(\alpha_i : i \in I_{k-1})$ of $X^{\epsilon_{k-1}}$ within the larger complex X^{ϵ_k} obtained by including the simplex σ_k . In fact, these coboundaries must be multiples of the elementary cocycle $\alpha = [\sigma_k]$ defined by $\alpha(\sigma_k) = 1$, and $\alpha(\sigma_j) = 0$ otherwise. We can write $d\alpha_i = c_i[\sigma_k]$. If all the c_i are zero, then we have one new cocycle: let $I_k = I_{k-1} \cup \{k\}$ and define $\alpha_k = [\sigma_k]$. Otherwise, we must lose a cocycle. Let $j \in I_{k-1}$ be the largest index for which $c_j \neq 0$. We delete α_j by setting $I_k = I_{k-1} \setminus \{j\}$, and we restore the earlier cocycles by setting $\alpha_i \leftarrow \alpha_i - (c_i/c_j)\alpha_j$. In this latter case, we write the persistence interval $[\epsilon_j, \epsilon_k)$ to the output.

At the end of the process, surviving cocycles are associated with semi-infinite intervals: $[\epsilon_i, \infty)$ for $i \in I_m$.

REMARK. The reader may be more familiar with persistence diagrams in homology rather than cohomology. In fact, the universal coefficient theorem [9] implies that the two diagrams are identical. The salient point is that cohomology is the vector-space dual of homology, when working with field coefficients. That said, we cannot simply use the usual algorithm for persistent homology: we are interested in obtaining explicit cocycles, whereas the classical algorithm [17] returns cycles.

We will establish the correctness of this algorithm in the archival version of this paper. The expert reader may regard this as an exercise in the theory of persistence.

2.4 Lifting to integer coefficients

We now have a simplicial complex $X = X^\delta$ and a cocycle $\alpha_p \in C^1(X; \mathbb{F}_p)$. The next step is to ‘lift’ α_p by constructing an integer cocycle α which reduces to α_p modulo p .

To show that this is (almost) always possible, note that the short exact sequence of coefficient rings $0 \rightarrow \mathbb{Z} \xrightarrow{\cdot p} \mathbb{Z} \rightarrow \mathbb{F}_p \rightarrow 0$ gives rise to a long exact sequence, called the Bockstein sequence (see Section 3.E of [9]). Here is the relevant section of the sequence:

$$\rightarrow H^1(X; \mathbb{Z}) \rightarrow H^1(X; \mathbb{F}_p) \xrightarrow{\beta} H^2(X; \mathbb{Z}) \xrightarrow{\cdot p} H^2(X; \mathbb{Z}) \rightarrow$$

By exactness, the Bockstein homomorphism β induces an isomorphism between the cokernel of $H^1(X; \mathbb{Z}) \rightarrow H^1(X; \mathbb{F}_p)$ and the kernel of $H^2(X; \mathbb{Z}) \xrightarrow{\cdot p} H^2(X; \mathbb{Z})$, and this kernel is precisely the set of p -torsion elements of $H^2(X; \mathbb{Z})$. If there is no p -torsion, then it follows immediately that the cokernel of the first map is zero. In other words $H^1(X; \mathbb{Z}) \rightarrow H^1(X; \mathbb{F}_p)$ is surjective; any cocycle $\alpha_p \in C^1(X; \mathbb{F}_p)$ can be lifted to a cocycle $\alpha \in C^1(X; \mathbb{Z})$.

If we are unluckily sabotaged by p -torsion, then we pick another prime and redo the calculation from scratch: it is enough to pick a prime that does not divide the order of the torsion subgroup of $H^2(X; \mathbb{Z})$, so almost any prime will do.

In practice, we construct α by taking the coefficients of α_p in $\mathbb{F}_p$ and replacing them with integers in the correct congruence class modulo p . The default choice is to choose coefficients close to zero. If $d_1\alpha = 0$ then we are done; otherwise it becomes necessary to do some repair work. Certainly $d_1\alpha \equiv 0$ modulo p , so we can write $d_1\alpha = p\eta$ for some $\eta \in C^2(X; \mathbb{Z})$. In the absence of p -torsion, we can then solve $\eta = d_1\zeta$ for $\zeta \in C^1(X; \mathbb{Z})$, and then the required lift is $\alpha - p\zeta$. Fortunately, this has not proved necessary in any of our examples.

REMARK. We expect that p -torsion is extremely rare in ‘real’ data sets, since it is symptomatic of rather subtle topological phenomena. For instance, the simplest examples which exhibit 2-torsion are the nonorientable closed surfaces (such as the projective plane and the Klein bottle).

2.5 Harmonic smoothing

Given an integer cocycle $\alpha \in C^1(X; \mathbb{Z})$, or indeed a real cocycle $\alpha \in C^1(X; \mathbb{R})$, we wish to find the ‘smoothest’ real cocycle $\bar{\alpha} \in C^1(X; \mathbb{R})$ cohomologous to α . It turns out that what we want is the harmonic cocycle representing the cohomology class $[\alpha]$.

We define smoothness. Each of the spaces $C^i(X; \mathbb{R})$ comes with a natural Euclidean metric:

$$\begin{aligned} \|f\|^2 &= \sum_{a \in X^0} |f(a)|^2, \\ \|\alpha\|^2 &= \sum_{ab \in X^1} |\alpha(ab)|^2, \\ \|A\|^2 &= \sum_{abc \in X^2} |A(abc)|^2. \end{aligned}$$

A circle-valued function θ is ‘smooth’ if its total variation across the edges of X is small. The terms $|\alpha(ab)|^2$ cap-

ture the variation across individual edges; therefore what we must minimize is $\|\bar{\alpha}\|^2$.

PROPOSITION 3. *Let $\alpha \in C^1(X; \mathbb{R})$. There is a unique solution $\bar{\alpha}$ to the least-squares minimization problem*

$$\operatorname{argmin}_{\bar{\alpha}} \{ \|\bar{\alpha}\|^2 \mid \exists f \in C^0(X; \mathbb{R}), \bar{\alpha} = \alpha + d_0 f \}. \quad (2)$$

Moreover, $\bar{\alpha}$ is characterized by the equation $d_0^ \bar{\alpha} = 0$, where d_0^* is the adjoint of d_0 with respect to the inner products on C^0, C^1 .*

PROOF. Note that if $d_0^* \bar{\alpha} = 0$ then for any $f \in C^0$ we have

$$\begin{aligned} \|\bar{\alpha} + d_0 f\|^2 &= \|\bar{\alpha}\|^2 + 2\langle \bar{\alpha}, d_0 f \rangle + \|d_0 f\|^2 \\ &= \|\bar{\alpha}\|^2 + 2\langle d_0^* \bar{\alpha}, f \rangle + \|d_0 f\|^2 \\ &= \|\bar{\alpha}\|^2 + \|d_0 f\|^2 \end{aligned}$$

which implies that such an $\bar{\alpha}$ must be the unique minimizer. For existence, note that

$$d_0^* \alpha + d_0^* d_0 f = 0$$

certainly has a solution f if $\operatorname{Im}(d_0^*) = \operatorname{Im}(d_0^* d_0)$. But this is a standard fact in finite-dimensional linear algebra: $\operatorname{Im}(A^T) = \operatorname{Im}(A^T A)$ for any real matrix A ; this follows from the singular value decomposition, for instance. $\square$

REMARK. It is customary to construct the Laplacian $\Delta = d_1^* d_1 + d_0 d_0^*$. The twin equations $d_1 \bar{\alpha} = 0$ and $d_0^* \bar{\alpha} = 0$ immediately imply (and conversely, can be deduced from) the single equation $\Delta \bar{\alpha} = 0$; in other words $\bar{\alpha}$ is *harmonic*.

2.6 Integration

The least-squares problem in equation (2) can be solved using a standard algorithm such as LSQR [12]. By Proposition 2 we can use the solution parameter f to define the circular coordinate θ on the vertices of X . This works because the original cocycle α has integer coefficients.

More generally, if $\bar{\alpha}$ is an arbitrary real cocycle such that $[\bar{\alpha}] \in \operatorname{Im}(H^1(X; \mathbb{Z}) \rightarrow H^1(X; \mathbb{R}))$, it is a straightforward matter to integrate $\bar{\alpha}$ to a circle-valued function θ on the vertex set X^0 . Suppose that X is connected (if not, each connected component can be treated separately) and pick a starting vertex x_0 and assign $\theta(x_0) = 0$. One can use Dijkstra’s algorithm to find shortest paths to each remaining vertex from x_0 . When a new vertex b enters the structure via an edge ab , we assign $\theta(b) = \theta(a) + \bar{\alpha}(ab)$ (or $\theta(a) - \bar{\alpha}(ba)$ if the edge is correctly identified as ba). If a vertex a is connected to x_0 by multiple paths then the different possible values of $\theta(a)$ differ by an integer; this is where we use the hypothesis that $\bar{\alpha}$ is cohomologous to an integer cocycle.

3. EXPERIMENTS

3.1 Software

The following experiments were carried out using the Java-based jPlex simplicial complex software [15], with high-level scripting and numerical analysis in MATLAB. We ran a development version of jPlex to obtain explicit persistent cohomology cocycles. We expect to include the code in the next release of jPlex. We used Paige and Saunders’ implementation of LSQR [11] for the least-squares problem in the harmonic smoothing step. Timings were determined using MATLAB’s built-in ‘tic’ and ‘toc’ commands, and are included for relative comparison against each other.

3.2 General procedure

We tested our methods on several synthetic data sets with known topology, ranging from the humble circle itself to a genus-2 surface (‘double torus’). Most of the examples were embedded in $\mathbb{R}^2$ or $\mathbb{R}^3$, with the exception of a sample from a complex projective curve (embedded in $\mathbb{CP}^2$) and a synthetic image-like data set (embedded in $\mathbb{R}^{120000}$).

In each case we selected vertices for the filtered simplicial complex: either the whole set, or a smaller well-distributed subset of ‘landmarks’ selected by iterative furthest-point sampling. We then built a Rips or witness complex, with maximum radius generally chosen to ensure around 10^5 simplices in the complex.

In most cases, we show the persistence diagram produced by the cocycle computation. The chosen value δ is marked on the diagonal, with its upper-left quadrant indicated in green lines. The persistent cocycles available at that parameter value are precisely those contained in that quadrant. Each of those cocycles then produces a circular coordinate.

There are various figures associated with each example. Most important are the correlation scatter plots: each scatter plot compares two circular coordinate functions. These may be functions produced by the computation (‘inferred coordinates’) or known parameters. These scatter plots are drawn in the unit square, which is of course really a torus $S^1 \times S^1$.

When the original data are embedded in $\mathbb{R}^2$ or $\mathbb{R}^3$, we also display the circular coordinates directly on the data set, plotting each point in color according to its coordinate value interpreted on the standard hue-circle. This works less well in grayscale reproductions, of course.

Finally, in certain cases we plot coordinate values against frequency, as a histogram. This distributional information can sometimes be useful in the absence of other information.

REMARK. When the goal is to infer the topology of a data set whose structure is unknown, we do not have any ‘known parameters’ available to us. We can still construct correlation scatter plots between pairs of inferred coordinates, and the distributional histograms for each coordinate individually. We exhort the reader to view the following examples through the lens of the topological inference problem: what structures can be distinguished using scatter plots and histograms (and persistence diagrams) alone?

3.3 Noisy circle

We begin with the circle itself, and its tautological circle-valued coordinate.

We picked 400 points distributed along the unit circle. We added a uniform random variable from $[0.0, 0.4]$ to each coordinate. A Rips complex was constructed with maximal radius 0.5, resulting in 23475 simplices. The computation of cohomology finished in 237 seconds.

Parametrizing at 0.4 yielded a single coordinate function, which very closely reproduces the tautological angle function. Parametrizing at 0.14 yielded several possible cocycles. We selected one of those with low persistence; this produced a parametrization which ‘snags’ around a small gap in the data.

See Figure 1. The left panel in each row shows the histogram of coordinate values; the middle panel shows the correlation scatter plot against the known angle function; the right panel displays the coordinate using color. The high-

persistence (‘global’) coordinate correlates with the angle function with topological degree 1. Variation in that coordinate is uniformly distributed, as seen in the histogram. In contrast, the low-persistence (‘local’) coordinate has a spiky distribution.

3.4 Trefoil torus knot

Another example with circle topology: see Figure 2. We picked 400 points distributed along the $(2, 3)$ torus knot on a torus with radii 2.0 and 1.0. We jittered them by a uniform random variable from $[0.0, 0.2]$ added to each coordinate. We generated a Rips complex up to radius 1.0, acquiring 36936 simplices. We computed persistent cohomology in 70 seconds. As expected, the inferred coordinate correlates strongly with the known parameter with topological degree 1. The histogram shows three ‘bulges’ corresponding to the three high-density regions of the sampled curve, which occur when the curve approaches the central axis of the torus.

3.5 Rotating cube

For a more elaborate data set with S^1 -topology, we generated a sequence of 657 rendered images of a colorful cube rotating around one axis. Each image was regarded as a vector in the Euclidean space $\mathbb{R}^{200 \cdot 200 \cdot 3}$. From this data we built a witness complex with 50 landmark points and constructed a single circular coordinate. Interpolating the resulting function linearly between the landmarks gave us coordinates for all the points in the family.

See Figure 3. The frequency distribution is comparatively smooth (by which we mean that there are no large spikes in the histogram), which indicates that the coordinate does not have large static regions. The correlation plot of the inferred coordinate against the original known sequence of the cube images shows a correlation with topological degree 1. We show the progression of the animation on an evenly-spaced sample of representative points around the circle.

3.6 Pair of circles

See Figure 4 for these two examples.

Conjoined circles: we picked 400 points distributed along circles in the plane with radius 1 and with centres at $(\pm 1, 0)$. The points were then jittered by adding noise to each coordinate taken uniformly randomly from the interval $[0.0, 0.3]$. A Rips complex was constructed with maximal radius 0.5, resulting in 76763 simplices. The cohomology was computed in 378 seconds.

Disjoint circles: 400 points were distributed on circles of radius 1 centered around $(\pm 2, 0)$ in the plane. These points were subsequently disturbed by a uniform random variable from $[0.0, 0.5]$. We constructed a Rips complex with maximum radius 0.5, which gave us 45809 simplices. The cohomology computation finished in about 117 seconds.

In both cases, our method detects the two most natural circle-valued functions. The scatter plots appear very similar. In the conjoined case, there is some interference between the two circles, near their meeting point.

3.7 Torus

See Figure 5. We picked 400 points at random in the unit square, and then used a standard parametrization to map the points onto a torus with inner and outer radii 1.0 and 3.0. These were subsequently jittered by adding a uni-




Figure 1: Noisy circle. Persistence diagram (left). Global coordinate (top row), local coordinate (bottom row). In each row: histogram of coordinate values (left), correlation scatter plot against known angle function (middle), inferred coordinate in color (right).




Figure 2: Trefoil torus knot. Persistence diagram (left), correlation scatter plot of inferred coordinate against known parametrization (middle), inferred coordinate in color (right).




Figure 3: Images of a rotating cube. Histogram of coordinate values (left); scatter plot against known angle function (middle); a selection of images matched to recovered circle coordinate (right).




Figure 4: Two conjoined circles (left); two disjoint circles (right). In each case we show the persistence diagram (top left), the two inferred coordinates (right column), the correlation scatter plot (bottom left).

form random variable from $[0.0, 0.2]$ to each coordinate. We constructed a Rips complex with maximal radius $\sqrt{3}$, resulting in 61522 simplices. The corresponding cohomology was computed in 209 seconds.

The two inferred coordinates in this (fairly typical) experimental run recover the original coordinates essentially perfectly: the first inferred coordinate correlates with the meridional coordinate with topological degree -1 , while the second inferred coordinate correlates with the longitudinal coordinate with degree 1.

When the original coordinates are unavailable, the important figure is the inferred-versus-inferred scatter plot. In this case the scatter plot is fairly uniformly distributed over the entire coordinate square (i.e. torus). In other words, the two coordinates are decorrelated. This is slightly truer (and more clearly apparent in the scatter plot) for the two original coordinates. Contrast these with the corresponding scatter plots for a pair of circles (conjoined or disjoint).

3.8 Elliptic curve

See Figure 6. For fun, we repeated the previous experiment with a torus abstractly defined as the zero set of a homogeneous cubic polynomial in three variables, interpreted as a complex projective curve. We picked 400 points at random on $S^5 \subset \mathbb{C}^3$, subject to the cubic equation

$$x^2y + y^2z + z^2x = 0.$$

To interpret these as points in $\mathbb{CP}^2$, we used the projectively invariant metric

$$d(\xi, \eta) = \cos^{-1}(|\bar{\xi} \cdot \eta|)$$

for all pairs $\xi, \eta \in S^5$. With this metric we built a Rips complex with maximal radius 0.15. The resulting complex had 44184 simplices, and the cohomology was computed in 56 seconds. We found two dominant coclasses that survived beyond radius 0.15, and we computed our parametrizations at the 0.15 mark.

The resulting correlation plot quite clearly exhibits the decorrelation which is characteristic of the torus.






Figure 6: Elliptic curve. Persistence diagram (left), correlation scatter plot between the two coordinates (right).

3.9 Double torus


See Figure 7. We constructed a genus-2 surface by generating 1600 points on a torus with inner and outer radii 1.0 and 3.0; slicing off part of the data set by a plane at distance 3.7 from the axis of the torus, and reflecting the remaining points in that plane. The resulting data set has 3120 points. Out of these, we pick 400 landmark points, and construct a witness complex with maximal radius 0.6. The landmark set yields a covering radius $r_{\max} = 0.9982$ and a complex with 70605 simplices. The computation took 748 seconds active computer time. We identified the four most significant cocycles.

Note that coordinates 1 and 4 are ‘coupled’ in the sense that they are supported over the same subtorus of the double torus. The scatter plot shows that the two coordinates appear to be completely decorrelated except for a large mass concentrated at a single point. This mass corresponds to the other subtorus, on which coordinates 1 and 4 are essentially constant. A similar discussion holds for coordinates 2 and 3.

The uncoupled coordinate pairs (1,2), (1,3), (2,4), (3,4) produce scatter plots reminiscent of two conjoined or disjoint circles.




(a) Persistence diagram (left); first inferred coordinate (middle); second inferred coordinate (right).




(b) Correlation scatter plots between the two original and two inferred coordinates.

Figure 5: Torus in $\mathbb{R}^3$.



(a) Persistence diagram: four cocycles detected.



(b) Correlation scatter plots between the four inferred coordinates.

Figure 7: Double torus in $\mathbb{R}^3$.

4. ACKNOWLEDGEMENTS

We are immensely grateful to Dmitriy Morozov: he has given us considerable assistance in implementing the algorithms in this paper. In particular we thank him for the persistent cocycle algorithm.

Thanks also to Jennifer Klope for sharing her analysis of a visual image data set; this example did not make the present version of this paper. Finally, we thank Gunnar Carlsson, for his support and encouragement as leader of the topological data analysis research group at Stanford; and Robert Ghrist, as leader of the DARPA-funded project Sensor Topology and Minimal Planning (SToMP).

5. REFERENCES

- [1] M. Belkin and P. Niyogi. Laplacian Eigenmaps and spectral techniques for embedding and clustering. In T. Diettrich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 585–591. MIT Press, Cambridge, Massachusetts, 2002.
- [2] T. F. Cox and M. A. A. Cox. *Multidimensional Scaling*. Chapman & Hall, London, 1994.
- [3] V. de Silva and G. Carlsson. Topological estimation using witness complexes. In M. Alexa and S. Rusinkiewicz, editors, *Eurographics Symposium on Point-Based Graphics*, ETH, Zürich, Switzerland, 2004.
- [4] M. Dixon, N. Jacobs, and R. Pless. Finding minimal parameterizations of cylindrical image manifolds. In *CVPRW '06: Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop*, page 192, Washington, DC, USA, 2006. IEEE Computer Society.
- [5] D. L. Donoho and C. Grimes. Hessian Eigenmaps: New locally linear embedding techniques for high-dimensional data. Technical Report TR 2003-08, Department of Statistics, Stanford University, 2003.
- [6] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
- [7] R. Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.
- [8] L. J. Guibas and S. Y. Oudot. Reconstruction using witness complexes. In *Proc. 18th ACM-SIAM Sympos. on Discrete Algorithms*, pages 1076–1085, 2007.
- [9] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
- [10] J. A. Lee and M. Verleysen. Nonlinear dimensionality reduction of data manifolds with essential loops. *Neurocomputing*, 67:29–53, 2005.
- [11] C. C. Paige and M. A. Saunders. LSQR: Sparse equations and least squares. <http://www.stanford.edu/group/SOL/software/lqr.html>.
- [12] C. C. Paige and M. A. Saunders. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Transactions on Mathematical Software*, 8(1):43–71, March 1982.
- [13] R. Pless and I. Simon. Embedding images in non-flat spaces. In *Conference on Imaging Science Systems and Technology*, pages 182–188, 2002.
- [14] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, Dec. 2000.
- [15] H. Sexton and M. Vejdemo-Johansson. jPlex simplicial complex library. <http://comptop.stanford.edu/programs/jplex/>.
- [16] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, Dec. 2000.
- [17] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005.