

Persistence Barcodes for Shapes

Gunnar Carlsson[†], Afra Zomorodian^{‡§}, Anne Collins[†], and Leonidas Guibas[‡]

[†] Department of Mathematics

[‡] Department of Computer Science
Stanford University

Abstract

In this paper, we initiate a study of shape description and classification via the application of persistent homology to two tangential constructions on geometric objects. Our techniques combine the differentiating power of geometry with the classifying power of topology. The homology of our first construction, the tangent complex, can distinguish between topologically identical shapes with different “sharp” features, such as corners. To capture “soft” curvature-dependent features, we define a second complex, the filtered tangent complex, obtained by parametrizing a family of increasing subcomplexes of the tangent complex. Applying persistent homology, we obtain a shape descriptor, called a barcode, that is a finite union of intervals. We define a metric over the space of such intervals, arriving at a continuous invariant that reflects the geometric properties of shapes. We illustrate the power of our methods through a number of detailed studies of parametrized families of mathematical shapes.

1. Introduction

In this paper, we initiate a study of shape description using a marriage of geometric and topological techniques. For us, a shape is any closed subset of a Euclidean space. Questions we wish to answer include:

- Which shapes are identical up to a rigid motion of the ambient Euclidean space?
- Which shapes belong to certain classes of shapes? For instance, we may wish to distinguish the class of rectangular prisms from the class of tetrahedra, or the class of ellipsoids from the class of spheres.
- Which shapes are identical up to a smooth diffeomorphism of the ambient Euclidean space? This is a coarse classification.
- What are the *features* of the shape? For example, these features could be singular points of various types.

We attempt to obtain information about shapes through different levels of classification, from fine to coarse. One can

view geometry as the finest level of classification as it focuses on *local* properties of shapes. In this sense, geometry has a *quantitative* nature and can answer low level questions about a shape. But all our questions have a *qualitative* feel and take a higher view of a shape. This prompts us to look at topological techniques which classify shapes according to the way they are connected *globally* – their *connectivity*. Unfortunately, this view is often too coarse to be useful. For example, the topological invariant *homology* classifies shapes according to the number of components, tunnels, and enclosed spaces. This classification cannot distinguish between circles and ellipses, between circles and rectangles, or even between Euclidean spaces of different dimensions. Further, it cannot identify *sharp* features, such as corners, edges, or cone points: their neighborhoods are all *homeomorphic* or connected the same way. Whereas geometry is too sensitive for shape description, topology seems to be too insensitive. We seek to enrich topological techniques with the geometric content of a shape so that our methods manifest moderate sensitivity. We need a method that combines local and global information to characterize a shape.

1.1. Approach

In this paper, we propose a robust method that combines the differentiating power of geometry with the classifying power

[†] Research supported, in part, by NSF under grant DMS 0101364.

[‡] Research supported, in part, by NSF/DARPA under grant CARGO 0138456 and NSF under grant ITR 0086013.

[§] Portion done while author was visiting the Max-Planck-Institut für Informatik in Saarbrücken, Germany.

of topology. We apply homology not to a space X itself, but to derived spaces that have richer geometric content. In this case, we construct spaces out of X using tangential information about X as a subset of \mathbb{R}^n . We define two tangential constructions: the tangent complex and the filtered tangent complex. The *tangent complex* is the closure of the space of all tangents to all points of X . The homology of this complex can detect many sharp features of X , such as edges and corners. However, if the two shapes differ only in *soft features*, their tangent complexes are topologically indistinguishable. For example, an ellipse differs from a circle in having a range of curvatures, a feature not reflected in the ellipse's tangents.

For even more differentiating power, we enrich our tangent complex with information about the curvature of the space at each point. The resulting *filtered tangent complex* is actually an increasing family of spaces, parametrized by curvature. Applying homology to this family, we obtain an invariant called a *persistence module* that yields subtler information about a shape, such as the existence of regions of low or high curvature. The filtered tangent complex can distinguish a circle from an ellipse, even though these spaces are homeomorphic and have homeomorphic tangent complexes. We provide a number of examples to illustrate the power of our tangential constructions.

We also define a simple shape descriptor that arises from persistent homology and carries information about a shape. We call this combinatorial invariant a *barcode* as it is simply a finite collection of intervals on the right half line. To compare barcodes, we define a metric on the set of barcodes. We may now use clustering techniques on the metric space of barcodes for shape recognition and classification.

In this paper, we deal with explicit calculations for geometric objects to demonstrate the sensitivity of our invariants to properties of shapes, and their effectiveness in shape recognition. Our main interest, however, is computational. We wish to obtain information about a shape when we only have a finite set of samples from that shape. Such a set of points, commonly called *point cloud data* or *PCD* for short, is a discrete topological space. We are faced, therefore, with the additional difficulty of recovering the underlying shape topology, as well as approximating the tangential spaces we define. These issues extend well beyond the scope of this paper. Our goal here is to examine the viability of our techniques in an ideal mathematical setting, where we have full information about the spaces considered. We make no claims, however, about the direct applicability of the results of this paper to the PCD domain, which we have addressed for PCDs of curves in [CZCG04].

1.2. Background and Prior Work

Distinguishing and recognizing shapes is a well-studied problem in many areas of computer science, such as vision, pattern recognition, and graphics [Fan90, Fis89]. Me-

dial axis techniques have been extensively used for this purpose [ZY96]. One particular aspect of this problem is understanding “shape spaces” in their entirety [Boo91, KBCL99]. For instance, one may study a shape space from a differential point of view [LPM01]. Broadly speaking, most methods try to coordinatize the shape space in a way that gives useful combinatorial information about the space.

The idea of applying homology to derived spaces is a familiar one within topology. Some applications include the classification of vector bundles [MS74], and *Donaldson theory*, which constructs invariants of smooth structures on homotopy equivalent four-dimensional manifolds, using moduli spaces of self-dual connections [DK90]. The tangential complexes we study are familiar in *geometric measure theory* [Fed69]. *Morse theory* analyzes the topology of a space by utilizing increasing families of spaces [Mil63]. *Persistent homology* was first defined for topological simplification of alpha shapes [ELZ02]. It was then extended to arbitrary d -dimensional spaces and arbitrary fields of coefficients using a classification result [ZC04]. The latter paper also introduced the notion of using intervals as homological invariants.

1.3. Overview

We organize the rest of the paper as follows. We describe a topological construction that we employ in this paper in Section 2. We also list some tools we utilized in our calculations. In Section 3, we define our first derived complex, the tangent complex, based on the notion of the tangent cone from geometric measure theory. We also discuss a number of examples that point out the strengths and weaknesses of the complex in shape recognition. This discussion motivates our second derived complex in Section 4, the filtered tangent complex, which is more sensitive and can detect curvature-dependent features. Once again, we study a number of examples in two and three dimensions. Section 5 introduces persistent homology and the classification it gives in terms of barcodes. It also defines a metric on the space of barcodes and gives an algorithm for computing this metric based on weighted bipartite matching. In Section 6, we revisit our examples from Sections 3 and 4 and compute their homology or persistent homology, drawing the barcodes in the latter cases. We devote Section 7 to an extended example, where we illustrate the effect of varying parameters on the barcodes and the distance between shapes by studying the parametrized families of a glass and a bottle. We conclude the paper with some preliminary results from our software to compute tangent complexes for PCDs.

2. Algebraic Topology

We assume that the reader is familiar with major concepts from algebraic topology, such as homotopy, homotopy equivalence, homology, and Meyer-Vietoris sequences. For

these definition see any standard text in the area, such as Greenberg and Harper [GH81] or Hatcher [Hat01].

We mention a space-level construction we will be using later in this paper. Intuitively, when a space is broken into a union of two pieces, we can express it via the construction of gluing the two pieces together along a subspace.

Definition 1 ($A \cup_X B$) Suppose we have a space X , together with inclusions $f: X \hookrightarrow A$ and $g: X \hookrightarrow B$. Then, we may construct a new space as the quotient of the disjoint union of A and B ,

$$A \cup_X B = A \dot{\cup} X \dot{\cup} B / \sim,$$

where \sim is the equivalence relation generated by the equivalences $x \sim f(x) \in A$, and $x \sim g(x) \in B$, for all $x \in X$. Similarly, if we have inclusions $f: X \hookrightarrow A$, $g: X \hookrightarrow B$, $h: Y \hookrightarrow B$, and $k: Y \hookrightarrow C$, we define

$$A \cup_X B \cup_Y C = (A \cup_X B) \cup_Y C.$$

Although this construction depends on the inclusions f and g , they are not indicated in the notation.

Example 1 (glass) Figure 1 shows the construction of a glass shape from two pieces. The side is a cylinder $A = S^1 \times [0, 1]$ and the bottom is a disc $B = D^2$. We use the dashed circle $X = S^1$ to glue the two spaces together, resulting in $D^2 \cup_{S^1} (S^1 \times [0, 1])$, which looks like a glass or the bottom of a bottle.

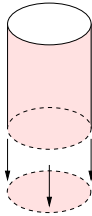


Figure 1: The cylinder $S^1 \times [0, 1]$ and the disc D^2 are glued together along the dashed circles. The resulting space $D^2 \cup_{S^1} (S^1 \times [0, 1])$ has the same geometry as a glass or bottom of a bottle.

3. The Tangent Complex

In this section, we introduce the first of our tangential constructions, the tangent complex for a subset of \mathbb{R}^n . Our construction is closely related to the concept of the *tangent cone*, defined in geometric measure theory [Fed69]. After defining the complex, we give a number of examples for intuition.

Definition 2 (tangent complex) Let X be any subset of \mathbb{R}^n . We define $T^0(X) \subseteq X \times S^{n-1}$ to be

$$T^0(X) = \left\{ (x, \zeta) \mid \lim_{t \rightarrow 0} \frac{d(x+t\zeta, X)}{t} = 0 \right\}.$$

The tangent complex of X is the closure of T^0 , $T(X) = \overline{T^0(X)} \subseteq X \times S^{n-1}$. $T(X)$ comes equipped with a projection $\pi: T(X) \rightarrow X$ that projects a point $(x, \zeta) \in T(X)$ in the tangent complex onto its basepoint x , and $T(X)_x = \pi^{-1}(x) \subseteq T(X)$ is the fiber at x .

Informally, (x, ζ) represents a tangent vector at a point $x \in X$, if the ray beginning at x and pointing in the direction prescribed by ζ approaches X faster than linearly. There is also a projection $\omega: T(X) \rightarrow S^{n-1}$ that projects all the fibers onto the $(n-1)$ -sphere. We have the following useful proposition concerning this construction.

Proposition 1 Suppose that $x \in X$ is a smooth point of X , so that there is a neighborhood $U \subseteq \mathbb{R}^n$ of x and a smooth function $f: U \rightarrow \mathbb{R}^m$, such that

- $U \cap X = f^{-1}(0)$, and
- the Jacobian of f , $Df(\xi)$, has rank m for every ξ in U .

Then $T(X)_x \cong S^{n-m-1}$

For intuition, we next give a number of examples. We will revisit these examples in Section 6 to discuss the homology of their tangent complexes. We begin with one-dimensional objects, as they are easier to work with, given the definition.

Example 2 (hyperplane) Let L be a line in the xy -plane, given by the equation $(x-x_0) \cdot \xi = 0$, for point x_0 and vector ξ . Then we have $\omega(T(L)) = \{\pm\eta\}$, where η is a unit vector perpendicular to ξ , and

$$T(L) \cong L \times \{\pm\eta\}.$$

We show the line and its tangent complex in Figure 2. The arrows on the tangent complex components give the directions of the tangents. The fiber at any point consists of two distinct points, or a *zero-dimensional sphere* S^0 . More generally, let $X \subseteq \mathbb{R}^n$ be the hyperplane determined by the equation $(x-x_0) \cdot \xi = 0$, where x_0 and ξ are n -vectors. Then $T(X) \cong X \times S(\xi)$, where $S(\xi)$ denotes the unit sphere in the plane of vectors perpendicular to ξ . This result holds with X replaced by any halfplane or quadrant in \mathbb{R}^n .

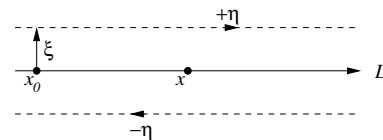


Figure 2: The solid line L is defined by point x_0 and normal ξ . It has a tangent complex that is composed of the two dashed components in the directions $\pm\eta$. The fiber at point x consists of two points.

Example 3 (vee) Consider a space $X \subseteq \mathbb{R}^2$ that looks like a ‘V’, as shown in Figure 3. Formally, $X = (\mathbb{R}_+ \times 0) \cup (0 \times \mathbb{R}_+)$. We evaluate the fibers $T(X)_x$ directly. The fiber at any smooth point is S^0 as in the previous example. For points along the x -axis, the two points will be $(x, (1, 0))$ and $(x, (-1, 0))$, and along the y -axis, they will be $(x, (0, 1))$ and $(x, (0, -1))$. At the origin O , however, the fiber $T(X)_O$ consists of four points, namely $(O, (\pm 1, 0))$ and $(O, (0, \pm 1))$. We can easily verify that the tangent complex is actually the union of two pieces, the tangent complexes of the two factors of the space:

$$T(\mathbb{R}_+ \times 0) = (\mathbb{R}_+ \times 0) \times \{(\pm 1, 0)\},$$

$$T(0 \times \mathbb{R}_+) = (0 \times \mathbb{R}_+) \times \{(0, \pm 1)\}.$$

So that the entire complex is:

$$T(X) = (\mathbb{R}_+ \times 0) \times \{(\pm 1, 0)\} \cup (0 \times \mathbb{R}_+) \times \{(0, \pm 1)\}.$$

It is easy to see that this space is a disjoint union of four distinct half lines, as shown in Figure 3.

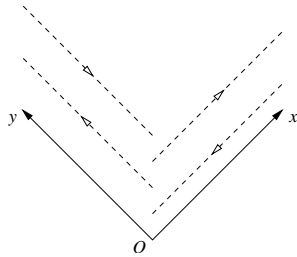


Figure 3: Solid ‘V’ space X has a tangent complex $T(X)$ that is composed of the four dashed components.

The features detected by homology of the tangent complex are *sharp features*: they are unaffected by smooth diffeomorphisms of the ambient space \mathbb{R}^n . However, we are often interested in *soft features* that are characterized by changes in curvature. We give two examples of the kinds of features we would like to capture.

Example 4 (circle vs. ellipse) Consider the difference between a circle and an ellipse. They are topologically the same and have no sharp features. There is, however, a difference in the lengths of the axes in the case of the ellipse that distinguishes it from the circle.

Example 5 (bottle vs. glass) Consider the difference between a bottle and a glass. Again, they are topologically the same and even have a common sharp feature: the circular rim at the bottom. But a bottle has a narrow neck that a glass does not have. Of course, bottles have many different shapes, but they all share this basic difference with a glass.

The distinctions above cannot be detected via homology directly, as the spaces are homeomorphic. They cannot be

detected by the tangent complex, either, as the objects are obtainable from each other by smooth deformations of the ambient space. Rather, they require more a sophisticated construction which we introduce in the next section.

4. The Filtered Tangent Complex

In this section, we define a second construction based on tangential information: the *filtered* tangent complex. We devote the bulk of this section to detailed examples, building the complex for interesting families of subspaces of the two- and three-dimensional Euclidean spaces. We will revisit these examples in Section 6, when we apply persistent homology to arrive at compact shape descriptors.

4.1. Definition

We begin by examining points in the tangent complex of a curve. Recall that a point (x, ζ) in $T^0(X)$ is composed of a basepoint $x \in X$ in the space and a tangent direction ζ , as shown for an example in Figure 4. For each direction, we compute the absolute value of the curvature, the inverse of the radius of the osculating circle shown in the figure. We formalize this concept next.

Definition 3 (second order contact) Let $X \subseteq \mathbb{R}^n$ and $(x, \zeta) \in T^0(X)$ (the tangent complex before the closure operation.) We say (x, ζ) has a circle of second order contact if there is a point $x_0 \in \mathbb{R}^n$, such that $(x - x_0) \cdot v = 0$ for all v for which $(x, v) \in T^0(X)$ and

$$\lim_{\theta \rightarrow 0} \frac{d(x_0 + \cos \theta \cdot (x - x_0) + \sin \theta \cdot |x - x_0| \cdot \zeta, X)}{\theta^2} = 0.$$

We define $\rho(x, \zeta) = |x - x_0|$ to be the radius of this circle.

Note that there exists at most one circle of second order contact at any point of $T(X)$, so the definition is well-defined. Here, we allow ρ to take the value ∞ for flat regions of a curve.

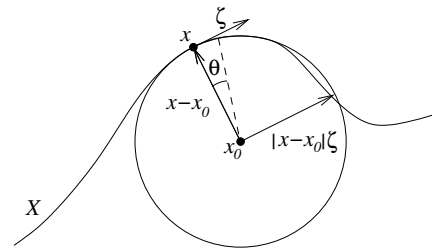


Figure 4: (x, ζ) is in $T^0(X)$, i.e. point x on curve X has unit tangent ζ . The osculating circle at x is centered at x_0 . The perpendicular vectors $x - x_0$ and $|x - x_0| \zeta$ and parameter θ provide coordinates for points on the circle.

Definition 4 (filtered) Let $T_\delta^0(X)$ to be the set of points $(x, \zeta) \in T^0(X)$ for which the circle of second order contact exists with $\delta \geq 1/\rho(x, \zeta)$. Let $T_\delta(X)$ be the closure of $T_\delta^0(X)$ in $\mathbb{R}^n \times S^{n-1}$. We call the δ -parametrized family of spaces $\{T_\delta(X)\}_{\delta \geq 0}$ the filtered tangent complex, denoted by $T^{filt}(X)$.

In other words, we allow tangent complex points to enter the complex according to the curvature at their associated basepoints. If $\delta \leq \delta'$, we have an inclusion $T_\delta(X) \subseteq T_{\delta'}(X)$, so the complexes are properly nested. As before, we denote $T_\delta(X)_x = \pi_\delta^{-1}(x)$, where $\pi_\delta: T_\delta(X) \rightarrow X$ is the natural projection. It is the topological structure of $T^{filt}(X)$ that carries information about the soft features of the underlying space.

4.2. Examples: Curves

We elucidate these concepts through examples in the rest of this section. We begin with one-dimensional subsets of \mathbb{R}^2 , or curves.

Example 6 (circle) Let X be the circle of radius R centered at the origin in the plane. The full tangent complex $T(X)$ is homeomorphic to $S^1 \times S^0$, as there are two tangent directions at each point. Every point $(x, \zeta) \in T^0(X)$ admits a circle of second order contact, namely X itself, so $\rho(x, \zeta) = R$ for all (x, ζ) . Therefore,

$$T_\delta(X) = \begin{cases} \emptyset, & \text{for } \delta < 1/R, \\ T(X), & \text{for } \delta \geq 1/R. \end{cases}$$

Example 7 (ellipse) Let X be the ellipse given by the equation $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$. Since X is a smooth manifold, every (x, ζ) admits a circle with second order contact. We recall that the formula for the radius of the osculating circle to a parametrized curve $\varphi(t) = (x(t), y(t))$ is given by

$$\frac{|\varphi'(t)|^3}{|x'(t)y''(t) - x''(t)y'(t)|}.$$

A parametrization for the ellipse X is given by $\varphi(t) = (a \cos(t), b \sin(t))$. Direct computation shows that at the point $\varphi(t)$, ρ is given by

$$\frac{(a^2 + (b^2 - a^2) \cos^2(t))^{\frac{3}{2}}}{ab},$$

for both tangent directions. Without loss of generality, we assume $b > a > 0$. Then, ρ attains its maximum value of b^2/a at $t \in \{0, \pi\}$, and its minimum value of a^2/b at $t \in \{\pi/2, 3\pi/2\}$.

Moreover, ρ is decreasing as t increases in intervals $(0, \frac{\pi}{2})$ and $(\pi, \frac{3\pi}{2})$, and increasing in intervals $(\frac{\pi}{2}, \pi)$ and $(\frac{3\pi}{2}, 2\pi)$. Therefore, the filtered tangent complex behaves as follows. For each case, there is a corresponding illustration in Figure 5.

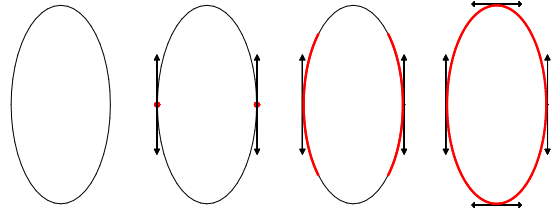


Figure 5: Ellipse $x^2 + y^2/4 = 1$. Points x whose fiber $T_\delta(X)_x$ has two components are indicated with thick lines. From left to right, we have $0 \leq \delta < \frac{a}{b^2}$, $\delta = \frac{a}{b^2}$, $\frac{a}{b^2} < \delta < \frac{b}{a^2}$, and $\delta \geq \frac{b}{a^2}$.

1. For $0 \leq \delta < \frac{a}{b^2}$, $T_\delta(X) = \emptyset$.
2. For $\delta = \frac{a}{b^2}$, $T_\delta(X)$ consists of four points, namely $(\pm a, 0) \times (0, \pm 1)$.
3. For $\frac{a}{b^2} < \delta < \frac{b}{a^2}$, $T_\delta(X)$ consists of four connected contractible components. These components have basepoints along two disjoint arcs on X , centered at the points $(\pm a, 0)$, and bounded away from the points $(0, \pm b)$.
4. For $\delta \geq \frac{b}{a^2}$, $T_\delta(X) = T(X)$, the full tangent complex, which is two copies of the ellipse, in clockwise and counterclockwise directions.

4.3. Examples: Surfaces

Before going on to examples involving surfaces, we review some aspects of the geometry of surfaces. Recall that on a curve, there was at most a single osculating circle. On a smooth surface, however, the tangent space at each point is a plane, giving us a whole circle of tangent directions. To construct the tangent complex, we need to determine the radius of the osculating circle in each direction.

Let $x \in X$ be a point in a smoothly embedded space X in Euclidean space. The best linear approximation to X at the point x , the *tangent plane*, determines a normal direction to the surface. Choosing one of the two unit vectors in the normal direction as axis z , we obtain a coordinate system for \mathbb{R}^3 centered at x with the surface tangent to the “ xy -plane”, as shown in Figure 6.

With respect to this coordinate system, the best approximating quadratic surface to X is given by an equation of the form $z = Q(x, y) = Ax^2 + Bxy + Cy^2$. When restricted to the unit circle on the tangent plane, Q attains its minimum and maximum values at pairs of antipodal points which determine a pair of perpendicular directions. These critical values are called the *principal curvatures* and their corresponding directions the *principal directions*. For an arbitrary direction ζ in the tangent plane, $2|Q(\zeta)| = 1/\rho(x, \zeta)$, the reciprocal of the radius $\rho(x, \zeta)$ of the osculating circle to the surface X in the direction ζ . Therefore, for a point $x \in X$, $T_\delta(X)_x$ consists

of points $(x, \zeta) \in T(X)_x$ where $2|Q(\zeta)| \leq \delta$. Let κ_1, κ_2 denote the minimum and maximum values of $2|Q|$ on $T(X)_x$. The following cases describe $T_\delta(X)_x$.

1. If $\delta < \kappa_1$, then $T_\delta(X)_x = \emptyset$.
2. If $\delta = \kappa_1 < \kappa_2$, then $T_\delta(X)_x$ consists of two antipodal points.
3. If $\kappa_1 < \delta < \kappa_2$, $T_\delta(X)_x$ consists of two disjoint closed antipodal arcs.
4. If $\delta \geq \kappa_2$, then $T_\delta(X)_x$ is the entire circle $T(X)_x$.

The fibers described here are similar to the basepoint projections of Example 7, illustrated in Figure 5, if we replace the ellipse with a unit circle. Given our analysis, we may now examine some surfaces in \mathbb{R}^3 .

Example 8 (surface of revolution) Consider the graph of any function $y = f(x)$, with $f(x) > 0$, and the surface of revolution X obtained by revolving the graph around the x -axis. Assuming that f is smooth, standard differential geometry tells us that at any point, the directions parallel and perpendicular to the circle of revolution are the principal directions [BG92]. Let $v = (x, y, z)$ be a point in X . We define $\rho^{rev}(v), \rho^{graph}(v)$ to be the radii of the circles of second order contact in the parallel and perpendicular directions, respectively, and $\kappa^{rev}(v) = 1/\rho^{rev}(v)$ and $\kappa^{graph}(v) = 1/\rho^{graph}(v)$ be the corresponding curvatures. It is easy to see that $\rho^{rev}(v)$ is the distance along the direction normal to the graph from the point $(x, f(x))$ to the axis of revolution, and $\rho^{graph}(v)$ is the radius of the osculating circle to the graph at $(x, f(x))$.

It is immediate that $T(X) \cong X \times S^1$. The projection of $T_\delta(X)$ is itself a surface of revolution, of the set of points that have directions with low enough curvature: $\{(x, f(x)) \mid \kappa^{rev} \leq \delta \text{ or } \kappa^{graph} \leq \delta\}$. A non-empty fiber at v consists of two antipodal points or two disjoint antipodal arcs if $(\kappa^{rev} \leq \delta < \kappa^{graph})$ or $(\kappa^{graph} \leq \delta < \kappa^{rev})$, and consists of a full circle, otherwise.

5. Persistent Homology and the Barcode Invariant

Our notion of a filtered tangent complex gives us a family of nested complexes, each with its own homology. The global topological properties of this family of spaces carries

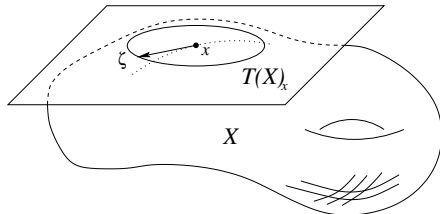


Figure 6: Surface X with the tangent plane at x and the unit tangent circle $T(X)_x$. We also show a dotted portion of the osculating circle in the direction ζ .

information about the shape of the object. As we showed in the last section, this family can often distinguish between topologically identical objects with different geometries. We will need invariants of this topological structure which can be readily computed. Evaluating homology on each of the spaces $T_\delta(X)$, we obtain homology groups $H_n(T_\delta(X))$ for each $\delta \geq 0$. While it is possible to glean information about the space from these groups, it would be computationally infeasible to do so. Furthermore, we would lose important information about the family, namely the homomorphisms $H_n(T_\delta(X)) \rightarrow H_n(T_{\delta'}(X))$ whenever $\delta \leq \delta'$, induced by the inclusion $T_\delta(X) \hookrightarrow T_{\delta'}(X)$.

We summarize the topological information about the filtered tangent complex in the notion of a *persistence module*, a directed system of Abelian groups parametrized on the ordered set $[0, \infty)$ [ZC04]. Persistence modules capture the information contained in the homomorphisms, are computable in the time required for computing a single homology group, and are classifiable in terms of a compact combinatorial object called a *barcode*. In this section, we introduce these algebraic and combinatorial notions for application to our geometric situation. We begin by defining the persistence modules and their classification as barcodes. We then introduce a metric on the space of barcodes. We end the section by giving an algorithm for computing this metric using maximum weighted bipartite matching.

5.1. Persistent Homology

In this section, we define the persistence module and show how filtered tangent complexes give rise to such modules naturally.

Definition 5 (persistence module) Let \mathcal{T} denote any totally ordered set, and let R be a ring. A R -persistence module parametrized by \mathcal{T} is a family of R -modules $\{M_t\}_{t \in \mathcal{T}}$ together with homomorphisms of R -modules $\varphi_{t,t'} : M_t \rightarrow M_{t'}$ for all $t \leq t'$, such that the homomorphisms are compatible,

$$\varphi_{t,t'} \cdot \varphi_{t',t''} = \varphi_{t,t''}, \tag{1}$$

whenever $t \leq t' \leq t''$

Example 9 (\mathbb{N}) Let $\mathcal{T} = \mathbb{N} = \{0, 1, 2, \dots\}$. Then an R -persistence module parametrized by \mathcal{T} is simply a family of modules $\{M_n\}_{n \geq 0}$, with homomorphisms

$$M_0 \rightarrow M_1 \rightarrow \dots \rightarrow M_{n-1} \rightarrow M_n \rightarrow M_{n+1} \rightarrow \dots$$

Such an object is often called an *inductive system* of R -modules. A straightforward variant is one where $\mathcal{T} = \mathbb{Z}$, the set of all integers, giving rise to a double-ended system of the type described above.

Example 10 (filtered tangent complex) Suppose that we have a family of spaces X_δ , parametrized by the real valued parameter δ , so that $X_\delta \subseteq X_{\delta'}$ whenever $\delta \leq \delta'$. Then the

family of Abelian groups $H_n(X_\delta, G)$, where G is any Abelian group, is a G -persistence module parametrized by \mathbb{R} . Note that if G is any field, then we obtain a G -persistence vector space parametrized by \mathbb{R} over the field G .

We also speak of a *persistence chain complex parametrized by \mathcal{T}* over a ring R , a family of chain complexes $\{C_*^t\}_{t \in \mathcal{T}}$ of G -modules together with chain maps $C_*^t \rightarrow C_*^{t'}$ whenever $t \leq t'$, satisfying the analogues of the compatibility relations (1) above. The homology of a persistence complex parametrized by \mathcal{T} over G is always a G -persistence module, parametrized by \mathcal{T} . There are analogous notions of homomorphism and isomorphism for persistence modules. We are mainly concerned with the classification, up to isomorphism, of persistence modules over fields.

5.2. Classification over Fields

We first consider persistence modules $\{V_t\}_{t \in \mathbb{N}}$ parametrized by \mathbb{N} over a field F . We are able to classify these modules, provided they are finite. Precisely, a persistence module $\{V_t\}_t$ parametrized by \mathbb{N} over F is of *finite type* if

1. each vector space V_t is finite dimensional, and
2. there is an integer N so that for all $t \geq N$, $\varphi_{N,t}$ are isomorphisms.

Given any pairs of integers (m, n) with $m \leq n$, we define the persistence module $Q(m, n)$ over \mathbb{N} by

$$Q(m, n)_t = \begin{cases} 0, & \text{if } t < m \text{ or } t \geq n, \\ F, & \text{otherwise,} \end{cases}$$

The homomorphisms φ_{ij} are the identity homomorphisms for $m \leq i \leq j \leq n$. We can trivially extend this definition to the case $n = +\infty$. The main result of [ZC04] is the following.

Proposition 2 (classification) *A persistence module parametrized by \mathbb{N} over F of finite type is isomorphic to one of the form*

$$\bigoplus_{s=1}^n Q(i_s, j_s),$$

where j_s can be $+\infty$, and the decomposition is unique up to the order of the pairs.

This result is a consequence of the fundamental theorem for finitely generated modules over the graded principal ideal domains (PID) $F[t]$, with t having degree 1. Using the Artin-Rees construction, we can show that persistence modules parametrized by \mathbb{N} over any ring R are characterized by an associated graded module (non-negatively graded) over $R[t]$, and the result follows by taking $R = F$. We note this simple classification does not extend to non-fields, as the graded $R[t]$ will not be a PID.

We can achieve similar results for other parameter spaces using suitable finiteness conditions. For example, if $\mathcal{T} = \mathbb{R}$, we say a persistence module $\{V_s\}_s$ over F is of *finite type* if there are a finite number of unique finite-dimensional vector spaces in the persistence module. Let I be an interval. We define a persistence vector space $Q(I)$ over F

$$Q(I)_s = \begin{cases} 0, & \text{if } s \notin I, \\ F, & \text{otherwise,} \end{cases}$$

where the homomorphism is the identity within each interval. We can now state an analog of Proposition 2.

Proposition 3 *A persistence module parametrized by \mathbb{R} over F of finite type is isomorphic to one of the form*

$$\bigoplus_{s=1}^n Q(I_s),$$

where each interval I_s is bounded from below, and the description is unique up to the order of the intervals.

5.3. Metric Space of Barcodes

In the last section, we saw that we can classify persistence modules parametrized by \mathbb{N} with using a number of pairs of integers $\{(i_s, j_s)\}_s$. If we view these pairs as half-open intervals $\{[i_s, j_s)\}_s$, the description matches that of modules parametrized by \mathbb{R} . This family of intervals is our shape descriptor.

Definition 6 (barcode) *A barcode is a finite set of intervals that are bounded below.*

Intuitively, the intervals denote the life-times of a non-trivial loop in a growing complex. The left endpoint signifies the birth of a new topological attribute, and the right endpoint signals its death. The longer the interval, the more important the topological attribute, as it insists on being a feature of the complex.

We next wish to form a metric space over the collection of all barcodes. We do so using a quasi-metric, a metric that has ∞ as a possible value, where $x + \infty = \infty$ and $\infty + \infty = \infty$. Let \mathcal{I} denote the collection of all possible barcodes. We wish to define a quasi-metric $\mathcal{D}(S_1, S_2)$ on all pairs of barcodes (S_1, S_2) , with $S_1, S_2 \in \mathcal{I}$, so that if we move the endpoint of any single interval in either set by a dissimilarity ϵ , $\mathcal{D}(S_1, S_2)$ changes by no more than ϵ . Let I, J be any two intervals in a barcode. We define their dissimilarity $\delta(I, J)$ to be their symmetric difference: $\delta(I, J) = \mu(I \cup J - I \cap J)$, where μ denotes one-dimensional measure. Note that $\delta(I, J)$ may be infinite. Given a pair of barcodes S_1 and S_2 , a *matching* is a set $M(S_1, S_2) \subseteq S_1 \times S_2 = \{(I, J) \mid I \in S_1 \text{ and } J \in S_2\}$, so that any interval in S_1 or S_2 occurs in at most one pair (I, J) . Let M_1, M_2 be the intervals from S_1, S_2 , respectively, that are matched in M , and let N be the non-matched intervals $N = (S_1 - M_1) \cup (S_2 - M_2)$. Given a matching M for S_1 and

S_2 , we define the *distance of S_1 and S_2 relative to M* to be the sum

$$\mathcal{D}_M(S_1, S_2) = \sum_{(I, J) \in M} \delta(I, J) + \sum_{L \in N} \mu(L).$$

We now look for the best possible matching to define the quasi-metric.

Definition 7 (metric) $\mathcal{D}(S_1, S_2) = \min_M \mathcal{D}_M(S_1, S_2)$.

For a pair (I, J) , we know that $\delta(I, J) = \mu(I) + \mu(J) - 2\mu(I \cap J)$. Now, for any matching M , we define the *similarity of S_1 and S_2 with respect to M* to be

$$\begin{aligned} S_M(S_1, S_2) &= \sum_{(I, J) \in M} \mu(I \cap J) \\ &= \frac{1}{2} \left(\sum_{S_1} \mu(I) + \sum_{S_2} \mu(J) - \mathcal{D}_M(I, J) \right). \end{aligned}$$

Minimizing \mathcal{D}_M is equivalent to maximizing S_M . We may do the latter by recasting the problem as a graph problem. Given sets S_1 and S_2 , we define $G(V, E)$ to be a weighted bipartite graph [CLRS01]. We place a vertex in V for each interval in $S_1 \cup S_2$. We place an edge in E for each pair $(I, J) \in S_1 \times S_2$ with weight $\mu(I \cap J)$. Maximizing S_M is equivalent to the well-known *maximum weight bipartite matching problem*. We may utilize one of several algorithms for computing the matching, such as the *Hungarian* algorithm with time complexity $O(|V||E|)$ [Kuh55], or the *Gabow-Tarjan* algorithm with complexity $O(\sqrt{|V||E|} \log |E|)$ [GT89].

6. Examples, Revisited

Having described persistent homology, in this section we revisit all of our examples from sections 3 and 4. We give precise descriptions of the homology groups, and in the case of the filtered complexes, we provide pictures of the barcodes that describe the persistence modules.

Example 11 (hyperplane) For the hyperplane X in Example 2, we found that the tangent complex has the form $X \times S^{n-1}$. Since X is contractible, we may use the Künneth formula to conclude that $H_*(T(X)) \cong H_*(S^{n-1})$. We list $H_i(T(X))$ in the following table, where the dimension is listed horizontally, and i maps to dimensions not accounted for.

$$\begin{array}{c|c|c} 0 & i & n-1 \\ \hline \mathbb{Z} & 0 & \mathbb{Z} \end{array}$$

The Betti numbers are the dimensions of these vector spaces, so $\beta_0(T(X)) = \beta_{n-1}(T(X)) = 1$ and $\beta_i(T(X)) = 0$, for $i \neq 0, n-1$. In the remaining examples, we do not list the Betti numbers explicitly, as they can be easily read off.

Example 12 (vee) In Example 3, we found that $T(X)$ for the vee-shaped space is a disjoint union of four rays, shown

in Figure 3. As each ray is contractible, $H_*(T(X))$ is isomorphic to the homology of four points. Explicitly,

$$\begin{array}{c|c} 0 & i \\ \hline \mathbb{Z}^4 & 0 \end{array}$$

We next analyze some cases of the filtered tangent complex. Since our description of persistent homology in terms of barcodes only works for homology with field coefficients, we will work only with homology with $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$ coefficients.

Example 13 (circle) In Example 6, we computed $T^{\text{filt}}(X)$ for a circle of radius R in the plane. We found that the entire tangent complex appeared when $\delta = 1/R$. Below, we list the homology groups, where we now use rows to delineate the range for δ . To reduce clutter, we omit the lower bound of δ in each row, as it is implicitly defined by a higher row or is zero. We draw the corresponding barcodes in Figure 7.

$$\begin{array}{c|c|c} \delta & 0 & 1 \\ \hline < 1/R & 0 & 0 \\ \geq 1/R & \mathbb{F}^2 & \mathbb{F}^2 \end{array}$$

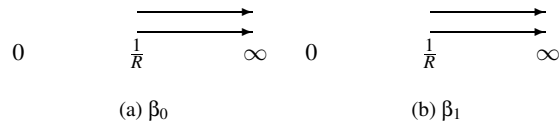


Figure 7: Barcodes for the circle.

Example 14 (ellipse) In Example 7, we computed $T^{\text{filt}}(X)$ for the ellipse $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$. While the full tangent complex $T(X)$ matches that of the circle in the previous example, we saw that $T_\delta(X)$ for the ellipse evolved through three stages, as shown in Figure 5. We draw the corresponding barcodes in Figure 8.

$$\begin{array}{c|c|c} \delta & 0 & 1 \\ \hline < a/b^2 & 0 & 0 \\ < b/a^2 & \mathbb{F}^4 & 0 \\ \geq b/a^2 & \mathbb{F}^2 & \mathbb{F}^2 \end{array}$$

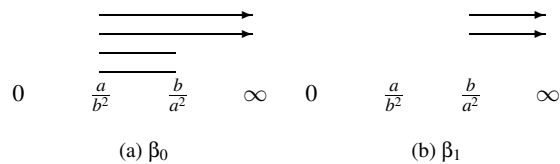


Figure 8: Barcodes for the ellipse.

7. Shape Recognition with Homology

In this section, we show how the homology of our tangential constructions is an effective tool for shape recognition and classification. In particular, we see how the method can classify objects in a continuous family of objects that have been modified by smooth deformations in the ambient space.

Example 15 (bottle vs. glass) In Example 5, we asked for a technique for distinguishing a bottle from a glass. The surfaces are homeomorphic and share sharp features. Therefore, we must look at the filtered tangent complex for distinguishing features. In this example, we compute the barcode invariant for both objects and show how we may compare them using our metric

Both a bottle and a glass may be described as surfaces of revolution, as shown in Figure 9. In both cases, we add a vertical line to the graph of a function. In order to study the filtered tangent complex, we need to make some assumptions about the shape of these objects. We begin with the bottle. Let $f: [0, H] \rightarrow \mathbb{R}$ be a twice-differentiable positive function whose graph is used to sweep out a bottle of height H . We make the following assumptions about f and the principle curvatures κ^{rev} and κ^{graph} (cf. Example 8):

1. f is constant on the closed intervals $[0, \xi_0]$ and $[\xi_H, H]$, with $f(0) > f(H)$, and is monotonically decreasing from ξ_0 to ξ_H , with a single inflection point at ξ .
2. κ^{rev} is monotonically increasing.
3. κ^{graph} is 0 at the inflection point and the intervals where f is constant, and has exactly two local maxima at $\xi_- \in (\xi_0, \xi)$ and $\xi_+ \in (\xi, \xi_H)$.

Let $\kappa_0 = \kappa^{rev}(0) = 1/f(0)$ and $\kappa_H = \kappa^{rev}(H) = 1/f(H)$ denote the inverse of the cross-sectional radii at $0, H$, respectively. Clearly, $\kappa_0 < \kappa_H$. Let $\kappa_- = \kappa^{graph}(\xi_-)$ and $\kappa_+ = \kappa^{graph}(\xi_+)$ denote the inverse of the radius of the osculating circle to the curve at ξ_-, ξ_+ , respectively. There are a number of different cases for analyzing $T_\delta(X)$. We will deal with the case $0 < \kappa_+ < \kappa_- < \kappa_0 < \kappa_H$, corresponding to a rather long, slowly tapering bottle, much like a Riesling wine bottle. Other cases may be treated similarly. We obtain the following results about $T_\delta(X)$.

- $\delta = 0$: Any point where the curvature κ^{graph} vanishes has non-empty fiber in $T_0(X)$. The projection of $T_0(X)$ decom-

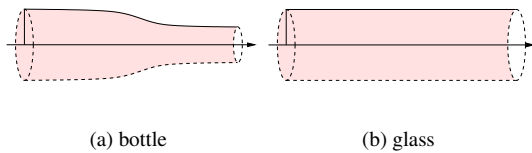


Figure 9: Surfaces of revolution around the x-axis.

poses into three components: W_1, W_2 and W_3 , corresponding to bottle’s body, a circle at f ’s inflection point, and the bottle’s neck, as shown in Figure 10(a). More precisely, W_1 is the surface of revolution generated by rotating the vertical line segment $0 \times [0, f(0)]$ together with the horizontal line segment $[0, \xi_0] \times f(0)$ around the x-axis. W_2 is generated by rotating the point $(\xi, f(\xi))$. W_3 is generated by rotating the segment $[\xi_H, H] \times f(H)$. W_2 is a circle, and W_3 is a cylinder (or annulus, when viewed from the top), which is homotopy equivalent to a circle. The fiber in $T_0(X)$ at each point of W_2 or W_3 is a pair of points, the zero-dimensional sphere S^0 , and we have

$$\pi^{-1}(W_2) \cong \pi^{-1}(W_3) \cong S^1 \times S^0.$$

W_1 is precisely the glass we constructed in Example 1. So, $\pi^{-1}(W_1)$ is the union of two pieces U and V , where U is the 0-tangent complex of the base disc, and V is the 0-tangent complex of the side. Since the base disc D^2 is flat, each fiber in U is a full circle, and we know that $U \cong D^2 \times S^1$ by Example 2. The side is a cylinder once again, so its 0-tangent complex is $V \cong S^1 \times S^0$ as for W_2 and W_3 above. Although the base disc and side intersect in a circle, the tangent directions from each piece do not intersect, so $\pi^{-1}(W_1)$ is the disjoint union of U and V :

$$\pi^{-1}(W_1) \cong (D^2 \times S^1) \dot{\cup} (S^1 \times S^0).$$

Putting it all together, $T_0(X)$ has the homotopy type

$$(D^2 \times S^1) \dot{\cup} (S^1 \times S^0) \dot{\cup} (S^1 \times S^0) \dot{\cup} (S^1 \times S^0).$$

- $0 < \delta < \kappa_+$: $T_\delta(X)$ does not change its homotopy type, while the inflection circle grows into an annulus, as shown in Figure 10(b).
- $\delta = \kappa_+$: The two components $\pi^{-1}(W_2)$ and $\pi^{-1}(W_3)$ become connected. So, the image of T_{κ_+} under π becomes the union of two pieces that correspond to the

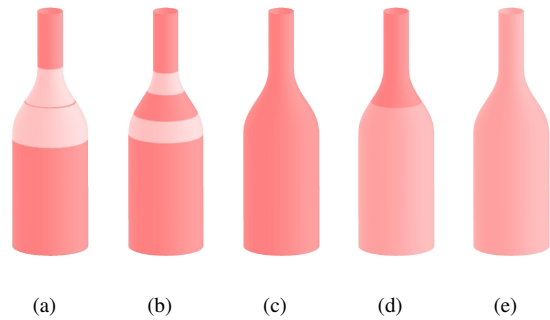


Figure 10: Bottles shaded by number of components in fiber. Light gray denotes an empty fiber, medium gray is the full circle, and dark gray is two arcs. In this example, $\kappa_- = \kappa_+ < \kappa_0$. We have the following δ values: (a) $\delta = 0$, (b) $0 < \delta < \kappa_+$, (c) $\kappa_- \leq \delta < \kappa_0$, (d) $\kappa_0 < \delta < \kappa_H$, and (e) $\delta \geq \kappa_H$.

body and the neck. The fiber of any point away from the base disc is a pair of points or disjoint antipodal intervals. Along the crease of the base, the fiber is a disjoint union of a circle and a pair of disjoint antipodal intervals. The complex description loses a term to become $T_{\kappa_+}(X) \cong (D^2 \times S^1) \dot{\cup} (S^1 \times S^0) \dot{\cup} (S^1 \times S^0)$.

- $\kappa_+ < \delta < \kappa_-$: The homotopy type does not change.
- $\delta = \kappa_-$: The neck and the body in the projection of $T_\delta(X)$ merge, as shown in Figure 10(c). Every fiber is now non-empty, and consists of two antipodal arcs or points. The homotopy type of $T_\delta(X)$ loses another term to become $(D^2 \times S^1) \dot{\cup} (S^1 \times S^0)$.
- $\kappa_- < \delta < \kappa_0$: The homotopy type does not change.
- $\delta = \kappa_0$: Points along the lower sides of the bottle now have full fiber. The individual fibers along the crease become a union of two circles, which now intersect in a pair of antipodal points. We now have

$$T_{\kappa_0}(X) \cong (D^2 \times S^1) \cup_{S^1 \times S^0} (S^1 \times S^1).$$

where $\cup_{S^1 \times S^0}$ is the construction in Definition 1.

- $\delta \geq \kappa_0$: The homotopy type remains unchanged, as shown in Figures 10(d) and 10(e).

Our analysis gives the following description of homology for $T_\delta(X)$.

δ	0	1	2
$< \kappa_+$	\mathbb{F}^7	\mathbb{F}^7	0
$< \kappa_-$	\mathbb{F}^5	\mathbb{F}^5	0
$< \kappa_0$	\mathbb{F}^3	\mathbb{F}^3	0
$\geq \kappa_0$	\mathbb{F}	\mathbb{F}^3	\mathbb{F}^2

We next perform the same calculation for a glass shown in Figure 9(b). We assume it is the surface of revolution $(0 \times [0, R]) \cup ([0, H] \times R)$ of radius R and height H . Observe that the glass is geometrically similar to the body of the bottle, the section called W_1 when $\delta = 0$ (or the surface in Example 1.) Let $\kappa = 1/R$. The analysis follows quickly from our observation:

- $\delta = 0$: We have $T_0(X) \cong (D^2 \times S^1) \dot{\cup} (S^1 \times S^0)$ as in the analysis of $\pi^{-1}(W_1)$ for the bottle.
- $0 < \delta < \kappa$: The homotopy does not change.
- $\delta = \kappa$: We have $T_\kappa(X) \cong (D^2 \times S^1) \cup_{S^1 \times S^0} (S^1 \times S^1)$ as in the analysis for the bottle when $\delta = \kappa_0$.
- $\delta > \kappa$: The homotopy does not change.

Our analysis gives the following description of homology for $T_\delta(X)$.

δ	0	1	2
$< \kappa$	\mathbb{F}^3	\mathbb{F}^3	0
$\geq \kappa$	\mathbb{F}	\mathbb{F}^3	\mathbb{F}^2

Note that the table for the glass is simply the bottom two rows of the table for the bottle. The bottle's neck is reflected in the first two rows, as well as shorter intervals, as shown in Figures 11 on the left. In the figure, we set $\kappa_0 = \kappa$ to compare a glass and bottle of the same curvature at the bottom.

Finally, we compute distances, in turn, between different bottles, different glasses, as well as bottles and glasses.

- Suppose we are given two bottles of the type analyzed here (i.e. $0 < \kappa_+ < \kappa_- < \kappa_0 < \kappa_H$) with parameters $(\kappa_+, \kappa_-, \kappa_0)$ and $(\kappa'_+, \kappa'_-, \kappa'_0)$. The distance between the β_0 -barcodes associated to these two bottles is

$$2(|\kappa_+ - \kappa'_+| + |\kappa_- - \kappa'_-| + |\kappa_0 - \kappa'_0|).$$

As expected, small changes in the parameters produce barcodes that are near each other in barcode space.

- Suppose we are given two glasses with curvatures κ, κ' , respectively. If we assume $|\kappa - \kappa'|$ is small, then the distance between the associated β_0 -barcodes, $2|\kappa - \kappa'|$, is also small.
- Suppose we have a bottle of the type analyzed with parameter vector $(\kappa_+, \kappa_-, \kappa_0)$ and a glass with curvature κ . There are three cases:

1. $0 < 2\kappa \leq \kappa_+ + \kappa_-$
2. $\kappa_+ + \kappa_- \leq 2\kappa \leq \kappa_- + \kappa_0$
3. $2\kappa \geq \kappa_- + \kappa_0$.

The corresponding distances are:

1. $2|\kappa_+ - \kappa| + 2\kappa_- + 2\kappa_0$
2. $2\kappa_+ + 2|\kappa_- - \kappa| + 2\kappa_0$
3. $2\kappa_+ + 2\kappa_- + 2|\kappa_0 - \kappa|$.

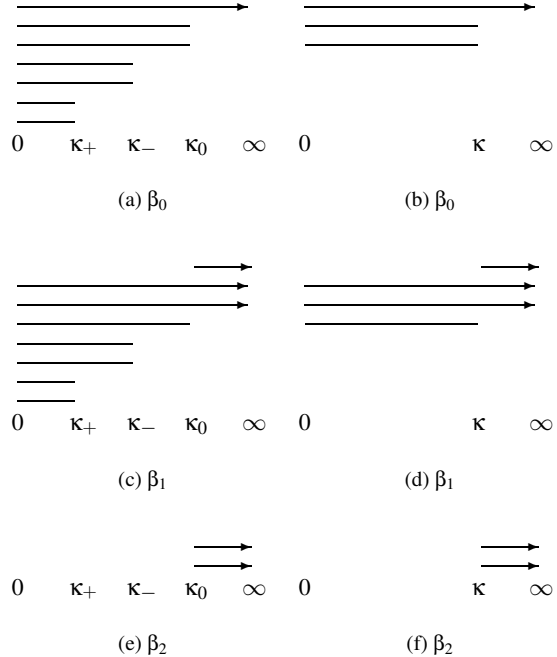


Figure 11: Barcodes for the bottle (left) and the glass (right.)

In the case of the bottle and the glass, the distances have dimensions and are not relative, so we can reasonably expect a good performance from our distance function in distinguishing between the two. We can also visualize how changes in the parameter values for the bottle will make it similar to a glass. As κ_+ , κ_- tend to 0, we get longer and flatter bottles, arriving at a glass in the limit. We may observe the same effect on the barcode, where the four shorter intervals become shorter as these parameters shrink, decreasing the bottle's distance to a glass.

8. Conclusion

In this paper, we propose a method for combining geometry and topology for shape recognition and classification. Our method applies persistent homology to tangential constructions to arrive at a simple and compact shape descriptor, the *barcode*, that captures both sharp and curvature-dependent features of a shape. We also define a metric over the space of barcodes, and provide an algorithm for computing this metric. Through detailed examples, we illustrate the viability of our method for shape recognition. In particular, we show that we may think of our homologically invariant barcodes as coordinates for shape spaces, as they lie in a metric space. These coordinates correspond to geometric attributes of shapes, such as the eccentricity of an ellipse or the curvature of the neck of a bottle. They enable us to distinguish between interesting categories of shapes, such as between the categories of bottles and glasses.

Our work suggests major avenues for future research:

- Application of our techniques to point cloud data (PCD). There are a number of interesting computational questions to be resolved, such as algorithms for constructing our derived complexes for the computation of barcodes. We have done so for PCDs of curves [CZCG04]. We plan to apply our techniques to PCDs of surfaces in the near future.
- Application of our techniques to noisy data. Our current experiments show that our techniques work well on idealized data, as well as data with superimposed Gaussian noise.
- Development of statistical methods for analyzing data in the metric space of barcodes. Such methods should permit the use of discriminant analysis and Bayesian decision procedures to develop automatic methods for shape classification.
- Utilization of barcodes for coordinatizing spaces, for example, gray-scale images of families of three-dimensional objects.
- Utilization of linear algebraic algorithms recently developed to locate features within data sets [CdS03].

We believe that the primary contribution of this paper is an approach of blending geometric and topological techniques that is grounded in theory. As high throughput scanning becomes commonplace, and more and more shapes are digitized, automatic qualitative shape analysis and classification

will be of critical value. Such analysis provides useful priors for shapes that may be exploited for operations on them, such as reconstruction. We believe that the type of research initiated here will be important in the geometry processing field in the years to come.

References

- [BG92] BERGER M., GASTIAUX B.: *Géométrie différentielle: variétés, courbes et surfaces*. Mathématiques Presses Universitaires de France, Paris, France, 1992.
- [Boo91] BOOKSTEIN F.: *Morphometric Tools for Landmark Data*. Cambridge Univ. Press, Cambridge, UK, 1991.
- [CdS03] CARLSSON G., DE SILVA V.: A geometric framework for sparse matrix problems. *Advances in Applied Mathematics* (2003). (To appear).
- [CLRS01] CORMEN T. H., LEISERSON C. E., RIVEST R. L., STEIN C.: *Introduction to Algorithms*. The MIT Press, Cambridge, MA, 2001.
- [CZCG04] COLLINS A., ZOMORODIAN A., CARLSSON G., GUIBAS L.: A barcode shape descriptor for curve point cloud data. *Computers and Graphics* (2004). (To appear).
- [DK90] DONALDSON S. K., KRONHEIMER P. B.: *The Geometry of Four-Manifolds*. The Clarendon Press, New York, NY, 1990.
- [ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete Comput. Geom.* 28 (2002), 511–533.
- [Fan90] FAN T.-J.: *Describing and Recognizing 3D Objects Using Surface Properties*. Springer-Verlag, New York, NY, 1990.
- [Fed69] FEDERER H.: *Geometric Measure Theory*, vol. 153 of *Die Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, New York, NY, 1969.
- [Fis89] FISHER R. B.: *From Surfaces to Objects: Computer Vision and Three-Dimensional Scene Analysis*. John Wiley and Sons, Inc., New York, NY, 1989.
- [GH81] GREENBERG M. J., HARPER J. R.: *Algebraic Topology: A First Course*, vol. 58 of *Mathematics Lecture Note Series*. Benjamin/Cummings Publishing Co., Reading, MA, 1981.
- [GT89] GABOW H. N., TARJAN R. E.: Faster scaling algorithm for network problems. *SIAM J. Comput.* 18 (1989), 1013–1036.

- [Hat01] HATCHER A.: *Algebraic Topology*. Cambridge Univ. Press, Cambridge, UK, 2001.
- [KBCL99] KENDALL D., BARDEN D., CARNE T., LE H.: *Shape and Shape Theory*. John Wiley and Sons, Inc., New York, NY, 1999.
- [Kuh55] KUHN H. W.: The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2 (1955), 83–97.
- [LPM01] LEE A. B., PEDERSEN K. S., MUMFORD D.: *The non-linear statistics of high contrast patches in natural images*. Tech. rep., Brown University, 2001. Available online.
- [Mil63] MILNOR J.: *Morse Theory*, vol. 51 of *Annals of Mathematical Studies*. Princeton Univ. Press, Princeton, NJ, 1963.
- [MS74] MILNOR J., STASHEFF J. D.: *Characteristic Classes*, vol. 76 of *Annals of Mathematical Studies*. Princeton Univ. Press, Princeton, NJ, 1974.
- [ZC04] ZOMORODIAN A., CARLSSON G.: Computing persistent homology. *Discrete Comput. Geom.* (2004). (To appear).
- [ZY96] ZHU S., YUILLE A.: Forms: A flexible object recognition and modeling system. *Int. J. Comp. Vision* 20, 3 (1996), 187–212.