International Journal of Shape Modeling
Vol. 11, No. 2 (2005) 149–187
© World Scientific Publishing Company

World Scientific
www.worldscientific.com

PERSISTENCE BARCODES FOR SHAPES

GUNNAR CARLSSON
Department of Mathematics, Stanford University
Stanford, CA 94305, USA
gunnar@math.stanford.edu

AFRA ZOMORODIAN
Department of Computer Science, Stanford University
Stanford, CA 94305, USA
afra@cs.stanford.edu

ANNE COLLINS
Department of Mathematics, Centre College
Danville, KY 40422, USA
collins@centre.edu

LEONIDAS J. GUIBAS
Department of Computer Science, Stanford University
Stanford, CA 94305, USA
guibas@cs.stanford.edu

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

In this paper, we initiate a study of shape description and classification via the application of persistent homology to 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 parameterizing 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 parameterized 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 reasonably nice 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 are identical up to a smooth diffeomorphism of the ambient Euclidean space? This is a coarser classification than the previous one.
- 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.
- What are the *features* of the shape? For example, these features could be singular points of various types.

We wish 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 most of our questions have a *qualitative* feel and take a higher view of a shape. This prompts us to look at topological techniques that 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 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 three tangential constructions: the tangent complex, the filtered tangent complex, and the tame tangent complex. The *tangent complex* is the closure of the space of all unit 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

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

indistinguishable. For example, an ellipse differs from a circle in having a range of curvatures, a feature not reflected in the homology of its tangent complex.

For even more differentiating power, we enrich our tangent complex with information about the curvature of a space at each point. The resulting *filtered tangent complex* is an increasing family of spaces, parameterized 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. Finally, the *tame tangent complex* is the union of the increasing family that constitutes the filtered tangent complex. It can detect certain singular features overlooked by the tangent complex, such as cone points. 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. To compare barcodes, we define a metric on the set of barcodes. The metric enables the use of 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 address the application of this method to the PCD domain for curves in another paper.[4]

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.[8,10] Medial axis techniques have been extensively used for this purpose.[21] One particular aspect of this problem is understanding "shape spaces" in their entirety.[2,15] For instance, one may study a shape space from a differential point of view.[18] 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[20] and *Donaldson theory*, which constructs invariants of smooth structures on homotopy equivalent four-dimensional manifolds using moduli spaces of self-dual connections.[6] The tan-

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

gential complexes we study are familiar in *geometric measure theory.*[9] *Morse theory* analyzes the topology of a space by utilizing increasing families of spaces.[19] *Persistent homology* was first defined for topological simplification of alpha shapes.[7] It was then extended to arbitrary d-dimensional spaces and arbitrary fields of coefficients using a classification result.[22] The latter paper also introduced the notion of using intervals as homological invariants.

We wish to make this paper mathematically rigorous and complete, while keeping it accessible to non-specialists. These goals often conflict. Therefore, we have augmented the mathematics with intuitive explanations and a large number of illustrated examples. The non-topologist should feel free to ignore the theoretical sections, and focus on the intuition necessary to understand the main results. Similarly, a specialist may skip the standard definitions without incurring any loss. We emphasize, however, that the bulk of this paper is dedicated to the study of examples that illustrate the strengths of our tangential constructions for shape classification and recognition.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

1.3. *Overview*

We organize the rest of the paper as follows. Section 2 is a brief tutorial on some techniques from algebraic topology. We do not assume that the reader will become conversant with these tools, but hope that this section will serve as a guide for the rest of the paper. In Section 3, we define our first derived space, 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 derived spaces in Section 4, the filtered and tame tangent complexes, that are more sensitive. 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 all of 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 two extended studies of families of shapes. We illustrate the effect of varying parameters on the barcodes and the distance between shapes by studying the parameterized families of a glass and a bottle.

2. A Brief Tutorial on Algebraic Topology

In this section, we will give a brief introduction to the ideas and terminology of algebraic topology. The discussion will be informal and is intended to give only an overview. For a more complete account, see a standard text in the area, such as Greenberg and Harper, or Hatcher.[12,13]

2.1. *Homotopy*

The central notion in algebraic topology is that of a *homotopy* of maps between topological spaces. Informally, two maps $f, g\colon X \to Y$ are said to be homotopic if there is a continuous one parameter family of continuous maps $\{h_t\}_{t\in[0,1]}\colon X \to Y$ such that $h_0 = f$ and $h_1 = g$. More precisely, f and g are homotopic if there is a continuous map $h\colon X \times [0,1] \to Y$, called a *homotopy between f and g*, so that $h(x,0) = f(x)$ and $h(x,1) = g(x)$. We write $f \simeq g$ whenever f is homotopic to g. We look at a few simple examples.

Example 2.1 (continuous maps). Any two continuous maps $f, g\colon X \to \mathbb{R}^n$ are homotopic, for any space X, since we can write down the explicit homotopy $h(x,t) = (1-t)\cdot f(x) + t\cdot g(x)$.

Example 2.2 (circle rotation). Let X and Y both be the unit circle in the plane, and let $f_\theta\colon S^1 \to S^1$ denote the map that rotates any point in S^1 by an angle θ. Then for any angles θ and η, we have $f_\theta \simeq f_\eta$, because one can continuously vary the angle θ until one obtains η, and simultaneously vary the map f_θ.

Example 2.3 (circle and annulus). Let X be the unit circle, and let Y denote the annulus in the plane given by $\{(x,y) \mid 1 \le x^2+y^2 \le 4\}$. We have two continuous maps f and g from S^1 into Y, the inclusion on the inner boundary $f(x,y) = (x,y)$, and the inclusion on the outer boundary $g(x,y) = (2x,2y)$. The two maps are homotopic via the map $h\colon S^1 \times [0,1] \to Y$ given by $h((x,y),t) = ((1+t)x,(1+t)y)$.

Example 2.4 (wrapping). Let $X = Y = S^1$, let $f\colon S^1 \to S^1$ denote the identity map, and let g be the map given in polar coordinates by $g(1,\theta) = (1,2\theta)$. This map wraps the circle around itself twice. When S^1 is interpreted in terms of complex numbers of length one, g is given by $g(z) = z^2$. We can show that $f \not\simeq g$.

Example 2.5 (punctured $\mathbb{R}^3$). Let X denote the two-sphere $S^2 \subseteq \mathbb{R}^3$, and let Y denote the subspace $\mathbb{R}^3 - \{(0,0,0)\} \subseteq \mathbb{R}^3$ that contains X. Let f denote the inclusion map $X \hookrightarrow Y$, and let g denote the constant map with value $(1,0,0)$. We can show that $f \not\simeq g$.

It is frequently not difficult to show that two maps are homotopic, since all that is required is to construct a homotopy. To show that two maps are not homotopic, on the other hand, can not typically be done with elementary methods and requires more advanced techniques of algebraic topology.

2.2. *Homotopy Equivalence*

We now define several related notions. A map $f\colon X \to Y$ is a *homotopy equivalence* if there is a map $g : Y \to X$ so that fg is homotopic to the identity map on Y and gf is homotopic to the identity map on X. This notion is a weakening of the

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

notion of *homeomorphism*, that requires the existence of a continuous map g so that fg and gf are *equal* to the corresponding identity maps. The less restrictive notion of homotopy equivalence is useful in understanding relationships between complicated spaces and spaces with simple descriptions. We say two spaces X and Y are *homotopy equivalent*, or have the same *homotopy type* if there is a homotopy equivalence from X to Y. This is denoted by $X \simeq Y$. If $X \subseteq Y$ is the inclusion of a subspace, then X is a *deformation retract* of Y if there is a map $h\colon Y \times [0,1] \to Y$, so that $h(y,0) = y$ for all $y \in Y$, $h(x,t) = x$ for all $x \in X, t \in [0,1]$, and $h(y,1) \in X$ for all $y \in Y$. Note that if X is a deformation retract of Y, then X and Y are homotopy equivalent: let i denote the inclusion of X into Y and $h_1(y) = h(y,1)$, then $h_1 i$ is equal to the identity map on X, and ih_1 is homotopic to the identity on Y. Finally, we say a space is *contractible* if it is homotopy equivalent to a point.

Example 2.6 (annulus). Let A be the annulus. For variety, we parameterize the annulus in polar coordinates this time. Let $1 \leq r \leq 2$, and let f be the map from A to the unit circle given by $f(r,\theta) = (1,\theta)$. Then f is a homotopy equivalence, since the inclusion of the circle into A provides the required map g.

Example 2.7 (S^n). The subspace $S^1 \subseteq \mathbb{R}^2 - \{(0,0)\}$ is a deformation retract of $\mathbb{R}^2 - \{(0,0)\}$. Consider the map $h\colon (\mathbb{R}^2 - \{(0,0)\}) \times [0,1] \to \mathbb{R}^2 - \{(0,0)\}$ given by $h(v,t) = \frac{v}{\|v\|^t}$. Similarly, S^{n-1} is a deformation retract of $\mathbb{R}^n - \{(0,\ldots,0)\}$. But according to Example 2.5, it is not contractible.

Example 2.8 ($\mathbb{R}^n$). The point $(0,\ldots,0)$ is a deformation retract of $\mathbb{R}^n$ since we have the map $h\colon \mathbb{R}^n \times [0,1] \to \mathbb{R}^n$ given by $h(v,t) = (1-t)\cdot v$. It follows that $\mathbb{R}^n$ is homotopy equivalent to the single point $(0,\ldots,0)$ and is contractible.

Example 2.9 (matrices). Let $SL(\mathrm{n})$ be the *Special Linear* group, that is, the space of all $n \times n$ matrices with determinant 1. Let $SO(\mathrm{n})$ be the *Special Orthogonal* group, that is the space of orthonormal matrices. Clearly, $SO(\mathrm{n})$ is a subspace of $SL(\mathrm{n})$. Using the Gram-Schmidt orthonormalization procedure, moreover, we can show that $SO(\mathrm{n})$ is a deformation retract of $SL(\mathrm{n})$.

2.3. *Homology*

Algebraic topology is a set of algebraic tools for recognizing when maps between spaces are not homotopic, when spaces are not homotopy equivalent, and at times, when spaces might indeed be equivalent. These tools are rather complicated to construct, so we will be content with a brief descriptive account. For any topological space X, we associate a family of Abelian groups $H_n(X)$ for every integer $n \geq 0$. They have the following key properties:

Functoriality: Each H_n is a *functor*, that is, for any continuous map $f\colon X \to Y$, there is an induced homomorphism $H_n(f)\colon H_n(X) \to H_n(Y)$, such that $H_n(fg) = H_n(f)H_n(g)$ and $H_n(i_X) = i_{H_n(X)}$, where i is the identity.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

Homotopy Invariance: If $f, g\colon X \to Y$ are homotopic, then $H_n(f) = H_n(g)$. If f is a homotopy equivalence, then $H_n(f)$ is an isomorphism.

The significance of these properties may not be clear to the reader. The original developers of algebraic topology did not deal with groups but with numbers. One collection of numbers, $\{\beta_n(X)\}$, called the *Betti numbers*, consists of the *ranks* of the Abelian groups $\{H_n(X)\}$ we discussed above. That is, the Betti number $\beta_n(X)$ is the minimum number of generators of the torsion-free part of the group $H_n(X)$. Another collection, the *torsion coefficients*, contains the sizes of the cyclic summands in the torsion part of the group.[17] While these numbers reflect the properties of spaces, simply dealing with numerical invariants led to theoretical and practical difficulties due to two important considerations:

(1) It is not obvious how to show the invariance of these numbers under homotopy equivalence. This is a severe limitation, as the major use for invariants is determining whether two spaces are homotopy equivalent.
(2) It is not possible to compute these numerical invariants without using groups. This will become clearer when we look at the *Mayer-Vietoris sequence* later.

It was Emmy Noether's insight that one should deal not with numbers, but with Abelian groups, that lead to the development of algebraic topology as a powerful method for studying topological spaces.

There are some variants of homology groups that are useful computationally. For any ring R, we can define *homology groups with coefficients in* R, denoted $H_n(X, R)$, with the original definition being equivalent to $H_n(X, \mathbb{Z})$. When the ring is a field F, such as the finite fields $\mathbb{Z}/p\mathbb{Z}$ for a prime p, the rational numbers $\mathbb{Q}$, or the real numbers $\mathbb{R}$, the homology groups become vector spaces over the field F. As such, they are determined up to isomorphism by their dimension. We write $\beta_n(X, F)$ for the dimension of $H_n(X, F)$.

2.4. *Intuitive Homology*

By itself, the definition of homology provides no intuition as to what it measures. It is rather difficult to formulate a precise answer to this question. Generally, we can say that homology measures the connectivity of a space. For torsion-free spaces in three-dimensions, the Betti numbers have intuitive meaning as a consequence of the *Alexander Duality*. β_0 measures the number of components of the complex. β_1 is the rank of a basis for the *tunnels*, loops that cannot be deformed to a point. β_2 counts the number of *voids*, enclosed spaces that look like a room. Before going any further, we look at a few examples for intuition.

Example 2.10 (circle). Let X be the unit circle in the plane. Its non-zero Betti numbers are $\beta_0 = \beta_1 = 1$. So, this space is connected and has a non-trivial loop that cannot be deformed to a point *while staying in the space.*

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

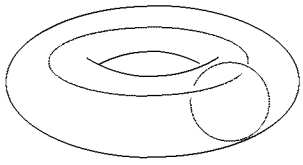


Fig. 1. Two independent non-trivial loops on the torus.

Example 2.11 (figure eight). Let X be a figure eight. Its non-zero Betti numbers are $\beta_0 = 1$ and $\beta_1 = 2$. So, this space is connected and has two independent loops, non-deformable to a point or to each other.

Example 2.12 (tunnel). Let X be the result of removing a line from $\mathbb{R}^3$. Its non-zero Betti numbers are $\beta_0 = \beta_1 = 1$. This reflects the connectedness of the space, as well as the presence of a loop that goes around the line, and cannot be deformed to a point because the line has been removed.

Example 2.13 (void). Let X be the result of removing an open unit ball from $\mathbb{R}^3$. Its non-zero Betti numbers are $\beta_0 = \beta_2 = 1$. The space is connected and the value of β_2 indicates the presence of a "two-dimensional" non-trivial loop. If we had removed k disjoint open balls from $\mathbb{R}^3$, β_2 would have been k.

Example 2.14 (torus). Let X be a torus (the surface of a donut). Its Betti numbers are $\beta_0 = \beta_2 = 1$ and $\beta_1 = 2$. The space is connected and has two independent loops, as shown in Figure 1. The value of β_2 indicates the presence of a two-dimensional enclosed space that the torus encloses.

2.5. *Mayer-Vietoris Sequence*

The main tool for computing the homological invariants of topological spaces is the *Mayer-Vietoris sequence.* To define it, we first need to discuss the notion of an *exact sequence.*

Definition 2.1 (long exact sequence). A sequence

$$A \xrightarrow{f} B \xrightarrow{g} C \tag{1}$$

of homomorphisms of Abelian groups (or linear transformations of vector spaces over a field F) is said to be *exact* if the kernel of g is equal to the image of f, as subgroups (or subspaces) of B. A sequence of homomorphisms (or linear transformations)

$$\cdots \to A_{i+1} \xrightarrow{f_{i+1}} A_i \xrightarrow{f_i} A_{i-1} \xrightarrow{f_{i-1}} A_{i-2} \xrightarrow{f_{i-2}} \cdots \tag{2}$$

is exact if each subsequence $A_{k+1} \to A_k \to A_{k-1}$ is exact. Such a sequence is called a *long exact sequence.*

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

Roughly, if one has information about two of every three consecutive groups in such a sequence, one can obtain information about the third. In the case of vector spaces over a field, we may obtain a very explicit statement.

Proposition 2.1. *Suppose we have a long exact sequence of vector spaces over a field F of the form*

$$\cdots \longrightarrow B_{i+1} \longrightarrow C_{i+1} \xrightarrow{f_{i+1}} A_i \longrightarrow B_i \longrightarrow C_i \xrightarrow{f_i} A_{i-1} \longrightarrow B_{i-1} \longrightarrow \cdots, \quad (3)$$

then we have the relationships

$$\dim(B_i) = \dim(C_i) + \dim(A_i) - \operatorname{rank}(f_i) - \operatorname{rank}(f_{i+1}). \quad (4)$$

In other words, the vector spaces B_i are recoverable from the vector spaces A_i and C_i, together with rank information about the transformations between them. Although the results cannot be stated as simply and explicitly, a similar proposition is true for Abelian groups. That is, information, such as rank of the B_i's, is recoverable from the A_i's and C_i's, together with the information about the homomorphisms between them. We may now state the main tool for homological computation.

Theorem 2.1 (Mayer-Vietoris). *Suppose we have a topological space X with two subspaces $Y, Z \subseteq X$, such that $X = \mathring{Y} \cup \mathring{Z}$. That is, the interiors of Y and Z cover X. Then, there exists a long exact sequence*

$$\cdots \to H_i(Y \cap Z) \to H_i(Y) \oplus H_i(Z) \to H_i(X) \to \\ \to H_{i-1}(Y \cap Z) \to H_{i-1}(Y) \oplus H_{i-1}(Z) \to \cdots. \quad (5)$$

Analogous sequences exist for homology with coefficients. The homomorphism $H_i(Y \cap Z) \to H_i(Y) \oplus H_i(Z)$ is given by $x \mapsto (H_i(i_Y)(x), -H_i(i_Z)(x))$, where $i_Y : Y \cap Z \hookrightarrow Y$ and $i_Z : Y \cap Z \hookrightarrow Z$ are the inclusions of subspaces.

There is also a formula for describing the homology of a product of two spaces.

Theorem 2.2 (Künneth Formula). *For any spaces X and Y, and any field F, there is a canonical isomorphism*

$$H_n(X \times Y, F) \cong \bigoplus_{i=0}^{n} H_{n-i}(X, F) \otimes H_i(Y, F), \quad (6)$$

where $\otimes$ denotes tensor product.[17] *It follows that*

$$\beta_n(X \times Y, F) = \sum_{i=o}^{n} \beta_{n-i}(X, F) \cdot \beta_i(Y, F). \quad (7)$$

These results, together with the homotopy invariance property mentioned above, provide an extremely effective method for computing the Betti numbers of topological spaces. We do not give examples here, but refer the reader to the standard textbooks.[12,13]

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

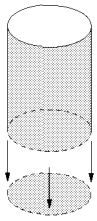


Fig. 2. The cylinder $S^1 \times [0,1]$ and the disc D^2 are glued together along the dashed circles. The resulting space $D^2 \cup_{S^1} (S^1 \times [0,1])$ has the same geometry as a glass or bottom of a bottle.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

We conclude this section by mentioning a space-level construction we will be using 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 2.2 ($A \cup_X B$). Suppose we have a space X, together with inclusions $f\colon X \hookrightarrow A$ and $g\colon X \hookrightarrow B$. Then, we may construct a new space as the quotient of the disjoint union of A and B,

$$A \cup_X B = A \dot{\cup} X \dot{\cup} B / \sim, \tag{8}$$

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 $f\colon X \hookrightarrow A$, $g\colon X \hookrightarrow B$, $h\colon Y \hookrightarrow B$, and $k\colon Y \hookrightarrow C$, we define

$$A \cup_X B \cup_Y C = (A \cup_X B) \cup_Y C. \tag{9}$$

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

Example 2.15 (glass). Figure 2 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 circle $X = S^1$, whose inclusion into A and B is indicated by dashed circles, to glue the two spaces together. The gluing gives $D^2 \cup_{S^1} (S^1 \times [0,1])$, which looks like a glass or the bottom of a bottle.

3. The Tangent Complex

In this section, we define 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.[9] After defining the complex, we give a number of examples for intuition.

Definition 3.1 (tangent complex). Let X be any subset of $\mathbb{R}^n$. We define $T^0(X) \subseteq X \times S^{n-1}$ to be

$$T^0(X) = \left\{ (x, \zeta) \;\middle|\; \lim_{t \to 0} \frac{\mathrm{d}(x + t\zeta, X)}{t} = 0 \right\}. \tag{10}$$

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

Proposition 3.1. *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 $f\colon U \to \mathbb{R}^m$, such that*

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

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

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.

Example 3.1 (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, normal to the line. Then we have $\omega(T(L)) = \{\pm\eta\}$, where η is a unit vector perpendicular to ξ, and

$$T(L) \approx L \times \{\pm\eta\}. \tag{11}$$

We show the line and its tangent complex in Figure 3. 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, as shown for the point x in the figure. 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 $T(X) \approx X \times S(\xi)$, 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 X.

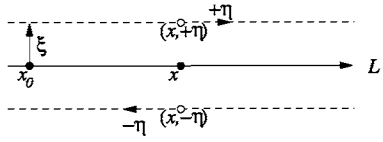


Fig. 3. The solid line L is defined by point x_0 and normal ξ. It has a tangent complex that is composed of the two dashed components in the directions $\pm\eta$. The fiber at point x consists of two points $(x, \pm\eta)$.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

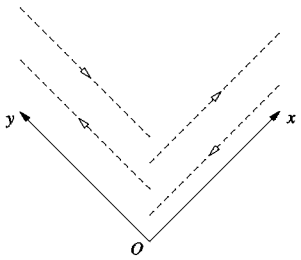


Fig. 4. Solid 'V' space X has a tangent complex $T(X)$ that is composed of the four dashed components.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

Example 3.2 (vee). Consider a space $X \subseteq \mathbb{R}^2$ that looks like a 'V', as shown in Figure 4. 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:

$$T(\mathbb{R}_+ \times 0) = (\mathbb{R}_+ \times 0) \times \{(\pm 1, 0)\}, \tag{12}$$

$$T(0 \times \mathbb{R}_+) = (0 \times \mathbb{R}_+) \times \{(0, \pm 1)\}. \tag{13}$$

So that the entire complex is:

$$T(X) = (\mathbb{R}_+ \times 0) \times \{(\pm 1, 0)\} \cup (0 \times \mathbb{R}_+) \times \{(0, \pm 1)\}. \tag{14}$$

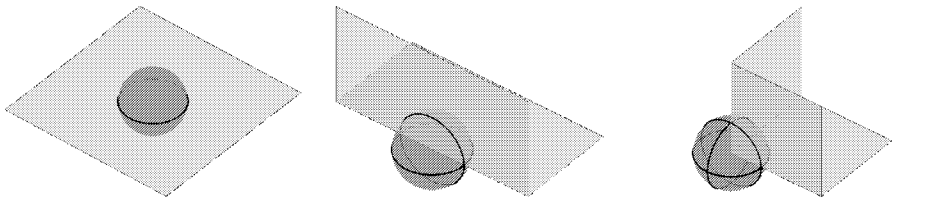
This space is a disjoint union of four distinct half lines, as shown in Figure 4.

Example 3.3 (boundary of positive octant). Let $X \subseteq R^3$ be the boundary of the positive octant,

$$X = \{(x, y, z) \mid x, y, z \geq 0, \text{ and at least one of } x, y, z \text{ is equal to zero}\}. \tag{15}$$

In this case, X can be decomposed into three pieces, its intersections with the three coordinate planes. We denote the intersection of X with the xy-plane by X_{xy} and the other two intersections by X_{yz} and X_{xz}, accordingly. Each of these intersections is a quadrant in the corresponding coordinate plane. From Example 3.1, we know that $T(X_{xy}) \approx X_{xy} \times S^1_{xy}$, where S^1_{xy} denotes the unit circle in the xy-plane. There are similar descriptions for the other coordinate planes. We now examine the fibers of $T(X) \to X$:

(1) Smooth point: for any point v not lying on a coordinate axis, the fiber $T(X)_v$ is a circle, as shown in Figure 5(a).
(2) Edge point: for any point v that lies on a coordinate axis, but is not at the *corner point* $(0, 0, 0)$, $T(X)_v$ is the union of two circles. The circles, shown in Figure 5(b), intersect at a pair of antipodal points.



(a) Smooth point: one circle (b) Edge point: two circles (c) Corner point: three circles

Fig. 5. $X \subset \mathbb{R}^3$ is the boundary of the positive octant. We show the fibers at three types of points.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

(3) Corner point: $T(X)_{(0,0,0)}$ is the union of three circles. The circles, shown in Figure 5(c), intersect pairwise at pairs of antipodal points. $T(X)$ contains the fiber at the corner point as a deformation retract, so they have the same homotopy type.

Example 3.4 (cone). Let X denote the cone given by $z^2 = x^2 + y^2$, and $z \leq 0$. We would like to distinguish the *cone point* at the origin from the other *regular points* using the tangent complex. Every regular point p is smooth, so the fiber $T(X)_p$ is a full circle. In fact, two regular points that lie on the same line through the origin have the same fiber, as shown in Figure 6. The fiber at the origin consists of the union of all circles occurring at the regular points, or equivalently, the set of points (x, y, z) in S^2 with $|z| \leq \frac{\sqrt{2}}{2}$. This set is homeomorphic to an annulus, which is homotopy equivalent to a circle, as shown in Example 2.3. Therefore, we cannot distinguish the cone point from the regular points using the tangent complex $T(X)$.

The features detected by homology of the tangent complex are *sharp features*: they are unaffected by smooth diffeomorphisms of the ambient space $\mathbb{R}^n$. How-

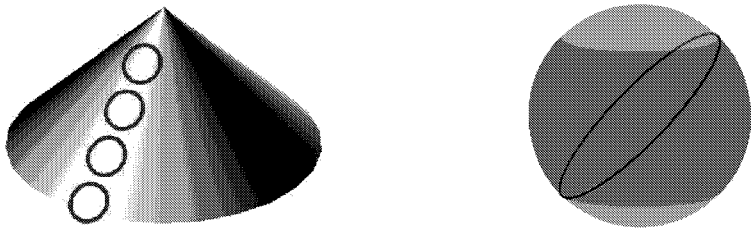


Fig. 6. Bottom half of cone $x^2 + y^2 = z^2$. The fiber $T_\delta(X)_p$ at each smooth point p is a full circle (left). These fibers sweep out an annulus (right), which is the fiber at the origin.

ever, 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 3.5 (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 3.6 (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 from 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 sophisticated constructions, which we introduce in the next section.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

4. The Filtered and Tame Tangent Complexes

In this section, we define two constructions based on tangential information: the *filtered* and *tame* tangent complexes. We devote the bulk of this section to detailed examples, building the complexes 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*

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 a curve example in Figure 7. 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 4.1 (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

$$\lim_{\theta \to 0} \frac{\mathrm{d}(x_0 + \cos\theta \cdot (x - x_0) + \sin\theta \cdot |x - x_0| \cdot \zeta, X)}{\theta^2} = 0. \tag{16}$$

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

Definition 4.2 (filtered, tame). Let $T^0_\delta(X)$ 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

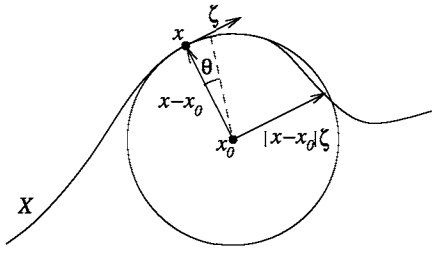


Fig. 7. (x, ζ) is in $T^0(X)$, i.e. point x on curve X has unit tangent ζ. The osculating circle at x is centered at x_0. The perpendicular vectors $x - x_0$ and $|x - x_0|\zeta$ and parameter θ provide coordinates for points on the circle.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

the closure of $T^0_\delta(X)$ in $\mathbb{R}^n \times S^{n-1}$. We call the δ-parameterized family of spaces $\{T_\delta(X)\}_{\delta \geq 0}$ the *filtered tangent complex*, denoted by $T^{filt}(X)$. We call the union $\bigcup_\delta T_\delta(X)$ the *tame tangent complex*, denoted by $T^{tame}(X)$.

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 $\pi_\delta \colon T_\delta(X) \to X$ is the natural projection. There is also a natural inclusion $T^{tame}(X) \hookrightarrow T(X)$ that is not surjective in general. When X is a smooth submanifold of $\mathbb{R}^n$, then $T^{tame}(X) = T(X)$. For non-smooth spaces, however, $T^{tame}(X)$ has additional power that allows us to detect cone points. The topological structure of $T^{filt}(X)$ carries information about the soft features of the underlying space.

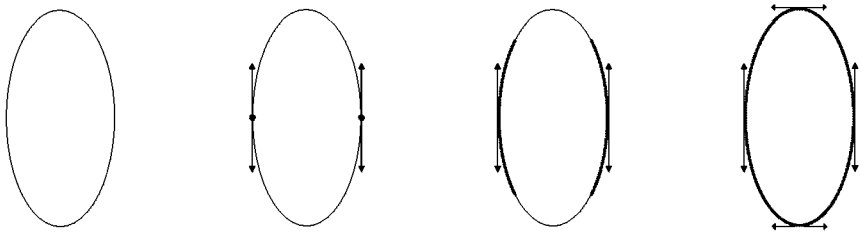
4.2. *Examples: Curves*

We elucidate these concepts through examples in the rest of this section. We begin with one-dimensional subsets of $\mathbb{R}^2$, or curves. Each point x of a smooth curve X has a circle of second order contact with radius $\rho(x)$. So, $T_\delta(X)$ has a projection onto the *excursion set* $X_\delta = \{x \in X \mid \delta \geq 1/\rho(x)\}$ and the fiber over a point of X_δ consists of a pair of points.

Example 4.1 (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,

$$T_\delta(X) = \begin{cases} \emptyset, & \text{for } \delta < 1/R, \\ T(X), & \text{for } \delta \geq 1/R. \end{cases} \tag{17}$$

Example 4.2 (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 parameterized



(a) $0 \leq \delta < \frac{a}{b^2}$ (b) $\delta = \frac{a}{b^2}$ (c) $\frac{a}{b^2} < \delta < \frac{b}{a^2}$ (d) $\delta \geq \frac{b}{a^2}$

Fig. 8. Ellipse $x^2/a^2 + y^2/b^2 = 1$. The points whose fibers have two components are indicated with thick lines. We show the case with $a = 1$ and $b = 2$.

curve $\varphi(t) = (x(t), y(t))$ is given by

$$\frac{|\varphi'(t)|^3}{|x'(t)y''(t) - x''(t)y'(t)|}. \tag{18}$$

A parameterization 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

$$\frac{(a^2 + (b^2 - a^2)\cos^2(t))^{\frac{3}{2}}}{ab}, \tag{19}$$

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 $\left(0, \frac{\pi}{2}\right)$ and $\left(\pi, \frac{3\pi}{2}\right)$, and increasing in intervals $\left(\frac{\pi}{2}, \pi\right)$ and $\left(\frac{3\pi}{2}, 2\pi\right)$. Therefore, the filtered tangent complex behaves as follows. For each case, there is a corresponding illustration in Figure 8.

(a) For $0 \leq \delta < \frac{a}{b^2}$, $T_\delta(X) = \emptyset$.
(b) For $\delta = \frac{a}{b^2}$, $T_\delta(X)$ consists of four points, namely $(\pm a, 0) \times (0, \pm 1)$.
(c) 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 $(0, \pm b)$.
(d) 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.

Example 4.3 (cubic). Consider the case where X is the graph of a cubic polynomial. Standard translation and scaling show that it suffices to consider the family of polynomials $f(x) = x^3 + ax$, where a is a constant. Recall that we have the formula for the radius of the osculating circle to the graph of a function f, given by

$$\rho = \frac{(1 + (f'(x))^2)^{\frac{3}{2}}}{|f''(x)|}. \tag{20}$$

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

To compute the filtered tangent complex, we need to analyze the behavior of this function for f. Observe that it suffices to consider only the values $x \geq 0$ as negative values of x may be obtained by reflection through the origin. We are interested in the critical points of ρ (maxima and minima) that divide the parameter space into intervals where ρ is monotonic. It is sufficient to consider to square of ρ, the rational function

$$\rho^2 = \frac{(1 + (f'(x))^2)^3}{f''(x)^2}. \tag{21}$$

We have $f'(x) = 3x^2 + a$ and $f''(x) = 6x$, so

$$\rho^2 = \frac{(1 + (3x^2 + a)^2)^3}{36x^2} \tag{22}$$

$$= \frac{1}{36}\left((1 + a^2)x^{-\frac{2}{3}} + 6ax^{\frac{4}{3}} + 9x^{\frac{10}{3}}\right)^3 \tag{23}$$

$$= \frac{g^3}{36}. \tag{24}$$

Clearly, it is adequate to analyze the behavior of the term g. Its derivative is given by

$$g' = -\frac{2}{3}(1 + a^2)x^{-\frac{5}{3}} + 8ax^{\frac{1}{3}} + 30x^{\frac{7}{3}} \tag{25}$$

$$= x^{-\frac{5}{3}}\left(30x^4 + 8ax^2 - \frac{2}{3}(1 + a^2)\right). \tag{26}$$

For positive x, this function has a single real root at

$$x_a = \sqrt{\frac{-2a + \sqrt{5 + 9a^2}}{15}}. \tag{27}$$

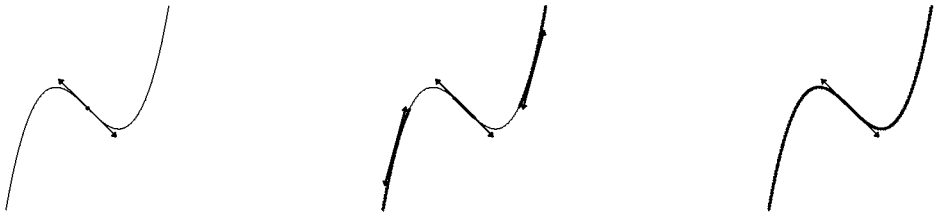
The function g' is negative to the left of x_a and positive to its right. Therefore, g and ρ attain a global minimum at $x = x_a$. Setting $\kappa_a = 1/\rho(x_a)$, we have the following cases for $T_\delta(X)$. As in the previous example, the cases are illustrated in Figure 9.

(a) $T_0(X)$ consists of two points, namely $(0, 0) \times \left(\pm\frac{1}{\sqrt{1+a^2}}, \pm\frac{a}{\sqrt{1+a^2}}\right)$.
(b) For $0 < \delta < \kappa_a$, $T_\delta(X)$ consists of 6 connected components, whose projections onto the curve X give three disjoint arcs in the plane, two unbounded and one bounded.
(c) For $\delta \geq \kappa_a$, $T_\delta(X) = T(X)$, the full tangent complex, which consists of two copies of the curve, one for each direction.

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 is at most a single osculating

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.



(a) $\delta = 0$ (b) $0 < \delta < \kappa_a$ (c) $\delta \geq \kappa_a$

Fig. 9. Function $f(x) = x^3 - x$. The points whose fibers have two components are indicated with thick lines.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

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 filtered 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 10.

With respect to this coordinate system, the best approximating quadratic surface to X is given by an equation of the form

$$z = Q(x, y) = Ax^2 + Bxy + Cy^2. \tag{28}$$

When restricted to the unit circle on the tangent plane, Q attains its minimum and maximum values at pairs of antipodal points that determine a pair of perpendicular directions called 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 *principal*

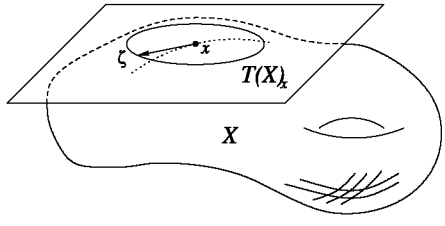


Fig. 10. Surface X with the tangent plane at x and the unit tangent circle $T(X)_x$. We also show a dotted portion of the osculating circle in the direction ζ.

curvatures, the minimum and maximum values of $2|Q|$ on $T(X)_x$. The following cases describe $T_\delta(X)_x$.

(a) If $\delta < \kappa_1$, then $T_\delta(X)_x = \emptyset$.
(b) If $\delta = \kappa_1 < \kappa_2$, then $T_\delta(X)_x$ consists of two antipodal points.
(c) If $\kappa_1 < \delta < \kappa_2$, $T_\delta(X)_x$ consists of two disjoint closed antipodal arcs.
(d) 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 4.2, illustrated in Figure 8, if we replace the ellipse with a unit circle. Given our analysis, we may now examine some surfaces in $\mathbb{R}^3$.

Example 4.4 (ellipsoid). Suppose that X is the ellipsoid given by the equation

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1 \tag{29}$$

with $0 < a < b < c$. In order to analyze $T_\delta(X)$, we must understand when this space undergoes a change in homology. We use a parameterization of the ellipsoid using *elliptic coordinates* from Berger and Gostiaux.[1] The parametrization, shown in Figure 11, maps (u, v) to

$$\left(\sqrt{\frac{a^2(a^2-u)(a^2-v)}{(a^2-b^2)(a^2-c^2)}}, \sqrt{\frac{b^2(b^2-u)(b^2-v)}{(b^2-a^2)(b^2-c^2)}}, \sqrt{\frac{c^2(c^2-u)(c^2-v)}{(c^2-a^2)(c^2-b^2)}} \right). \tag{30}$$

Here, the vector (u, v) ranges over all values with $a^2 \leq v \leq b^2 \leq u \leq c^2$. This is a bijective parameterization of the surface in the positive octant. By making different choices of square root sign for each of the coordinates, we may coordinatize the surface in all octants. The boundary of the positive portion intersects the three positive quadrants. The resulting boundary curves are parameterized as follows, as shown in Figure 11.

- xy-plane: (c^2, a^2) to (c^2, b^2)
- yz-plane: (b^2, a^2) to (c^2, a^2)
- xz-plane: (b^2, a^2) to (b^2, b^2) to (c^2, b^2) (note the resulting cusp in Figure 11(b).)

The squares of the minimum κ_1 and maximum κ_2 principal curvatures on the ellipsoid at the point parameterized by (u, v) are given by

$$\kappa_1^2(u, v) = \frac{a^2 b^2 c^2}{u^3 v}, \tag{31}$$

$$\kappa_2^2(u, v) = \frac{a^2 b^2 c^2}{u v^3}. \tag{32}$$

The globally minimum values of both κ_1 and κ_2 occur when $(u, v) = (c^2, b^2)$, while their globally maximum values occur when $(u, v) = (b^2, a^2)$. These are the points $(\pm a, 0, 0)$ and $(0, 0, \pm c)$ on the ellipsoid, respectively. It turns out that the homology of $T_\delta(X)$ changes as δ passes through these four critical curvature values, as well as the two principal curvatures at the points $(0, \pm b, 0)$, corresponding to the

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.



(a) (u,v) grid

(b) Parameterized patch

Fig. 11. Elliptic coordinates for the portion of ellipsoid in the positive octant. The parameterization is not orientation preserving, and maps the point (b^2, b^2) to the corner of the cusp in (b).

parameters $(u, v) = (c^2, a^2)$.[19] Therefore, we define $\lambda_1 < \mu_1 < \nu_1$ to be κ_1 at the extreme points of the ellipsoid, $(\pm a, 0, 0)$, $(0, \pm b, 0)$ and $(0, 0, \pm c)$, respectively. We easily compute these to be $\frac{a}{c^2} < \frac{b}{c^2} < \frac{c}{b^2}$. Similarly, we define $\lambda_2 < \mu_2 < \nu_2$ to be κ_2 at the extreme points, respectively. They are $\frac{a}{b^2} < \frac{b}{a^2} < \frac{c}{a^2}$.

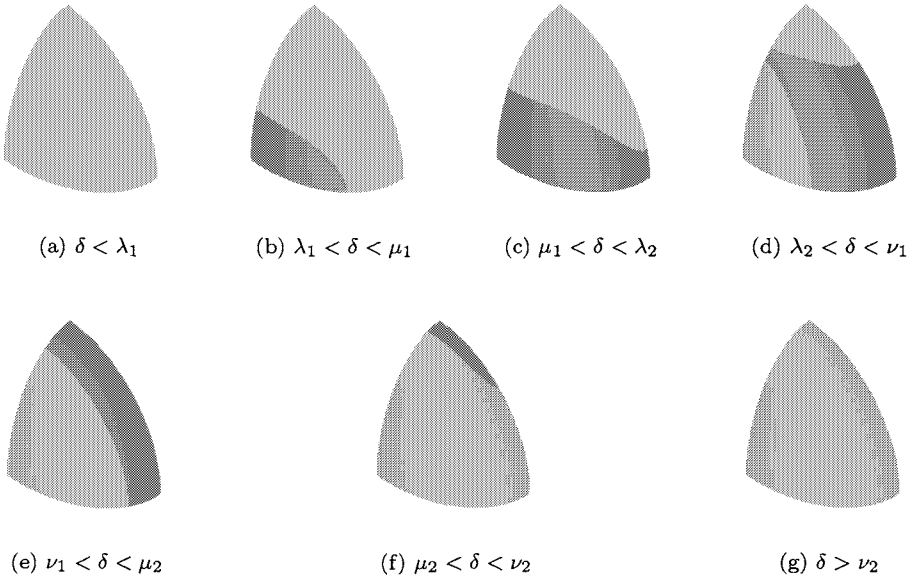
The evolution of $T_\delta(X)$ as δ increases depends upon the ordering of these six critical values, which are in terms of a, b, and c. There are 5 cases: (1) $a^2c < b^3 < ac^2$, (2) $b^3 < a^2c$, (3) $b^3 > ac^2$, (4) $b^3 = a^2c$, and (5) $b^3 = ac^2$. Here, we will treat the case (1) explicitly. The other cases may be dealt with in a similar fashion. For this case, the critical curvature values are ordered as follows.

$$\lambda_1 < \mu_1 < \lambda_2 < \nu_1 < \mu_2 < \nu_2. \tag{33}$$

We now detail the shape of $T_\delta(X)$ for the critical values and the ranges they bound. Figure 12 shows the basepoints of T_δ for various ranges of δ.

- $\delta < \lambda_1$: $T_\delta(X) = \emptyset$ (Figure 12(a)).
- $\delta = \lambda_1$: $T_\delta(X)$ consists of four points, whose basepoints are the points $(\pm a, 0, 0)$. At each such point, we have two vectors tangent to the curve obtained by intersecting X with the xz-plane.
- $\lambda_1 < \delta < \mu_1$: $T_\delta(X)$ consists of a set of points whose projection onto X consists of two disjoint contractible sets surrounding the points $(\pm a, 0, 0)$. The fiber at each of these points consists of either two tangent vectors or two disjoint antipodal arcs. The homotopy type of $T_\delta(X)$ is still that of four discrete points (Figure 12(b)).
- $\delta = \mu_1$: $T_\delta(X)$ is a set of points whose projection onto X consists of the union of two closed contractible discs that contain the points $(\pm a, 0, 0)$, respectively, and intersect in the two points $(0, \pm b, 0)$. This projected subset has the homotopy type of a circle. The fiber of every point is either two antipodal points (such as at $(0, \pm b, 0)$), or a union of two disjoint antipodal intervals. The homotopy type of $T_{\mu_1}(X)$ is that of $S^1 \times S^0$.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.



(a) $\delta < \lambda_1$ (b) $\lambda_1 < \delta < \mu_1$ (c) $\mu_1 < \delta < \lambda_2$ (d) $\lambda_2 < \delta < \nu_1$

(e) $\nu_1 < \delta < \mu_2$ (f) $\mu_2 < \delta < \nu_2$ (g) $\delta > \nu_2$

Fig. 12. Positive octant of ellipsoid $\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1$, colored by number of components in fiber $T_\delta(X)_x$, for various ranges of δ. Light gray is the empty fiber, medium gray the full circle, and dark gray two arcs. In the case where $a < b < c$ and $a^2c < b^3 < ac^2$, the homology of $T_\delta(X)$ changes only at the critical values $\lambda_1 < \mu_1 < \lambda_2 < \nu_1 < \mu_2 < \nu_2$.

- $\mu_1 < \delta < \lambda_2$: $T_\delta(X)$ is a set whose projection onto X consists of the union of two contractible sets that contain the points $(\pm a, 0, 0)$, respectively, and intersect in a disjoint union of two contractible sets containing the points $(0, \pm b, 0)$, respectively. The homotopy type of the projection is that of S^1, and the homotopy type of $T_\delta(X)$ remains $S^1 \times S^0$ (Figure 12(c)).
- $\delta = \lambda_2$: The projection of $T_\delta(X)$ is as in the previous case. The fiber of all except two basepoints consists of two antipodal points or two antipodal disjoint arcs. At the two points $(\pm a, 0, 0)$, the fiber is now a full circle. The homotopy type of this space is that of the space

$$\left(S^1 \times S^0\right) \cup \left(S^0 \times S^1\right) \subseteq S^1 \times S^1. \tag{34}$$

  Note that T_δ has now become connected.
- $\lambda_2 < \delta < \nu_1$: Again, the projection remains the same. The set of points whose fibers consist of full circles, however, is now a disjoint union of two contractible sets containing $(\pm a, 0, 0)$. The homotopy type remains unchanged from the previous case (Figure 12(d)).
- $\delta = \nu_1$: $T_\delta(X)$ is now a set whose projection is the entire ellipsoid X. The fibers are as usual: pairs of antipodal points, pairs of disjoint antipodal arcs, or full circles. The only points where the fiber is two antipodal points are

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

$(0, 0, \pm c)$. The set of points with full circles as fibers also remains as before. The homotopy type of $T_\delta(X)$, however, is more complicated. We define two maps $\varphi, \psi\colon S^1 \times S^0 \to S^1 \times S^1$ by

$$\varphi(\theta, (\pm 1, 0)) = (\theta, \pm\theta), \tag{35}$$

$$\psi(\theta, (\pm 1, 0)) = (\theta, \mp\theta). \tag{36}$$

We also let $i : S^1 \times S^1 \to C(S^1) \times S^1$ denote the inclusion of S^1 into the base of the cone $C(S^1)$. Then the homotopy type of $T_\delta(X)$ is that of the space

$$\left(C(S^1) \times S^1\right) \cup_{S^1 \times S^0} \left(C(S^1) \times S^1\right), \tag{37}$$

based on the inclusions $i\varphi$ and $i\psi$ and Definition 2.2

- $\nu_1 < \delta < \mu_2$: All fibers are now pairs of disjoint antipodal intervals or full circles. The homotopy type remains unchanged (Figure 12(e)).
- $\delta = \mu_2$: The projection and fibers stay as before, but the set of points with full fiber is now connected. The homotopy type of $T_\delta(X)$ changes to be that of

$$\left(C(S^1) \times S^0\right) \cup_{S^1 \times S^0} \left(S^1 \times S^1\right) \cup_{S^1 \times S^0} \left(C(S^1) \times S^0\right), \tag{38}$$

  where $S^1 \times S^0$ is included in $C(S^1) \times S^0$ via the inclusion as the base of the cone, and $S^1 \times S^0$ is included into $S^1 \times S^1$ via the inclusion $S^0 \hookrightarrow S^1$.
- $\mu_2 < \delta < \nu_2$: The homotopy type of $T_\delta(X)$ remains the same as in the previous case (Figure 12(f)).
- $\delta \geq \nu_2$: All fibers are now full circles and $T_\delta(X) = T(X)$ (Figure 12(g)).

Table 1 summarizes our results.

Table 1. Summary of the evolution of $T^{filt}(X)$ for the ellipsoid. A vertical bar indicates no change. To reduce clutter, we omit lower bound of δ, as it is implicitly defined by the higher row.

δ	$T_\delta(X) \simeq \cdots$	Figure
$< \lambda_1$	$\emptyset$	12(a)
$= \lambda_1$	$S^0 \cup S^0$	
$< \mu_1$	|	12(b)
$= \mu_1$	$S^1 \times S^0$	
$< \lambda_2$	|	12(c)
$= \lambda_2$	$\left(S^1 \times S^0\right) \cup \left(S^0 \times S^1\right)$	
$< \nu_1$	|	12(d)
$= \nu_1$	$\left(C(S^1) \times S^1\right) \cup_{S^1 \times S^0} \left(C(S^1) \times S^1\right)$	
$< \mu_2$	|	12(e)
$= \mu_2$	$\left(C(S^1) \times S^0\right) \cup_{S^1 \times S^0} \left(S^1 \times S^1\right) \cup_{S^1 \times S^0} \left(C(S^1) \times S^0\right)$	
$< \nu_2$	|	12(f)
$\geq \nu_2$	$T(X)$	12(g)

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

Example 4.5 (cone). In Example 3.4 we failed to detect a cone point using the tangent complex. We now utilize our new constructions, the filtered tangent complex $T^{filt}(X)$ and the tame tangent complex $T^{tame}(X)$. Each regular point p has a tangent direction ξ that points to the origin. The cone has zero curvature in this direction, so $(p, \xi) \in T_0(X)_p$. It follows that the fiber at the origin consists of vectors that are the directions of lines through the origin that lie on the cone. This is the set $(x, y, z) \in S^2$ with $|z| = \frac{\sqrt{2}}{2}$, homeomorphic to a disjoint union of two circles. Therefore, the full space $T_0(X)$ looks like two disjoint cylinders:

$$S^1 \times S^0 \times [0, 1], \tag{39}$$

and has the same homotopy type as the fiber at the origin. The map from $T_0(X) \to T_0(X)_{(0,0,0)}$ is projection on the factor $S^1 \times S^0$.

For each $p \in X$, the radius $\rho(p, \zeta)$ of the osculating circle in the direction ζ attains its minimum in the direction perpendicular to the line through p and the origin. Moreover, this minimum,

$$c(p) = \min_{\zeta \in T(X)_p} \rho(p, \zeta), \tag{40}$$

depends only on the distance from p to the origin, and increases strictly with the distance. Therefore, the further away we are from the origin, the lower the maximum curvature, and the lower the δ at which the fibers enter T_δ. For a given δ, if $c(p) \geq 1/\delta$, the entire fiber $T(X)_p$ (a circle) is included in $T_\delta(X)$. Otherwise, the fiber contracts into T_δ and has no interesting topology. So, we use a ball of radius $\lambda = |p|$ to cut out the portion of $T(X)$ that does not have full fibers. Let B_λ be an open ball of radius λ. Then,

$$T_\delta(X) \approx T_0(X) \cup \pi^{-1}(\mathbb{R}^3 - B_\lambda) \tag{41}$$

where π is the projection. As δ increases, the infinite annulus we cut out widens, as shown in Figure 13. The space $\pi^{-1}(\mathbb{R}^3 - B_\lambda)$ has the homotopy type of $S^1 \times S^1$, as each point in the annulus has a full circle as fiber. So, $T_\delta(X)$ has this homotopy type for all positive δ. Similarly, $T^{tame}(X)$ also has the same homotopy type and is not equivalent to the space $T(X)$, whose homotopy type was that of S^1. We can

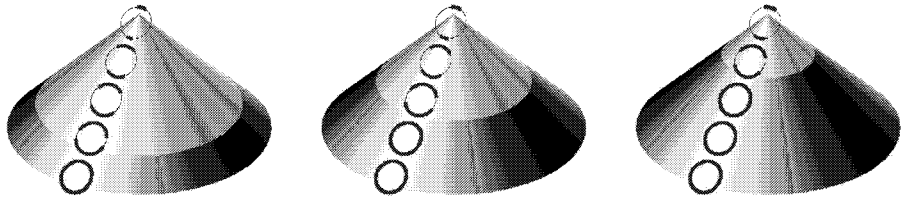


Fig. 13. Bottom half of cone $x^2 + y^2 = z^2$. The fiber $T_\delta(X)_p$ consists of two arcs when $c(p) < 1/\delta$, for points p near the cone point (light), or a full circle when $c(p) \geq 1/\delta$, for points away from cone point (dark). We show the fibers at various heights with thicker lines.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

identify the natural map $T^{tame}(X) \to T(X)$ up to homotopy equivalence of the two spaces with the projection of $S^1 \times S^1$ on the second factor.

Example 4.6 (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.[1] 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 $\kappa^{graph}(v) = 1/\rho^{graph}(v)$ to 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 $T(X) \approx X \times S^1$. The projection of $T_\delta(X)$ is itself a surface of revolution, of the set of points that have directions with low enough curvature:

$$\left\{(x, f(x)) \mid \kappa^{rev} \leq \delta \quad \text{or} \quad \kappa^{graph} \leq \delta\right\}. \tag{42}$$

A non-empty fiber at v consists of two antipodal points or two disjoint antipodal arcs if

$$\left(\kappa^{rev} \leq \delta < \kappa^{graph}\right) \text{ or } \left(\kappa^{graph} \leq \delta < \kappa^{rev}\right), \tag{43}$$

and consists of a full circle, otherwise.

We end this section with an example of surfaces in higher dimensions.

Example 4.7 (linear subspaces of $\mathbb{R}^n$). Suppose that X is a union of linear subspaces in $\mathbb{R}^n$. Then it is clear from the definition that $T_\delta(X) = T(X)$ for all δ.

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 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 that 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 $H_n(T_\delta(X)) \to H_n(T_{\delta'}(X))$ 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 parameterized on the ordered set $[0, \infty)$.[22] Persistence modules capture the information

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

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

$$\varphi_{t,t'} \cdot \varphi_{t',t''} = \varphi_{t,t''}, \tag{44}$$

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

Example 5.1 ($\mathbb{N}$). Let $\mathcal{T} = \mathbb{N} = \{0, 1, 2, \ldots\}$. Then an R-persistence module parameterized by $\mathcal{T}$ is simply a family of R-modules $\{M_n\}_{n\geq 0}$, with homomorphisms

$$M_0 \to M_1 \to \cdots \to M_{n-1} \to M_n \to M_{n+1} \to \cdots . \tag{45}$$

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 5.2 (filtered tangent complex). Suppose that we have a family of spaces X_δ, parameterized by the real valued parameter δ, so that $X_\delta \subseteq X_{\delta'}$ whenever $\delta \leq \delta'$. Then the family of Abelian groups $H_n(X_\delta, R)$, where R is a ring, is a R-persistence module parameterized by $\mathbb{R}$. Note that if R is any field, then we obtain a R-persistence *vector space* parameterized by $\mathbb{R}$ over the field R.

We also speak of a *persistence chain complex parameterized by $\mathcal{T}$* over a ring R, a family of chain complexes $\{C_*^t\}_{t\in\mathcal{T}}$ of R-modules together with chain maps $C_*^t \to C_*^{t'}$ whenever $t \leq t'$, satisfying the analogues of the compatibility relations (44) above. The homology of a persistence complex parameterized by $\mathcal{T}$ over R is always a R-persistence module, parameterized 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.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

5.2. *Classification over Fields*

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

$$Q(m,n)_t = \begin{cases} 0, \text{ if } t < m \text{ or } t \geq n, \\ F, \text{ otherwise,} \end{cases} \tag{46}$$

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 [22] is the following.

Proposition 5.1 (classification). *A persistence module parameterized by* $\mathbb{N}$ *over* F *of finite type is isomorphic to one of the form*

$$\bigoplus_{s=1}^{n} Q(i_s, j_s), \tag{47}$$

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 parameterized by $\mathbb{N}$ over any ring R are characterized by an associated graded module (non-negatively graded) over the polynomial ring $R[t]$, 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

$$Q(I_s) = \begin{cases} 0, \text{ if } s \notin I, \\ F, \text{ otherwise,} \end{cases} \tag{48}$$

where the homomorphism is the identity within each interval. We can now state an analog of Proposition 5.1.

Proposition 5.2. *A persistence module parameterized by* $\mathbb{R}$ *over* F *of finite type is isomorphic to one of the form*

$$\bigoplus_{s=1}^{n} Q(I_s), \tag{49}$$

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

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 parameterized by $\mathbb{N}$ 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 parameterized by $\mathbb{R}$. This family of intervals is our shape descriptor.

Definition 5.2 (barcode). A *barcode* is a finite multiset of intervals that are bounded below.

Intuitively, the intervals denote the life-times of non-trivial cycles in a growing complex. The left endpoint of an interval 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 persists in 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 pseudo-metric (also called quasi-metric in the literature).

Definition 5.3 (pseudo-metric). Let X be a set. A *pseudo-metric* $\mathrm{d}\colon X \times X \to [0, +\infty]$ satisfies the following properties.

(i) $\mathrm{d}(x, y) \geq 0$, for all $x, y \in X$ (positivity),
(ii) $\mathrm{d}(x, x) = 0$, for all $x \in X$ (non-degeneracy),
(iii) $\mathrm{d}(x, y) = \mathrm{d}(y, x)$ for all $x, y \in X$ (symmetry),
(iv) $\mathrm{d}(x, z) \leq \mathrm{d}(x, y) + \mathrm{d}(y, z)$, for all $x, y, z \in X$ (triangle inequality).

We stipulate $x + \infty = \infty$ and $\infty + \infty = \infty$, so that ∞ is a possible value of the pseudo-metric.

Let $\mathcal{I}$ denote the collection of all possible barcodes. We wish to define a pseudo-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 between S_1 and S_2 relative to M* to be the sum

$$\mathcal{D}_M(S_1, S_2) = \sum_{(I,J) \in M} \delta(I, J) + \sum_{L \in N} \mu(L). \tag{50}$$

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

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

Definition 5.4 (barcode 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

$$\mathcal{S}_M(S_1, S_2) = \sum_{(I,J)\in M} \mu(I \cap J) \tag{51}$$

$$= \frac{1}{2}\left(\sum_{S_1} \mu(I) + \sum_{S_2} \mu(J) - \mathcal{D}_M(S_1, S_2)\right). \tag{52}$$

Minimizing $\mathcal{D}_M$ is equivalent to maximizing $\mathcal{S}_M$. 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.[5] 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 $\mathcal{S}_M$ 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 $O(|V||E|)$,[16] or the *Gabow-Tarjan* algorithm with complexity $O(\sqrt{|V|}|E| \log |E|)$.[11]

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 6.1 (hyperplane). For the hyperplane X in Example 3.1, 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.

0	i	$n-1$
$\mathbb{Z}$	0	$\mathbb{Z}$

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 6.2 (vee). In Example 3.2, we found that $T(X)$ for the vee-shaped space is a disjoint union of four rays, shown in Figure 4. As each ray is contractible, $H_*(T(X))$ is isomorphic to the homology of four points. Explicitly,

0	i
$\mathbb{Z}^4$	0

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

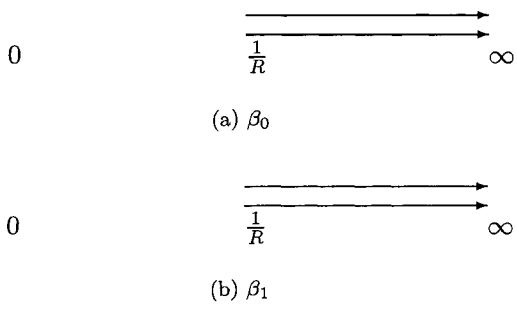


Fig. 14. Barcodes for the circle.

Example 6.3 (boundary of positive octant). We analyzed $T(X)$ of the boundary of the positive octant in Example 3.3. The topology of the space is that of the fiber at the corner point: three circles that intersect pairwise in pairs of antipodal points, as shown in Figure 5(c). We use the Mayer-Vietoris technique to compute the homology.

0	1	i
$\mathbb{Z}$	$\mathbb{Z}^7$	0

We next analyze some cases of the filtered tangent complex. Since our description of persistent homology in terms of barcodes requires homology with field coefficients, we use $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$ as our coefficients set.

Example 6.4 (circle). In Example 4.1, we computed $T^{filt}(X)$ 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 14.

δ	0	1
$< 1/R$	0	0
$\geq 1/R$	$\mathbb{F}^2$	$\mathbb{F}^2$

Example 6.5 (ellipse). In Example 4.2, we computed $T^{filt}(X)$ 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 8. As for Table 1, the lower bound for δ is implied by the value on the preceding row. We draw the corresponding barcodes in Figure 15.

δ	0	1
$< a/b^2$	0	0
$< b/a^2$	$\mathbb{F}^4$	|
$\geq b/a^2$	$\mathbb{F}^2$	$\mathbb{F}^2$

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

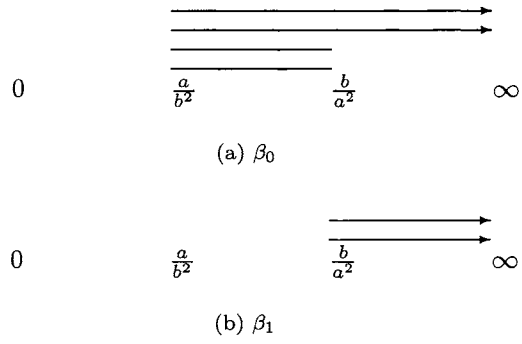


Fig. 15. Barcodes for the ellipse.

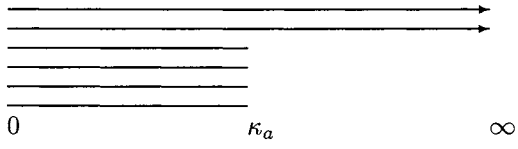


Fig. 16. β_0-barcode for the cubic. There is no β_1-barcode.

Example 6.6 (cubic). In Example 4.3, we computed $T^{filt}(X)$ for the graph of the cubic $f(x) = x^3 + ax$. As X is smooth, the full tangent complex is $X \times S^0$. We saw that $T_\delta(X)$ undergoes three phases, however, as shown in Figure 9. Below, κ_a is the curvature defined in the example. We draw the corresponding barcodes in Figure 16.

δ	0	1
$= 0$	$\mathbb{F}^2$	0
$< \kappa_a$	$\mathbb{F}^6$	
$\geq \kappa_a$	$\mathbb{F}^2$	

Example 6.7 (ellipsoid). In Example 4.4, we computed $T^{filt}(X)$ for the ellipsoid defined by

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1, \tag{53}$$

when $a < b < c$ and $a^2c < b^3 < ac^2$, summarizing the results in Table 1. In this case, the full tangent complex is homeomorphic to the unit tangent bundle to the sphere. This bundle can be expressed as a *fiber bundle* with base space S^2 and fiber S^1, but is not the trivial bundle $S^2 \times S^1$.[14] This fact is reflected in its integer

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

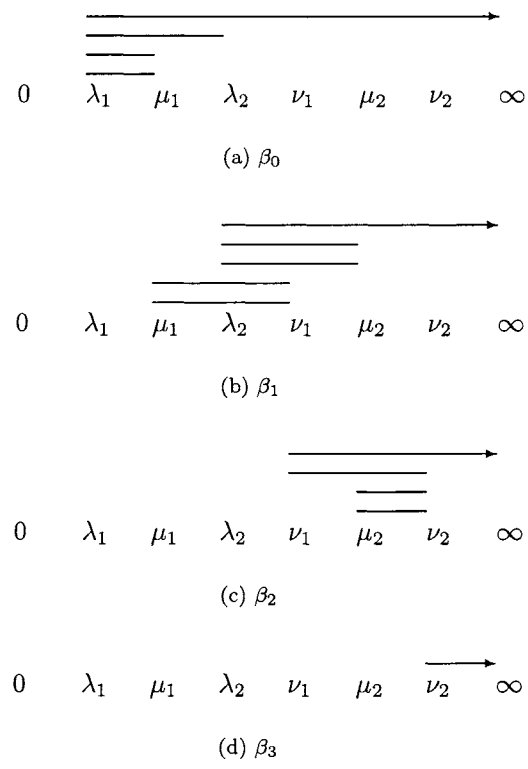


Fig. 17. Barcodes for the ellipsoid.

homology that we do not show here. Using Serre spectral sequences, we have

$$H_i(T(X)) \cong \begin{cases} \mathbb{F}, \, 0 \leq i \leq 3, \\ 0, \text{ otherwise.} \end{cases} \tag{54}$$

Our computations, summarized in Table 1, give us the following description of the homology of $T_\delta(X)$. We draw the non-empty barcodes in Figure 17.

δ	0	1	2	3	i
$< \lambda_1$	0	0	0	0	0
$< \mu_1$	$\mathbb{F}^4$	|	|	|	|
$< \lambda_2$	$\mathbb{F}^2$	$\mathbb{F}^2$	|	|	|
$< \nu_1$	$\mathbb{F}$	$\mathbb{F}^5$	|	|	|
$< \mu_2$	$\mathbb{F}$	$\mathbb{F}^3$	$\mathbb{F}^2$	|	|
$< \nu_2$	$\mathbb{F}$	$\mathbb{F}$	$\mathbb{F}^4$	|	|
$\geq \nu_2$	$\mathbb{F}$	$\mathbb{F}$	$\mathbb{F}$	$\mathbb{F}$	|

Example 6.8 (cone). In Example 4.5, we computed $T^{filt}(X)$ for the cone. We saw that $T_0(X)$ has the homotopy type of $S^1 \times S^0$, and that for every positive δ,

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

Table 2. The Betti numbers for $T^{tame}(X)$ for various geometric surfaces.

surface	β_1	β_2	β_3
sphere	1	1	1
cone	2	2	0
cube	24	0	0
tetrahedron	13	0	0
pyramid	21	0	0
cylinder	4	3	0

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

$T_\delta(X)$ has the homotopy type of $S^1 \times S^1$. We obtained the following description for the homology of T_δ. As the intervals are not particularly interesting, we do not draw them.

δ	0	1	2	i
$= 0$	$\mathbb{F}$	$\mathbb{F}$	0	0
> 0	$\mathbb{F}$	$\mathbb{F}^2$	$\mathbb{F}$	

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, or placed in alternate coordinate systems. We devote the rest of this section to two examples that focus on sharp and soft features, respectively.

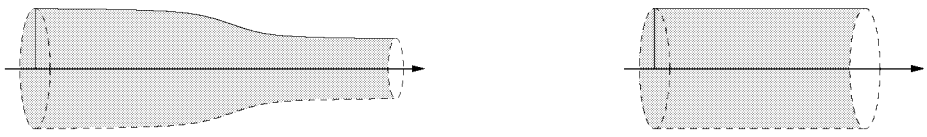
Example 7.1 (geometric surfaces). Suppose that our task is to devise a system that is able to recognize geometric surfaces from the list {sphere, cone, rectangular prism, tetrahedron, pyramid, cylinder}. We would like our system to recognize them even if they are given in spherical, cylindrical, hyperbolic, or other coordinate systems, without any knowledge of the system, and without making any transformations. Also, we would like our system to recognize deformed versions of the objects on the list, and classify them correctly. To do this, we first construct the tame tangent complex $T^{tame}(X)$ for a given object X. Our invariant will be the vector of Betti numbers for this complex. We show this vector for our list objects in Table 2. We compute these values using the Mayer-Vietoris techniques described in Section 2. It is clear that we can distinguish between these objects through the vectors of Betti numbers. This vector is unaffected by smooth deformations of the ambient Euclidean space, and does not depend on the coordinate system used. Note that this vector captures sharp features of the surfaces, and this is enough for differentiating between the objects.

Example 7.2 (bottle vs. glass). In Example 3.6, 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 18. 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 $f\colon [0, H] \to \mathbb{R}$ 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 4.6).

(i) 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.
(ii) κ^{rev} is monotonically increasing.
(iii) κ^{graph} is 0 at the inflection point and on 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 $T_0(X)$ decomposes 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 19(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



(a) bottle (b) glass

Fig. 18. Surfaces of revolution around the x-axis.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

$[\xi_H, H] \times f(H)$. 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 by Example 2.3. 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

$$\pi^{-1}(W_2) \simeq \pi^{-1}(W_3) \simeq S^1 \times S^0. \tag{55}$$

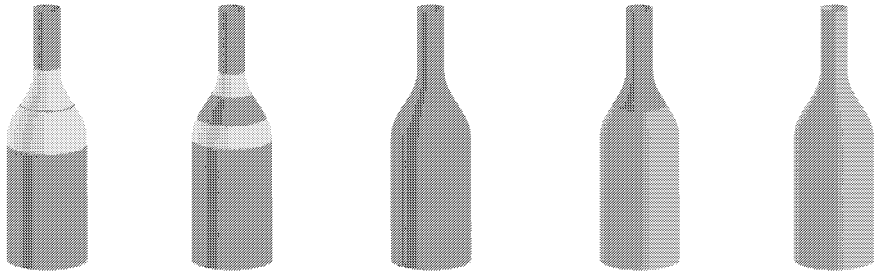
W_1 is precisely the glass we constructed in Example 2.15. So, $\pi^{-1}(W_1)$ is the union of two pieces U and V, where U is the 0-tangent complex of the base disc, and V is the 0-tangent complex of the side. Since the base disc D^2 is flat, each fiber in U is a full circle, and we know that $U \simeq D^2 \times S^1$ by Example 3.1. The side is a cylinder once again, so its 0-tangent complex is $V \simeq S^1 \times S^0$ 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:

$$\pi^{-1}(W_1) \simeq (D^2 \times S^1) \dot{\cup} (S^1 \times S^0). \tag{56}$$

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

$$(D^2 \times S^1) \dot{\cup} (S^1 \times S^0) \dot{\cup} (S^1 \times S^0) \dot{\cup} (S^1 \times S^0). \tag{57}$$

- $0 < \delta < \kappa_+$: $T_\delta(X)$ does not change its homotopy type, while the inflection circle grows into an annulus, as shown in Figure 19(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 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 $T_{\kappa_+}(X) \simeq (D^2 \times S^1) \dot{\cup} (S^1 \times S^0) \dot{\cup} (S^1 \times S^0)$.
- $\kappa_+ < \delta < \kappa_-$: The homotopy type does not change.



(a) $\delta = 0$ (b) $0 < \delta < \kappa_+$ (c) $\kappa_- \leq \delta < \kappa_0$ (d) $\kappa_0 < \delta < \kappa_H$ (e) $\delta \geq \kappa_H$

Fig. 19. 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, $\kappa_- = \kappa_+ < \kappa_0$.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

- $\delta = \kappa_-$: The neck and the body in the projection of $T_\delta(X)$ merge, as shown in Figure 19(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, that now intersect in a pair of antipodal points. We now have

$$T_{\kappa_0}(X) \simeq (D^2 \times S^1) \cup_{S^1 \times S^0} (S^1 \times S^1), \tag{58}$$

  where $\cup_{S^1 \times S^0}$ is the construction in Definition 2.2.
- $\delta \geq \kappa_0$: The homotopy type remains unchanged, as shown in Figures 19(d) and 19(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
$< \kappa_-$	$\mathbb{F}^5$	$\mathbb{F}^5$	|
$< \kappa_0$	$\mathbb{F}^3$	$\mathbb{F}^3$	|
$\geq \kappa_0$	$\mathbb{F}$	$\mathbb{F}^3$	$\mathbb{F}^2$

We next perform the same calculation for a glass shown in Figure 18(b). We assume it is the surface of revolution obtained by revolving $(0, [0, R]) \cup ([0, H'], R)$ around the x-axis, giving us a glass 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 2.15). Let $\kappa = 1/R$. The analysis follows quickly from our observation:

- $\delta = 0$: We have $T_0(X) \simeq (D^2 \times S^1) \dot{\cup} (S^1 \times S^0)$ as in the analysis of $\pi^{-1}(W_1)$ for the bottle.
- $0 < \delta < \kappa$: The homotopy does not change.
- $\delta = \kappa$: We have $T_\kappa(X) \simeq (D^2 \times S^1) \cup_{S^1 \times S^0} (S^1 \times S^1)$ 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, which correspond to the shorter intervals shown in Figure 20 on the left. In the figure, we set $\kappa_0 = \kappa$ to compare a glass and bottle of the same bottom radius.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

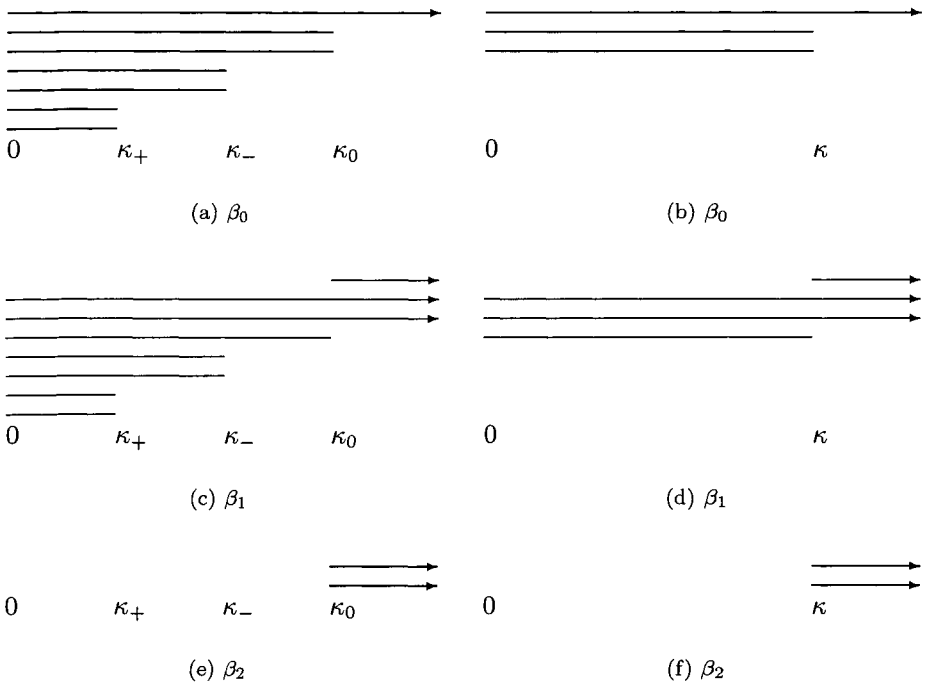


Fig. 20. Barcodes for the bottle (left) and the glass (right).

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

$$2\left(\left|\kappa_+ - \kappa'_+\right| + \left|\kappa_- - \kappa'_-\right| + \left|\kappa_0 - \kappa'_0\right|\right). \tag{59}$$

  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) $2\kappa \geq \kappa_- + \kappa_0$.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

The corresponding distances are:

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

Note that the distances are not relative, so we can reasonably expect a good performance from our distance function in distinguishing between a bottle and a glass. We can also visualize how changes in the parameter values for a 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 a bottle's distance to a glass.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

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 suggestions 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. Already, we have addressed these issues for PCDs of curves.[4] 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.[3]

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.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.

Acknowledgments

The first and third authors were supported in part by NSF under grant DMS 0101364. The second and fourth authors were supported in part by NSF/DARPA under grant CARGO 0138456. A portion of the work was done while the second author was at the Max-Planck-Institut für Informatik, Saarbrücken, Germany, and the third author was at Stanford University.

References

1. M. Berger and B. Gostiaux. *Géométrie différentielle: variétés, courbes et surfaces.* Mathématiques Presses Universitaires de France, Paris, France, 1992.
2. F. Bookstein. *Morphometric Tools for Landmark Data.* Cambridge Univ. Press, Cambridge, UK, 1991.
3. G. Carlsson and V. de Silva. A geometric framework for sparse matrix problems. *Advances in Applied Mathematics*, 33(1):1–25, 2004.
4. A. Collins, A. Zomorodian, G. Carlsson, and L. Guibas. A barcode shape descriptor for curve point cloud data. *Computers and Graphics*, 28:881–894, 2004.
5. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* The MIT Press, Cambridge, MA, 2001.
6. S. K. Donaldson and P. B. Kronheimer. *The Geometry of Four-Manifolds.* The Clarendon Press, New York, NY, 1990.
7. H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.
8. T.-J. Fan. *Describing and Recognizing 3D Objects Using Surface Properties.* Springer-Verlag, New York, NY, 1990.
9. H. Federer. *Geometric Measure Theory*, volume 153 of *Die Grundlehren der Mathematischen Wissenschaften.* Springer-Verlag, New York, NY, 1969.
10. R. B. Fisher. *From Surfaces to Objects: Computer Vision and Three-Dimensional Scene Analysis.* John Wiley and Sons, Inc., New York, NY, 1989.
11. H. N. Gabow and R. E. Tarjan. Faster scaling algorithm for network problems. *SIAM J. Comput.*, 18:1013–1036, 1989.
12. M. J. Greenberg and J. R. Harper. *Algebraic Topology: A First Course*, volume 58 of *Mathematics Lecture Note Series.* Benjamin/Cummings Publishing Co., Reading, MA, 1981.
13. A. Hatcher. *Algebraic Topology.* Cambridge Univ. Press, Cambridge, UK, 2001.
14. D. Husemoller. *Fiber Bundles*, volume 20 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, NY, 1994.

15. D. Kendall, D. Barden, T. Carne, and H. Le. *Shape and Shape Theory.* John Wiley and Sons, Inc., New York, NY, 1999.
16. H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
17. S. Lang. *Algebra*, volume 211 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, NY, third edition, 2002.
18. A. B. Lee, K. S. Pedersen, and D. Mumford. The non-linear statistics of high contrast patches in natural images. Technical report, Brown University, 2001. Available online.
19. J. Milnor. *Morse Theory*, volume 51 of *Annals of Mathematical Studies.* Princeton Univ. Press, Princeton, NJ, 1963.
20. J. Milnor and J. D. Stasheff. *Characteristic Classes*, volume 76 of *Annals of Mathematical Studies.* Princeton Univ. Press, Princeton, NJ, 1974.
21. S. Zhu and A. Yuille. Forms: A flexible object recognition and modeling system. *Int. J. Comp. Vision*, 20(3):187–212, 1996.
22. A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.

Int. J. Shape Model. 2005.11:149-187. Downloaded from www.worldscientific.com
by SIMON FRASER UNIVERSITY on 03/14/15. For personal use only.