Topology Proceedings

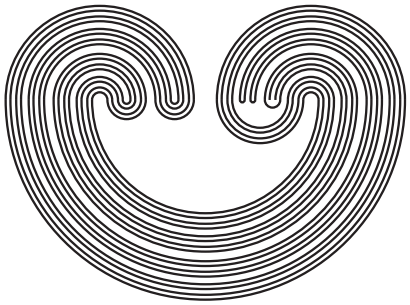

Web: http://topology.auburn.edu/tp/
Mail: Topology Proceedings
Department of Mathematics & Statistics
Auburn University, Alabama 36849, USA
E-mail: topolog@auburn.edu
ISSN: 0146-4124

Topology Proceedings
Volume 24, Summer 1999, 503–532

TOWARDS COMPUTING HOMOLOGY FROM FINITE APPROXIMATIONS

V. Robins

Abstract

We consider the problem of extrapolating the homology of a compact metric space from a finite point-set approximation. Our approach is based on inverse systems of ϵ-neighborhoods and inclusion maps. We show that the inclusion maps are necessary to identify topological features in an ϵ-neighborhood that persist in the limit as $\epsilon \to 0$. We outline a possible algorithm for computer implementation. As test examples, we present data for some iterated function system relatives of the Sierpinski triangle.

1. Introduction and Basic Definitions

A fundamental task in the study of dynamical systems is to identify the topological and geometric structure of attractors and other invariant sets. Such information can contribute to our understanding of global properties of the dynamics. An example where topological structure is important comes from nearly-integrable area-preserving twist maps of the cylinder. In these maps, invariant circles trap chaotic orbits and imply some degree of stability that is destroyed when the nonlinearity is increased and the invariant circles become Cantor sets [13].

Mathematics Subject Classification: Primary: 55-04; Secondary: 58F12, 28A80.

Key words: Computational homology, Betti numbers, fractal geometry, shape theory.

The geometry of an attractor, as measured by its fractal dimension, is related to its Lyapunov exponents, which measure the exponential rate of divergence of nearby orbits [17]. The various fractal dimensions are the most popular tools for quantifying geometric structure. However, the family of examples in Section 4 show that such information says nothing about topological structure. Topological information, such as whether or not an attractor is connected or how many and what type of holes it has, is more fundamental but more difficult to extract from data. Previous applications of computational topology in the context of dynamical systems include the computation of homology from flows that lie on smooth manifolds, [15], and the application of Conley index theory to time-series data, [14].

The standard computer representation of an attractor is a finite set of points generated along the orbit of a typical initial point. Our goal, therefore, is to develop computational techniques that allow us to extrapolate topological properties of the underlying attractor, given only the finite approximation. The basic trick is to coarse-grain the data at a sequence of resolutions that tend to zero. Of course, the extrapolation will always be constrained by the accuracy of the data and we address this problem by identifying a cutoff resolution.

This multiscale approach is applied to computing connectedness in [19, 20]. In the current paper, we investigate the more difficult problem of deducing the homology of a compact metric space from a finite amount of data. We derive sound mathematical foundations for the heuristic of the previous paragraph, using ideas from shape theory [12]. Although the applications we describe are in dynamical systems, our approach is quite general and is valid for any situation where a finite set of points is viewed as an approximation to an underlying space. A particular advantage over previous work on computational homology is that it allows for fractal data.

The abstract setting for our analysis is as follows. We assume that the underlying space, X, is a compact subset of a metric space, (M, d) and that the finite set of points, $S \subset M$, approximates X in a metric sense, i.e., each point of X is within distance ρ of some point in S, and vice versa. In other words, ρ is the Hausdorff distance between X and S. A small value of ρ implies S is a good approximation to X. This assumption holds, for example, in the situation where X is a compact attractor for a diffeomorphism $f : M \to M$. In this case, for each point in the basin of attraction, $x \in B$, the ω-limit set of x is contained in X [21]. Thus, given $\rho > 0$, and a point $x \in B$, there is an integer n, such that for $j > n$, $d(f^j(x), X) < \rho$. If we take S to be the finite number of iterates, $S = \{f^j(x) \text{ for } n < j < N\}$, then each point of S is within ρ of X. Conversely, if the attractor has an ergodic measure, then a typical orbit (in a measure-theoretic sense) must fill out the entire attractor. Given ρ, we can therefore choose N large enough so that every point of X is within ρ of S. In other contexts, ρ could represent a discretization error or the magnitude of noise present in the data.

To give the compact space, X, and its finite approximation, S, comparable topological structure, we form their closed ϵ-neighborhoods:

$$X_\epsilon = \{x \in M \mid d(x, X) \leq \epsilon\} \quad \text{and} \quad S_\epsilon = \{x \in M \mid d(x, S) \leq \epsilon\}.$$

Roughly speaking, since X and S are within ρ, their ϵ-neighborhoods should have similar properties for $\epsilon > \rho$. The homology of X_ϵ converges to the homology of X as $\epsilon \to 0$, in an inverse limit sense. Therefore, if ρ is small enough, we hope to have some confidence in extrapolating this limit from the structure of S_ϵ for $\epsilon > \rho$. We formalize these ideas in Section 2 using the machinery of inverse systems and some results from shape theory.

There is a second advantage to using ϵ-neighborhoods. The standard simplicial homology groups are only defined for spaces that have a finite triangulation. The attractors of chaotic dynamical systems frequently have fractal structure, and are not homeomorphic to a finite polytope. However, an ϵ-neighborhood of a compact metric space is an absolute neighborhood retract (ANR) and is therefore homotopy equivalent to a finite polytope [12, p.45]. This means simplicial homology groups are well-defined and effectively computable for the ϵ-neighborhoods. The appropriate homology theory for fractal sets is Čech homology, which is defined in terms of an inverse system of nerves of covers [9].

For notational purposes, we recall some definitions for simplicial homology, following Munkres [16]. The basic building block is an *oriented k-simplex*, σ^k —the convex hull of $k+1$ geometrically independent points, $\{x_0, x_1, \ldots, x_k\}$. The orientation is defined by an arbitrary but fixed ordering of the vertices. For example, a 0-simplex is just a point, a 1-simplex is a line segment, a 2-simplex a triangle, and a 3-simplex is a tetrahedron. A *simplicial complex*, $\mathcal{C}$, is a collection of oriented simplices with the property that the non-empty intersection of two simplices in $\mathcal{C}$ must itself be a simplex in $\mathcal{C}$. The union of simplices from $\mathcal{C}$, when viewed as a subset of $\mathbb{R}^n$, is called a *polytope*. If X is homeomorphic to a polytope, we say X is *triangulated* by $\mathcal{C}$.

We now define the group structures associated with a space, X, triangulated by the simplicial complex, $\mathcal{C}$. A *k-chain* is the formal sum of a finite number of oriented k-simplices: $c_k = \sum_i a_i \sigma_i^k$. The coefficients a_i are typically integers, but in general they can be elements of any abelian group, G. The resulting chain group is denoted $C_k(X;G)$, or simply C_k when the space and the coefficient group are understood. The chain group is free abelian and the oriented k-simplices form a basis. For $k \geq 1$, the *boundary operator*, $\partial_k : C_k \to C_{k-1}$, is a linear operator that maps a k-simplex onto the oriented sum of $(k-1)$-simplices in its boundary (e.g., the boundary of a triangle is its three edges).

The image of ∂_k is a subgroup of C_{k-1} and is denoted B_{k-1}. The kernel of ∂_k is the group of k-cycles, Z_k. The fundamental property of the boundary operator is that $\partial_k \partial_{k+1} = 0$. This implies that B_k is a subgroup of Z_k, so we can form their quotient group, which is precisely the *homology group*, $H_k = Z_k/B_k$. Its elements are equivalence classes of non-bounding k-cycles, loosely referred to here as k-dimensional holes. The fundamental theorem of finitely generated abelian groups implies that H_k is isomorphic to the direct sum of a free group and a finite number of finite cyclic groups. The rank of the free group is called the kth Betti number β_k and the orders of the finite cyclic groups are the torsion coefficients. If the coefficient group is the rational $\mathbb{Q}$ or real $\mathbb{R}$ numbers then there is no torsion and the Betti numbers are simply the dimensions of the homology groups. The kth Betti number effectively counts the number of k-dimensional holes in X, so is exactly the information we seek. When $k = 0$, the Betti number counts the number of path-connected components of X.

In Section 2 we give definitions for the inverse system of ϵ-neighborhoods of X and the corresponding inverse systems of homology groups. We would like to quantify the topological structure by computing Betti numbers as functions of ϵ. There is a problem with this however, since ϵ-neighborhoods can have holes that do not exist in X. We resolve this problem by introducing the concept of a persistent Betti number, which counts the number of holes in X_ϵ that correspond to a hole in X. When X has fractal structure, it is possible to see unbounded growth in the persistent Betti numbers as $\epsilon \to 0$. We characterize this growth by assuming an asymptotic power law. The final step of this section is to obtain formal relationships between the ϵ-neighborhoods of X and a finite point-set approximation.

In Section 3, we turn to the practical problem of how to implement these ideas computationally. There are a number of existing approaches for computing the Betti numbers of a given simplicial complex. We give a brief overview of an incremental algorithm due to Delfinado and Edelsbrunner [3] for computing

the Betti numbers of a sequence of simplicial complexes in $\mathbb{R}^2$ or $\mathbb{R}^3$. We also show that, to a finite resolution, the persistent Betti numbers are computable from dimensions of intersections of ranges and null spaces of matrices, and outline an algorithm to implement the inverse limit system approach.

We use a family of iterated function system attractors as simple test examples in Section 4. These fractals illustrate a problem with the use of the Hausdorff dimension to characterize structure. Each of these examples has exactly the same similarity (and therefore Hausdorff) dimension, and yet they have dramatically different topological structure. Since we have not yet implemented algorithms for computing the persistent Betti numbers, we give data for the regular Betti numbers and show that this can be misleading.

2. Mathematical Foundations

In this section, we use an inverse system framework from shape theory to clarify the sense in which the ϵ-neighborhoods of S relate to those of X. We introduce the notion of persistent Betti number to quantify the topological structures of ϵ-neighborhoods that persist in the limit as $\epsilon \to 0$.

2.1. The inverse system of ϵ-neighborhoods

To describe the limiting behavior of ϵ-neighborhoods, we use an inverse system. Recall that an *inverse system* of topological spaces consists of a collection of spaces, X_λ, indexed by a directed set $(\Lambda, \succeq)$, and continuous maps, called *bonding morphisms*,

$$p_{\lambda\mu} : X_\mu \to X_\lambda, \quad \text{for each pair} \quad \mu \succeq \lambda.$$

The bonding morphisms must satisfy the following two conditions (note that throughout this paper the composition of two functions, f and g, is written as fg).

$$p_{\lambda\lambda} = 1_{X_\lambda}, \qquad \text{the identity map on } X_\lambda, \text{ and} \tag{1}$$

$$p_{\lambda\mu}p_{\mu\nu} = p_{\lambda\nu}, \qquad \text{for any choice of } \nu \succeq \mu \succeq \lambda. \tag{2}$$

We write $\mathbf{X} = (X_\lambda, p_{\lambda\mu}, \Lambda)$ or $(X_\lambda, p_{\lambda\mu})$ when the index space is clear from context. The *inverse limit space*, $\varprojlim \mathbf{X}$, is the subspace of $\Pi_\Lambda X_\lambda$ consisting of all threads in $\mathbf{X}$:

$$\varprojlim \mathbf{X} = \{(x_\lambda) \mid x_\lambda \in X_\lambda \text{ and } p_{\lambda\mu}(x_\mu) = x_\lambda \text{ for } \mu \succeq \lambda\}.$$

The *projections*, $p_\mu : \varprojlim \mathbf{X} \to X_\mu$, are the continuous maps $p_\mu((x_\lambda)) = x_\mu$.

The concept of inverse system also holds in the category of groups. The definition is similar: simply replace the topological spaces by groups, the continuous maps by homomorphisms and the direct product by direct sum.

The inverse system of ϵ-neighborhoods has terms

$$X_\epsilon = \{x \mid d(x, X) \leq \epsilon\} \quad \text{for} \quad 0 < \epsilon \leq \epsilon_0,$$

with the order relation $\lambda \succeq \epsilon$ if $\lambda < \epsilon$ as real numbers. Since $X_\lambda \subset X_\epsilon$ when $\lambda < \epsilon$, the bonding morphisms are simply inclusion maps $p_{\epsilon\lambda} : X_\lambda \to X_\epsilon$, which clearly satisfy the two conditions above. The inverse limit space is homeomorphic to X and the projections $p_\epsilon : X \to X_\epsilon$ are again inclusion maps.

Each ϵ-neighborhood has the homotopy type of a finite polyhedron, so the kth homology groups $H_k(X_\epsilon)$ are well-defined for $\epsilon > 0$. The inclusion maps $p_{\epsilon\lambda} : X_\lambda \to X_\epsilon$ induce homomorphisms on the homology groups in the standard way [16]. We write $p_{\epsilon\lambda *} : H_k(X_\lambda) \to H_k(X_\epsilon)$ for these induced homomorphisms. The homology groups together with the inclusion-induced homomorphisms yield inverse systems of groups, denoted by $H_k(\mathbf{X})$. The inverse limit of $H_k(\mathbf{X})$ is isomorphic to the kth Čech homology group $\check{H}_k(X)$; for details see [12, p.121]. When X has the homotopy type of a finite polyhedron, the Čech homology groups are isomorphic to the finite simplicial homology groups. If X has the infinitely detailed structure of a fractal, some of its Čech homology groups may involve infinite products of the coefficient group. Note that in Čech homology, the choice of coefficient group is restricted to a field or compact topological group, such as $\mathbb{Q}$ or $\mathbb{R}$, [9].

2.2. Persistent Betti numbers

We would like to quantify the structure of X by looking at the Betti numbers $\beta_k(X_\epsilon) = \operatorname{rank} H_k(X_\epsilon)$ as $\epsilon \to 0$. In general though, it is not the case that $\beta_k(X_\epsilon) \to \beta_k(X)$ as $\epsilon \to 0$. As an example, consider Antoine's necklace [9]. This Cantor set, A, is constructed by taking the intersection of a sequence of nested and linked solid tori. The nth term in the sequence, A_n, consists of 4^n tori, hence $\beta_1(A_n) = 4^n \to \infty$. However, since A is a Cantor set, $\beta_1(A) = 0$. The problem stems from ignoring the role of the bonding morphisms. We now describe how to incorporate this information.

We want to detect holes that have preimages under the bonding morphisms. More formally, for $\lambda < \epsilon$, we say that an equivalence class of cycles $[z_\epsilon] \in H_k(X_\epsilon)$ *persists in* $H_k(X_\lambda)$ if it is in the image of the bonding homomorphism $[z_\epsilon] \in p_{\epsilon\lambda *}(H_k(X_\lambda))$. The number of holes in X_ϵ that persist in X_λ is therefore just the rank of the image subgroup:

$$\beta_k^\lambda(X_\epsilon) = \operatorname{rank}(p_{\epsilon\lambda *}(H_k(X_\lambda))). \tag{3}$$

Since we want to know the topology of X, we are also interested in the quantity

$$\beta_k^0(X_\epsilon) = \operatorname{rank}(p_{\epsilon *}(H_k(X))). \tag{4}$$

We refer to $\beta_k^\lambda(X_\epsilon)$ for $\lambda \geq 0$, as the *persistent Betti numbers*.

The persistent Betti number is an integer-valued function of two real numbers, $\lambda < \epsilon$. In order to understand the properties of this function, we give some elementary bounds on $\beta_k^\lambda(X_\epsilon)$. From the definition, it follows that the persistent Betti number is less than the regular Betti number for both X_λ and X_ϵ:

$$\beta_k^\lambda(X_\epsilon) \leq \beta_k(X_\epsilon) \qquad \text{and} \tag{5}$$

$$\beta_k^\lambda(X_\epsilon) \leq \beta_k(X_\lambda) \qquad \text{for } \lambda \geq 0. \tag{6}$$

The next two inequalities say that for a fixed ϵ-neighborhood, $\beta_k^\lambda(X_\epsilon)$ is a monotonic non-decreasing function of λ; while for a fixed λ, the persistent Betti number is a non-increasing function of ϵ.

$$\text{For} \quad \nu < \lambda < \epsilon, \quad \beta_k^\nu(X_\epsilon) \leq \beta_k^\lambda(X_\epsilon). \tag{7}$$

$$\text{For} \quad \lambda < \epsilon < \mu, \quad \beta_k^\lambda(X_\mu) \leq \beta_k^\lambda(X_\epsilon). \tag{8}$$

Roughly speaking, increasing the difference between λ and ϵ decreases the number of persistent holes.

The above inequalities are a first step towards understanding the continuity properties of $\beta_k^\lambda(X_\epsilon)$, but more work remains to be done. Most importantly, we want conditions on X that guarantee

$$\beta_k^0(X_\epsilon) \to \beta_k(X) \quad \text{as} \quad \epsilon \to 0. \tag{9}$$

If $\beta_k(X)$ is finite and (9) holds, then there must be an $\epsilon_0 > 0$ such that $\beta_k^0(X_\epsilon) = \beta_k(X)$ for $\epsilon \in [0, \epsilon_0)$. In addition to this, our computational work will be most effective for spaces where there is a λ_0 such that $\beta_k^\lambda(X_\epsilon) = \beta_k^0(X_\epsilon)$ for $\lambda \in [0, \lambda_0]$. We do not expect these conditions to hold for an arbitrary compact space. However, for special cases such as the iterated function system attractors in Section 4, it may be possible to say something more concrete about the continuity of persistent Betti numbers as λ and ϵ tend to zero.

For fractal sets, it is possible to see unbounded growth in the persistent Betti numbers as ϵ decreases to zero. If $\beta_k^0(X_\epsilon) \to \infty$ as $\epsilon \to 0$, we can quantify the rate of divergence by assuming an asymptotic power law, $\beta_k^0(X_\epsilon) \sim \epsilon^{\gamma_k}$. The exponent γ_k can be computed as the following limit when it exists (and as the limsup or liminf otherwise):

$$\gamma_k = \lim_{\epsilon \to 0} \frac{\log(\beta_k^0(X_\epsilon))}{\log(1/\epsilon)}. \tag{10}$$

For $k = 0$, the Betti number is just the number of path-connected components and the persistent Betti number is the same as the regular Betti number: $\beta_0^\lambda(X_\epsilon) = \beta_0(X_\epsilon)$. The exponent γ_0 is therefore the disconnectedness index of [19, 20]. For self-similar Cantor sets, γ_0 is equivalent to the similarity dimension. The examples in Section 4 suggest that an analogous result may hold for connected self-similar fractals that contain holes.

2.3. Finite Approximations

We now turn to analyzing the finite set of points, S, that approximate the compact set X in the metric space (M, d). As described in the introduction, we let ρ be the Hausdorff distance $d_H(X, S) = \rho$. In this section, we show how to relate the inverse systems of ϵ-neighborhoods, $(X_\epsilon, p_{\lambda\epsilon})$ and $(S_\epsilon, q_{\lambda\epsilon})$, and derive bounds on the persistent Betti numbers of X_ϵ.

Since $X \subset S_\rho$ and $S \subset X_\rho$, we have the following inclusion maps for any $\epsilon > 0$, $i_\epsilon : S_\epsilon \to X_{\epsilon+\rho}$ and $j_\epsilon : X_\epsilon \to S_{\epsilon+\rho}$. With appropriate bonding morphisms, this gives us a commuting diagram for any $\lambda \geq \rho$ and $\epsilon \geq \lambda + 2\rho$:

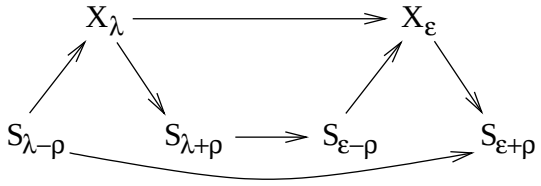

(11)

This diagram implies that if an element of $H_k(X_\epsilon)$ has a preimage in $H_k(X_\lambda)$, there must be an element of $H_k(S_{\epsilon-\rho})$ with a preimage in $H_k(S_{\lambda+\rho})$. Similarly, if $S_{\epsilon+\rho}$ has a hole that persists in $S_{\lambda-\rho}$, there must be a corresponding hole in X_ϵ that persists in X_λ. In terms of Betti numbers, we have that for $\lambda \geq \rho$ and $\epsilon \geq \lambda + 2\rho$,

$$\beta_k^{\lambda-\rho}(S_{\epsilon+\rho}) \leq \beta_k^\lambda(X_\epsilon) \leq \beta_k^{\lambda+\rho}(S_{\epsilon-\rho}). \tag{12}$$

We can swap the roles of X and S in (11) to obtain analogous bounds on $\beta_k^\lambda(S_\epsilon)$:

$$\beta_k^{\lambda-\rho}(X_{\epsilon+\rho}) \leq \beta_k^\lambda(S_\epsilon) \leq \beta_k^{\lambda+\rho}(X_{\epsilon-\rho}). \tag{13}$$

Since X maps into S_ρ, we can only hope to get information about X from holes in S_ϵ that persist in S_ρ. Setting $\lambda = \rho$ in (13), we have that for $\epsilon \geq \rho$,

$$\beta_k^0(X_{\epsilon+\rho}) \leq \beta_k^\rho(S_\epsilon). \tag{14}$$

Thus, if $\beta_k^0(X_\epsilon) \sim \epsilon^{\gamma_k}$ as $\epsilon \to 0$ and ρ is small enough, we should see at least that order of growth in $\beta_k^\rho(S_\epsilon)$.

In one sense, ρ is the optimal resolution for coarse-graining the data to estimate the topological structure of the underlying space. This could imply that $\beta_k(S_\rho)$ is the best approximation to $\beta_k(X)$, which would render our multiresolution approach redundant. For very simple spaces this may be the case. However, we are interested in more complicated settings. In general, the inverse system of ϵ-neighborhoods and the persistent Betti numbers offer two advantages. First, the cutoff resolution ρ is typically not known in advance and must be determined from the data; examining S_ϵ at many ϵ-values helps us estimate ρ. Second, computing the Betti numbers $\beta_k(S_\rho)$ at a single resolution does not distinguish between holes due to the topology of X and holes in S_ρ induced by the geometry. The sequence of persistent Betti numbers $\beta_k^\rho(S_\epsilon)$ for $\epsilon > \rho$ give a more accurate basis than $\beta_k(S_\rho)$ from which to extrapolate topological information about X (c.f., the example of Antoine's necklace). Such an extrapolation must always be given with respect to the cutoff resolution ρ, however, since it is possible that the topological properties of X change at resolutions below ρ.

3. Computation

We now address the implementation of these ideas. There are four parts to the overall process:

1. For a sequence of ϵ-values, generate simplicial complexes that triangulate the ϵ-neighborhoods of the data.
2. Estimate the value of the cutoff resolution ρ.
3. Compute the persistent Betti numbers, $\beta_k^\rho(\epsilon)$, for $\epsilon > \rho$.
4. If appropriate, compute the growth rate, γ_k.

This approach is implemented in [20] for the case $k = 0$, when $\beta_0^\rho(\epsilon) = \beta_0(\epsilon)$ is just the number of connected components of S_ϵ. For this case, a simplicial complex is unnecessary — all the information about ϵ-connected components of S is encoded in the Euclidean minimal spanning tree of the data.

In the present context, steps 1 and 3 are the most computationally intensive. We describe algorithms for their implementation below. We use a criterion for Step 2 that we derive in [20] for approximations to perfect spaces. Since a perfect space has no isolated points, we estimate the cutoff resolution as the largest value of ϵ for which S_ϵ has at least one isolated point. This underestimates the value for which $S_\rho \supset X$, but the examples in Section 4 show it to be a reasonably good approximation. Isolated points are straightforward to detect numerically [20] — a point is ϵ-isolated if the distance from it to every other point in the set is greater than ϵ. We write $I(\epsilon)$ for the number of ϵ-isolated points. The computation of growth rates in Step 4 is straightforward once the persistent Betti numbers are found.

3.1. Simplicial Complexes for ϵ-neighborhoods

Given an ϵ-neighborhood, S_ϵ, the first problem is to generate a simplicial complex, $\mathcal{C}_\epsilon$, whose underlying space is at least homotopy equivalent to S_ϵ. Since we are interested in the inverse system of ϵ-neighborhoods, we need simplicial complexes for a

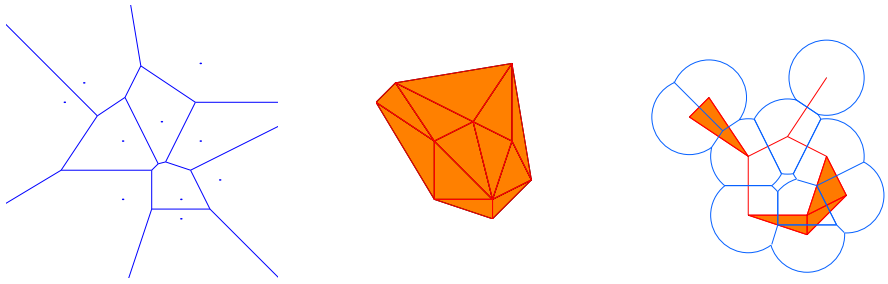

(a) (b) (c)

Fig. 1. (a) The Voronoi diagram of a set of ten points in the plane. (b) The corresponding Delaunay triangulation. (c) An ϵ-neighborhood and corresponding ϵ-complex.

sequence of numbers $\epsilon_i \to 0$. In order to have inclusion maps that are well defined, we need $\mathcal{C}_{\epsilon_i}$ to be either a subcomplex or a subdivision of $\mathcal{C}_{\epsilon_j}$ when $\epsilon_i < \epsilon_j$.

One approach to this problem, due to Edelsbrunner *et al.* [6], is to use *alpha shapes*. This construction makes use of two fundamental graphs from computational geometry — the *Voronoi diagram* and the *Delaunay triangulation*; see Figure 1 for an example. For a finite set of points $S \subset \mathbb{R}^d$, the Voronoi diagram associates a closed cell with each point. The Voronoi cell of $p \in S$ is the set of all points in $\mathbb{R}^d$ that are closer to p than to any other point in S, i.e.,

$$V(p) = \{x \mid d(x,p) \leq d(x,s), s \in S \backslash \{p\}\}.$$

Two cells can intersect only at a common boundary; their interiors are disjoint; and the union of all Voronoi cells is the whole of $\mathbb{R}^d$. The Delaunay triangulation is the geometric dual of the Voronoi diagram. When the points of S are in general position (i.e., no $d+2$ points lie on a d-sphere) the Delaunay triangulation is a simplicial complex whose underlying space is the convex hull of S. To make an analogy with the Čech

construction, the Voronoi diagram is a cover of $\mathbb{R}^d$ and (for points in general position) the Delaunay triangulation is the nerve of this cover.

The resolution parameter, ϵ, is introduced to this setting by taking the intersection of the Voronoi cells with S_ϵ, the closed ϵ-neighborhood. This gives a closed cover of S_ϵ and the nerve of this cover is a subset of the Delaunay triangulation, T_ϵ. The underlying space of an ϵ-complex is an "alpha shape" (our ϵ is the same as Edelsbrunner's α). Note that the ϵ-complex is homotopy equivalent to S_ϵ [5]. In [6], there are rules for determining the subset, T_ϵ, directly from the simplices in the Delaunay triangulation, without constructing the Voronoi cells or S_ϵ. Since the complexes T_ϵ are subsets of a finite simplicial complex, this implies that there are only a finite number of homologically distinct ϵ-neighborhoods.

Software that constructs these complexes for subsets of $\mathbb{R}^2$ or $\mathbb{R}^3$ is publicly available from NCSA [23]. We use these programs to generate the figures and compute Betti numbers for the examples of Section 4. If $S \subset \mathbb{R}^2$, the ϵ-complexes are built with a time cost that is subquadratic in the number of points of S.

For the type of data we are interested in, there is a large degree of redundancy in the Delaunay triangulations (i.e. an unnecessarily large number of simplices) for the resolutions $\epsilon \geq \rho$. This is likely to be the case whenever we construct $\mathcal{C}_{\epsilon_i}$ as a subcomplex of $\mathcal{C}_{\epsilon_j}$ for $\epsilon_i < \epsilon_j$. An alternative approach is to use subdivisions, so that coarse resolutions have fewer simplices than fine resolutions. A subdivision approach might be more efficient in terms of space (i.e., storing the lists of simplices) and computation (of persistent Betti numbers for example) than the ϵ-subcomplexes described above. We are still working to determine the most efficient implementation for our applications.

3.2. Computing Persistent Betti Numbers

Before discussing persistent Betti numbers, we give the standard formula for computing the regular Betti numbers [16]:

$$\beta_k = \operatorname{rank} H_k = \operatorname{rank} Z_k - \operatorname{rank} B_k.$$

Recall that Z_k is the kernel of the boundary operator, ∂_k, and B_k is the image of ∂_{k+1}. Since ∂_k is a linear operator and the oriented k-simplices form a basis for the chain group C_k, it is possible to represent ∂_k as a matrix, D_k, with entries in $\{-1, 0, 1\}$. Then the dimension of the null space of D_k is the rank of Z_k and the dimension of the range of D_{k+1} is the rank of B_k. The computation of Betti numbers is thus reduced to elementary numerical linear algebra. Although this algorithm is simple, for large simplicial complexes such as those for the examples in Section 4, the computations take a significant amount of computer time and are susceptible to roundoff errors. There are a number of faster algorithms for computing Betti numbers from a given complex; we describe some of these after deriving a formula for the persistent Betti number.

Recall from Section 2.3 that we want to compute ρ-persistent Betti numbers for ϵ-neighborhoods of the finite approximation, S_ϵ:

$$\beta_k^\rho(\epsilon) = \operatorname{rank} q_{\epsilon\rho *}(H_k(S_\rho)).$$

We start by simplifying notation a little and write q_* for the homomorphism of homology groups, $q_* : H(\rho) \to H(\epsilon)$. From elementary theorems about homomorphisms of groups, we know that

$$q_*(H(\rho)) \simeq H(\rho)/\ker q_*.$$

The kernel of q_* consists of equivalence classes of cycles in $H(\rho)$ that map to boundaries in $B(\epsilon)$. Let $q_\sharp : C(\rho) \to C(\epsilon)$ be the map on the chain complexes induced by the inclusion map

on the ϵ-neighborhoods. The homomorphism $q_\sharp$ commutes with the boundary operator [16], so $q_\sharp(B(\rho)) \subset B(\epsilon)$ and $q_\sharp(Z(\rho)) \subset Z(\epsilon)$. It follows that

$$q_*(H(\rho)) \simeq Z(\rho)/[Z(\rho) \cap q_\sharp^{-1}(B(\epsilon))].$$

The persistent Betti number is therefore

$$\beta_k^\rho(\epsilon) = \text{rank}[Z(\rho)] - \text{rank}[Z(\rho) \cap q_\sharp^{-1}(B(\epsilon))]. \qquad (15)$$

For computation, we can again use linear algebra. The first term is the dimension of the null space of $D_k(\rho)$ as before. For the second term we need to find the intersection of two spaces: $\text{null}\, D_k(\rho)$ and $q_\sharp^{-1}[\text{range}\, D_{k+1}(\epsilon)]$. Finding the intersection of two linear subspaces requires some information about their bases, and is therefore a more difficult problem than computing dimensions. There are standard algorithms for this [8], but we have not yet implemented them. This is because the implementation depends on how we construct the sequence of simplicial (or cubical) complexes for the ϵ-neighborhoods and this is work in progress. Instead, for the examples of Section 4 we use pre-existing software for computing the regular Betti numbers from subsets of the Delaunay triangulation described in Section 3.1. These examples illustrate why the regular Betti numbers are insufficient for extrapolation.

The algorithm for computing Betti numbers of the ϵ-subcomplexes of the Delaunay triangulation is due to Delfinado and Edelsbrunner [3]. Their algorithm is an incremental one that uses central results from homology theory such as the Meyer-Vietoris sequence and Poincaré duality. For subsets of $\mathbb{R}^2$ and $\mathbb{R}^3$ only, the running time is $O(n\alpha(n))$, where n is the number of simplices in the triangulation and $\alpha(n)$ is the extremely slow growing inverse of Ackermann's function [2]. Since the algorithm is incremental, given the Delaunay triangulation as input, it automatically computes Betti numbers of all the subcomplexes, T_ϵ, in a single run.

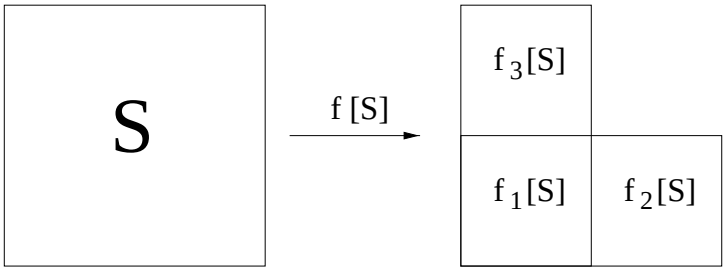

Fig. 2. Template for the iterated function system that generates the Sierpinski triangle relatives.

The Delfinado/Edelsbrunner implementation goes a long way towards carrying out our desired program. However, it is not clear how to incorporate persistent Betti number computations easily. It is possible that an incremental algorithm for finding the persistent Betti numbers exists. However, finding the persistent Betti numbers requires some explicit information about the cycles and boundaries, and the Delfinado/Edelsbrunner algorithm does not generate or record this information.

Finally, we note that the development of fast algorithms for computational homology is an active area of research and refer the interested reader to some recent papers [4, 7, 11].

4. Examples

The fractals we use here are attracting fixed sets of iterated function systems (the IFS framework and a proof of the existence of such fixed sets was first published in Hutchinson [10]). The Sierpinski triangle relatives are generated by a family of iterated function systems, illustrated by the template in Figure 2:

$$S = f[S] = f_1[S] \cup f_2[S] \cup f_3[S].$$

In each case, the functions f_i are similarity transformations of the unit square with contraction ratio $\frac{1}{2}$, i.e. $|f_i(x) - f_i(y)| = \frac{1}{2}|x - y|$. The functions that generate the Sierpinski triangle, shown in Figure 3, are simple contractions composed with a translation; the generators of the other examples involve additional rotation or reflection symmetries of the square. There are 232 different fractals in this family [18]. Their topology ranges from simply connected (Figure 7) to connected (Figure 3) to totally disconnected (Figure 5) to a class of examples with infinitely many connected components of non-zero diameter (Figure 9). This range of topological structure makes them ideal test cases for our techniques.

It is easy to generate a finite number of points on the attractor of an iterated function system. One way (Barnsley's "chaos game" [1]) is to chose an initial point x_0 in the domain of the IFS and then record its trajectory under the iteration $x_{n+1} = f_{i_n}(x_n)$, setting $i_n = 1, 2$ or 3 with probability p_1, p_2 and p_3 respectively. If x_0 is in the attractor then its entire orbit is in the attractor; if not, the iterates converge to it. Thus, the orbit can be viewed as a random sampling of the attractor by a finite number of points. When $p_1 = p_2 = p_3 = \frac{1}{3}$, the data cover the fractal uniformly; if the probabilities are not equal, the distribution of points is nonuniform and their density approximates a multifractal measure. See [20] for a discussion of the way in which the cutoff resolution, ρ, depends on the number of points and the choice of p_1, p_2 and p_3. In the examples below, we choose an initial point in each attractor, and generate 10^4 points with $p_1 = p_2 = p_3 = \frac{1}{3}$.

For each of the following fractals, we present data for the number of components, $\beta_0(\epsilon)$, the number of holes, $\beta_1(\epsilon)$, and the number of isolated points, $I(\epsilon)$. Where appropriate, we calculate the growth rates, γ_0 and γ_1; the results are compiled in Table I.

TABLE I
Computed values of γ_0 and γ_1. The results are the slope of a linear least-squares fit to the data and error bounds are estimated by varying the scaling range. The exact values for the non-zero exponents are the same as the similarity dimension, $\log 3/\log 2 \approx 1.585$.

Data Set		γ_0	γ_1
Sierpinski triangle	fig. 3	0	1.59 ± 0.02
Cantor set	fig. 5	1.40 ± 0.05	0
Simply connected set	fig. 7	0	0
Disconnected set	fig. 9	1.42 ± 0.05	0

The Sierpinski Triangle

The generating functions for the Sierpinski triangle are:

$$\begin{aligned} f_1(x,y) &= \tfrac{1}{2}(x,y) \\ f_2(x,y) &= \tfrac{1}{2}(x+1,y) \\ f_3(x,y) &= \tfrac{1}{2}(x,y+1). \end{aligned} \tag{16}$$

A finite point-set approximation to the triangle, an ϵ-neighborhood and the corresponding subset of the Delaunay triangulation are shown in Figure 3. The underlying set is perfect and connected with infinitely many holes, so we should see $\beta_0(\epsilon) = 1$ and $\beta_1(\epsilon) \to \infty$ as $\epsilon \to 0$. The exact self-similarity of the Sierpinski triangle allows us to deduce that for

$$\begin{aligned} \epsilon_0/2^{n+1} < \epsilon_n < \epsilon_0/2^n, \\ \beta_1(\epsilon_n) = \sum_{k=0}^{n} 3^k = \tfrac{1}{2}(3^{n+1}-1). \end{aligned}$$

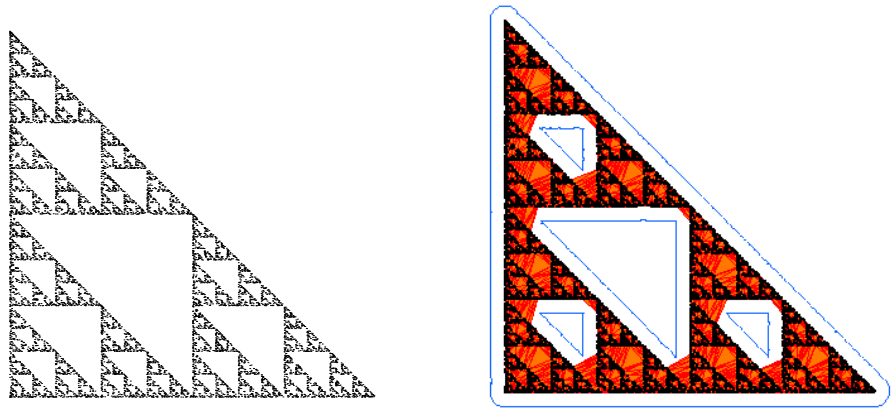

(a) (b)

Fig. 3. (a) 10^4 points on the Sierpinski triangle. (b) An ϵ-neighborhood (grey outline) and corresponding subset of the Delaunay triangulation (solid grey).

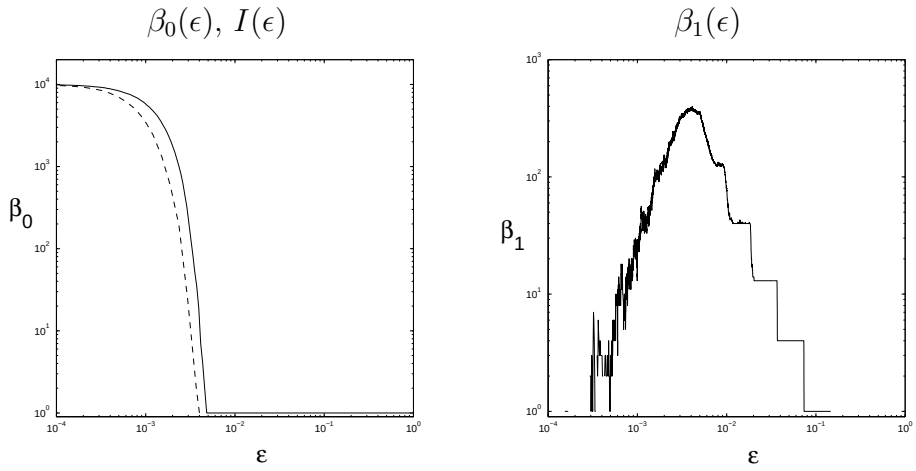

Fig. 4. Number of components, $\beta_0(\epsilon)$, and number of holes, $\beta_1(\epsilon)$, for 10^4 points uniformly distributed over the Sierpinski triangle. The dashed line in the graph of $\beta_0(\epsilon)$ is the number of isolated points, $I(\epsilon)$. All axes are logarithmic. The horizontal axis range is $10^{-4} < \epsilon < 1$.

In the Euclidean metric, ϵ_0 is the radius of the largest circle that is inscribed by the triangle with vertices at $(\frac{1}{2}, 0)$, $(0, \frac{1}{2})$ and $(\frac{1}{2}, \frac{1}{2})$, so $\epsilon_0 = \frac{1}{2} - \frac{\sqrt{2}}{4} \approx 0.146$. The growth rate, $\gamma_1 = \log 3/\log 2 \approx 1.585$, is the same as the similarity dimension.

These expectations are reflected by the computations of $\beta_0(\epsilon)$, $\beta_1(\epsilon)$ and $I(\epsilon)$, graphed in Figure 4, for a 10^4 point approximation to the triangle. We see that for ϵ above a threshold value, the computed values of $\beta_0(\epsilon)$ and $\beta_1(\epsilon)$ are in agreement with the theory. The point at which $\beta_0(\epsilon)$ and $\beta_1(\epsilon)$ "blur" is approximately $\epsilon = 0.003$, close to the value at which the number of isolated points, $I(\epsilon)$, becomes positive. This ϵ value is, of course, the cutoff resolution ρ discussed in section 3. At finer resolutions — i.e., $\epsilon < \rho$ — we see a sharp transition in the graph of β_0 from one to the number of points in the set, as each point becomes isolated. The graph of β_1 shows that the holes are destroyed as ϵ decreases. This is because the edges that form the loops are eventually deleted from the triangulation. We estimate the slope of the staircase by a linear, least squares fit and find $\gamma_1 \approx 1.59$. This is very close to the value derived above.

A Cantor Set Relative

Figure 5 shows the attractor for the iterated function system generated by

$$\begin{aligned} f_1(x, y) &= \tfrac{1}{2}(-y + 1, x) \\ f_2(x, y) &= \tfrac{1}{2}(y + 1, x) \\ f_3(x, y) &= \tfrac{1}{2}(y, -x + 2). \end{aligned} \tag{17}$$

This fractal is a Cantor set, and therefore perfect and totally disconnected, so we should see $\beta_0(\epsilon) \to \infty$ as $\epsilon \to 0$ and $\beta_1^0(\epsilon) = 0$. In [19] we derived the following form for $\beta_0(\epsilon)$:

$$\epsilon_0/2^{n+1} < \epsilon_n < \epsilon_0/2^n$$

$$\beta_0(\epsilon_n) = \begin{cases} 3^n + 2 \cdot 3^{(n-1)/2} & \text{if } n \text{ is odd} \\ 3^n + 3^{n/2} & \text{if } n \text{ is even.} \end{cases}$$

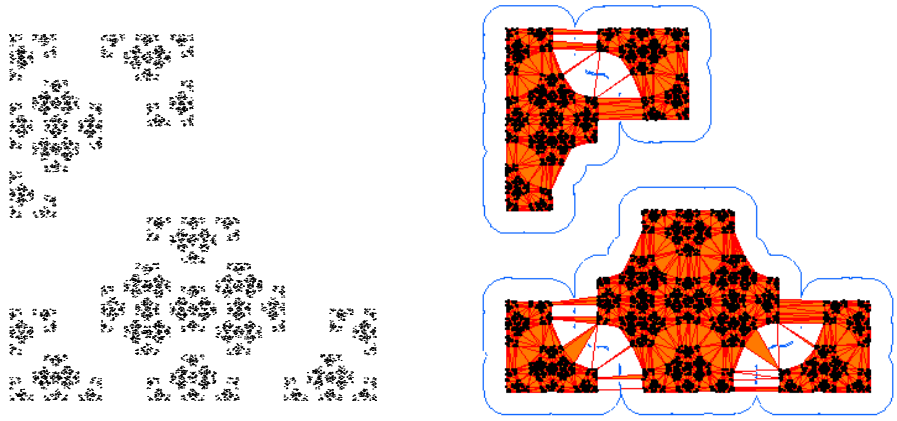

Fig. 5. (a) 10^4 points on the Cantor set relative. (b) An ϵ-neighborhood and corresponding subset of the Delaunay triangulation.

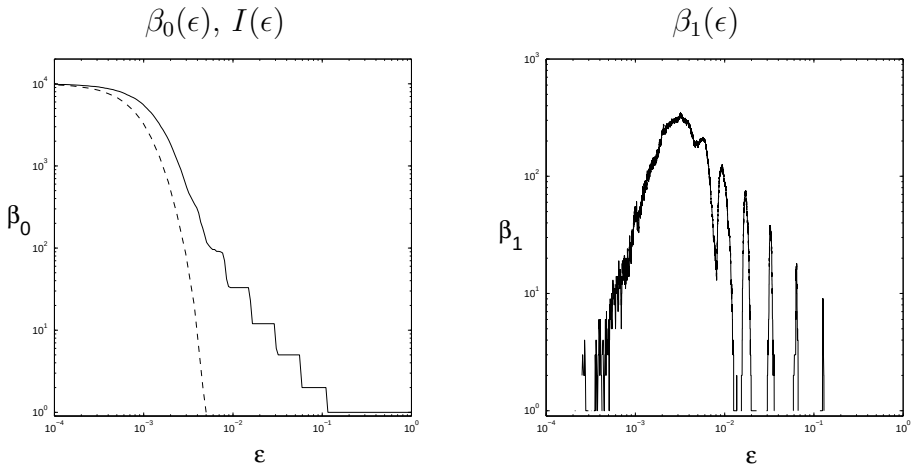

Fig. 6. Number of components, $\beta_0(\epsilon)$, and number of holes, $\beta_1(\epsilon)$, for 10^4 points uniformly distributed over the Cantor set relative. The dashed line in the graph of $\beta_0(\epsilon)$ is the number of isolated points, $I(\epsilon)$. All axes are logarithmic. The horizontal axis range is $10^{-4} < \epsilon < 1$.

Here, ϵ_0 is the smallest value of ϵ for which S_ϵ is connected. Using (10), we calculate the value of γ_0 to be $\log 3/\log 2$.

We can see in Figure 6 that the numerical computations of $\beta_0(\epsilon)$ agree very well with the theoretical values above, when ϵ is greater than the cutoff resolution $\rho \approx 0.005$. Using a least squares fit, we estimate the slope of the staircase to be $\gamma_0 \approx 1.40$. This is significantly lower than the true limiting value given above because of second-order effects at the relatively large values of ϵ for which the $\beta_0(\epsilon)$ data are valid.

The graph of $\beta_1(\epsilon)$ has regularly spaced spikes. The reason for this is seen in Figure 5(b), where the ϵ-neighborhood has two connected components. Data points in the two components are separated by a distance of at least $2\epsilon_0$. The value of ϵ in Figure 5(b) is just greater than $\epsilon_0/2$ — the value at which the ϵ-neighborhood breaks into five components. Holes appear in the triangulation at the value of ϵ displayed because of the few remaining edges that bridge the "gap". These edges are deleted at slightly smaller values of ϵ and the holes disappear. Graphing the persistent Betti number, $\beta_1^\rho(\epsilon)$, would remove these spikes.

A Simply Connected Relative

A simply connected relative of the Sierpinski triangle, shown in Figure 7, is generated by:

$$\begin{aligned} f_1(x,y) &= \tfrac{1}{2}(-x+1,-y+1) \\ f_2(x,y) &= \tfrac{1}{2}(-y+2,x) \\ f_3(x,y) &= \tfrac{1}{2}(x,y+1). \end{aligned} \tag{18}$$

We expect to see a single connected component and no persistent holes. However, holes do appear in the ϵ-neighborhoods, as seen in Figure 8, due to the geometry of the fractal. Again, these spikes in the graph of $\beta_1(\epsilon)$ would disappear in a graph of the persistent Betti number.

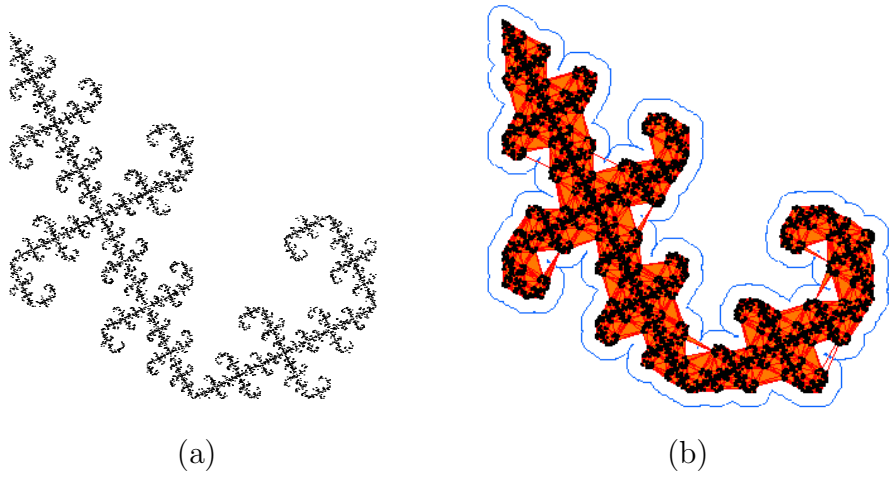

Fig. 7. (a) 10^4 points on a simply connected set. (b) An ϵ-neighborhood and corresponding subset of the Delaunay triangulation.

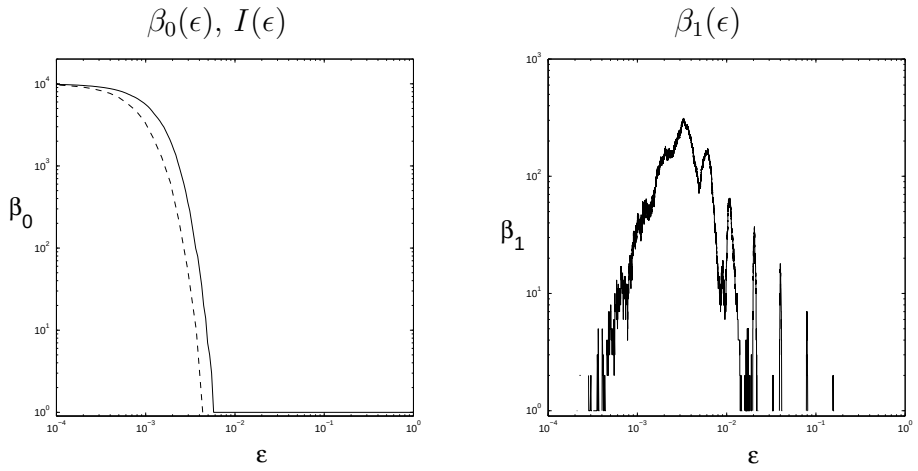

Fig. 8. Number of components, $\beta_0(\epsilon)$, and number of holes, $\beta_1(\epsilon)$, for 10^4 points uniformly distributed over a simply connected fractal. The dashed line in the graph of $\beta_0(\epsilon)$ is the number of isolated points, $I(\epsilon)$. All axes are logarithmic. The horizontal axis range is $10^{-4} < \epsilon < 1$.

A Relative with Infinitely Many Connected Components

A fourth triangle relative, shown in Figure 9, is generated by the following similarities:

$$\begin{aligned} f_1(x,y) &= \tfrac{1}{2}(x,y) \\ f_2(x,y) &= \tfrac{1}{2}(y+1,-x+1) \\ f_3(x,y) &= \tfrac{1}{2}(x,y+1). \end{aligned} \tag{19}$$

The attractor for this system has infinitely many connected components, yet is not totally disconnected because the components have positive diameters. Thus, we expect $\beta_0(\epsilon) \to \infty$ and $\beta_1^0(\epsilon) = 0$. Again, self-similarity means that for:

$$\begin{aligned} \epsilon_0/2^{n+1} < \epsilon_n < \epsilon_0/2^n, \\ \beta_0(\epsilon_n) = \tfrac{1}{2}(3^{n+1}+1), \end{aligned}$$

giving a growth rate of $\gamma_0 = \log 3/\log 2$. We estimate the slope of the graph of $\beta_0(\epsilon)$ as before and find $\gamma_0 \approx 1.42$. This is lower than the limiting value because of the small range of ϵ for which the computed values of $\beta_0(\epsilon)$ reflect those of the underlying fractal.

This example is another good illustration of why the computation of persistent Betti numbers is important. Topologically, the set is composed of disconnected line segments, so there can be no non-bounding cycles in the first Čech homology group. However, the geometry of the set creates holes in the ϵ-neighborhoods, as seen in Figure 9(b). In the previous example, it is clear that the holes do not persist as ϵ decreases because $\beta_1(\epsilon) = 0$ between the spikes. In Figure 10 we see an apparent growth in $\beta_1(\epsilon)$. The difference is that smaller holes appear before the larger ones disappear, resulting in an accumulation.

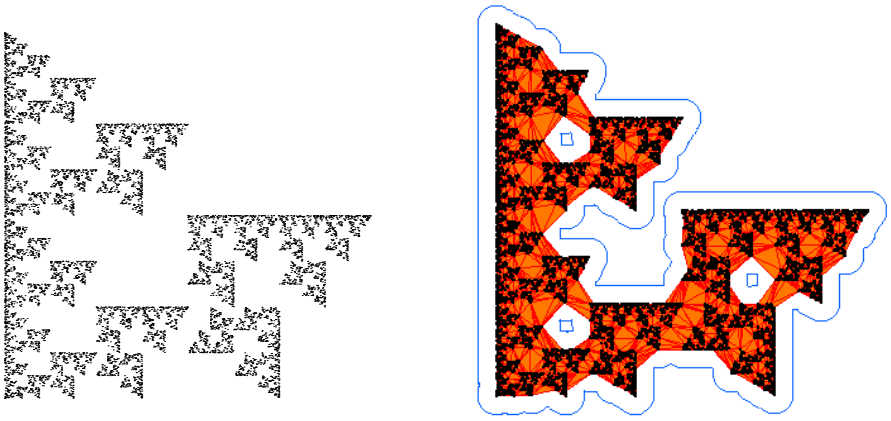

Fig. 9. (a) 10^4 points on a disconnected set. (b) An ϵ-neighborhood and corresponding subset of the Delaunay triangulation.

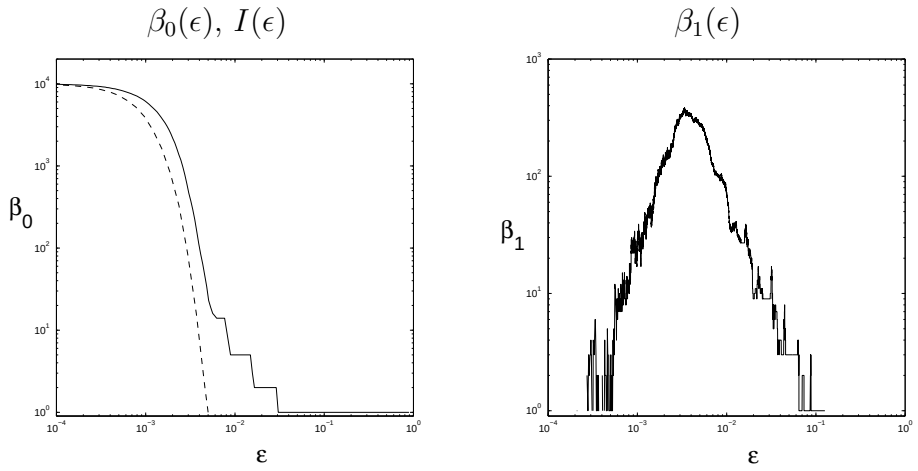

Fig. 10. Number of components, $\beta_0(\epsilon)$, number of holes, $\beta_1(\epsilon)$, for 10^4 points uniformly distributed over the disconnected relative. The dashed line in the graph of $\beta_0(\epsilon)$ is the number of isolated points, $I(\epsilon)$. All axes are logarithmic. The horizontal axis range is $10^{-4} < \epsilon < 1$.

Notice that this set has the same values of γ_0 and γ_1 as the Cantor set relative. This implies that we still do not have enough information to completely distinguish their different topological structure. The problem was solved in [20] by examining the diameters of the connected components. For a Cantor set, the diameters go to zero, while for the example in Figure 9, the largest diameter converges to one.

5. Concluding Remarks

This paper demonstrates that it is possible to extract information about the topological structure of a compact space from a finite approximation to it. The examples given here are subsets of the plane where it is easy to see the structure of the data. In higher-dimensional spaces, it is extremely difficult to visualize the underlying topology of a cloud of points, so computational techniques are necessary. In terms of analyzing data, our goal is to use the persistent Betti numbers of ϵ-neighborhoods to extrapolate the topology of the space the data approximate. In Section 4, we computed only the regular Betti numbers, $\beta_k(\epsilon)$, of each ϵ-neighborhood. The examples given highlight the need for an efficient implementation for computing the persistent Betti numbers from data.

The persistent Betti numbers of a compact space, $\beta_k^0(X_\epsilon)$, reflect the topology of the limit space X, while the regular Betti numbers, $\beta_k(\epsilon)$, characterize the topology of the ϵ-neighborhood. In applications, both sets of numbers may prove useful, since the Betti numbers of the ϵ-neighborhoods give geometric information about how their limit space is embedded. Questions about the behavior of $\beta_k^\lambda(X_\epsilon)$ as a function of λ and ϵ remain, specifically with regards to continuity properties as $\epsilon, \lambda \to 0$ for different classes of spaces.

Our work also gives a new perspective on fractal dimensions. In [20] we conjecture that for Cantor sets of zero Lebesgue measure, the disconnectedness index, γ_0, takes the same value as the

box-counting dimension. This is definitely the case for Cantor subsets of $\mathbb{R}^2$, as implied by results in [22] (in fact γ_0 is equivalent to the fat fractal exponent). It is possible that a similar relationship holds for γ_k when $k \geq 1$.

The computational implementation is far from complete. Although we show the ρ-persistent Betti numbers of data to be computable via linear algebra techniques, faster algorithms are necessary for the large data sets typically encountered in the study of dynamical systems. We anticipate that efficient computational tools will assist in the numerical investigation of chaotic systems by giving a deeper understanding of the structure of attractors and other invariant sets.

Acknowledgments

I would like to thank J.D. Meiss for many useful discussions, and E. Bradley and R. Easton for their helpful suggestions. This work was supported by EB's NSF National Young Investigator award #CCR-9357740. I am indebted to the anonymous referee for correcting an error in the original version of this paper. I also thank J. Kennedy for inviting me to participate in the 14th Summer Topology conference at the C.W. Post campus of Long Island University. My attendance was made possible by the generosity of the organizers who waived registration and accommodation costs, and a travel grant from the University of Colorado Graduate School.

References

[1] M.F. Barnsley. *Fractals Everywhere*. Academic Press, Boston, second edition, 1993.

[2] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.

[3] C.J.A. Delfinado and H. Edelsbrunner. An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere. *Computer Aided Geometric Design*, 12:771–784, 1995.

[4] T.K. Dey and S. Guha. Computing homology groups of simplicial complexes in $\mathbf{R^3}$. *Journal of the ACM*, 45:266–287, 1998.

[5] H. Edelsbrunner. The union of balls and its dual shape. *Discrete and Computational Geometry*, 1995.

[6] H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. *ACM Transactions on Graphics*, 13:43–72, 1994.

[7] J. Friedman. Computing Betti numbers via combinatorial Laplacians. *Algorithmica*, 21:331–346, 1998.

[8] G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins, third edition, 1996.

[9] J.G. Hocking and G.S. Young. *Topology*. Addison-Wesley, 1961.

[10] J.E. Hutchinson. Fractals and self similarity. *Indiana University Mathematics Journal*, 30:713–747, 1981.

[11] W.D. Kalies, K. Mischaikow, and G. Watson. Cubical approximation and computation of homology. *Banach Center Publications*, 47:115–131, 1999.

[12] S. Mardešić and J. Segal. *Shape Theory*. North-Holland, 1982.

[13] J.D. Meiss. Symplectic maps, variational principles, and transport. *Reviews of Modern Physics*, 64:795–847, 1992.

[14] K. Mischaikow, M. Mrozek, J. Reiss, and A. Szymczak. Construction of symbolic dynamics from experimental time series. *Physical Review Letters*, 82(6):1144–1147, 1999.

[15] M.R. Muldoon, R.S. MacKay, J.P. Huke, and D.S. Broomhead. Topology from a time series. *Physica D*, 65:1–16, 1993.

[16] J.R. Munkres. *Elements of Algebraic Topology*. Benjamin Cummings, 1984.

[17] E. Ott, T. Sauer, and J.A. Yorke. *Coping with Chaos*. Wiley, 1994.

[18] H.-O. Peitgen, S. Jürgens, and D. Saupe. *Chaos and Fractals: New Frontiers of Science*. Springer-Verlag, 1992.

[19] V. Robins, J.D. Meiss, and E. Bradley. Computing connectedness: An exercise in computational topology. *Nonlinearity*, 11:913–922, 1998.

[20] V. Robins, J.D. Meiss, and E. Bradley. Computing connectedness: Disconnectedness and discreteness. *Physica D*, 139:276–300, 2000.

[21] C. Robinson. *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos*. CRC Press, Boca Raton, 1995.

[22] C. Tricot. *Curves and Fractal Dimension*. Springer-Verlag, 1995.

[23] Software available from NCSA via anonymous ftp from ftp.ncsa.uiuc.edu/Visualization/Alpha-shape/.

Department of Applied Mathematics, Research School of Physical Sciences and Engineering, Australian National University, Canberra ACT 0200 AUSTRALIA

E-mail address: Vanessa.Robins@anu.edu.au