

Simplicial Multivalued Maps and the Witness Complex for Dynamical Analysis of Time Series*

Zachary Alexander[†], Elizabeth Bradley[†], James D. Meiss[†], and Nicole F. Sanderson[†]

Abstract. Topology-based analysis of time-series data from dynamical systems is powerful: it potentially allows for computer-based proofs of the existence of various classes of regular and chaotic invariant sets for high-dimensional dynamics. Standard methods are based on a cubical discretization of the dynamics and use the time series to construct an outer approximation of the underlying dynamical system. The resulting multivalued map can be used to compute the Conley index of isolated invariant sets of cubes. In this paper we introduce a discretization that uses instead a simplicial complex constructed from a witness-landmark relationship. The goal is to obtain a natural discretization that is more tightly connected with the invariant density of the time series itself. The time-ordering of the data also directly leads to a map on this simplicial complex that we call the witness map. We obtain conditions under which this witness map gives an outer approximation of the dynamics and thus can be used to compute the Conley index of isolated invariant sets. The method is illustrated by a simple example using data from the classical Hénon map.

Key words. computational topology, simplicial homology, Conley indices, Hénon chaotic attractor

AMS subject classifications. 37M10, 65P20, 55U10, 37C70

DOI. 10.1137/140971415

1. Introduction. Our goal in this paper is to develop some new computational topology techniques to characterize some aspects of discrete or continuous dynamical systems. We assume that the only knowledge we have of the dynamics is a finite time series $\Gamma = \{x_0, x_1, \dots, x_{T-1}\}$ taken from a state-space trajectory of the system. If the system is a map, $f : X \rightarrow X$, then Γ is simply the iterates of the map: $x_{t+1} = f(x_t)$. If the system is a flow, then Γ is a sequence of samples of the continuous trajectory $x(t)$, and we are effectively studying the evolution operator that maps the system forward in time. In either case, given this information, we cannot hope to approximate the dynamics on all of X ; instead, we assume that Γ lies close to Λ , a bounded invariant set of f . For example, any orbit in the basin of an attractor will eventually approach it, so in this case, Γ can be taken to be the trajectory after a transient is removed. Thus our goal is to develop tools that will allow us to characterize properties of $f|_{\Lambda}$, such as the number and types of periodic orbits, compute topological entropy, etc.

*Received by the editors June 3, 2014; accepted for publication (in revised form) by J. van den Berg May 6, 2015; published electronically July 16, 2015. This material is based on work supported by the National Science Foundation under grants CMMI-1245947, CNS-0720692, and DMS-1211350. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

<http://www.siam.org/journals/siads/14-3/97141.html>

[†]Departments of Computer Science, Applied Mathematics, and Mathematics, University of Colorado, Boulder, CO 80309-0526 (zaalexan@microsoft.com, lizb@colorado.edu, james.meiss@colorado.edu, nicole.sanderson@colorado.edu).

The tool that we use is the discrete Conley index (we recall the definition of this index and related concepts in Appendix B) [Con78, Eas98, KMM04]. The Conley index characterizes some dynamical properties of isolated invariant sets of f —for example it may establish the existence of periodic orbits or give a lower bound on the topological entropy.

Since we only have access to the trajectory Γ , we must use it to obtain an approximation of the underlying dynamics before we can compute the Conley index. There are two categories of maps that can serve this purpose. The first, as we recall in section 2, is multivalued maps. These are set-valued and are typically defined on a finite covering of a neighborhood of the invariant set Λ ; they capture how the images of the cover map across other elements of the cover. Multivalued maps have most commonly been defined on cubical grids [MMSR97, MMRS99], but as we discuss in section 2.1, more general grids can also be used. A multivalued map on a grid, which we call a *cellular multivalued map* (CMM) in section 2.2, is defined to be constant on the interior of each cell as well as on subsets of the boundary where groups of cells intersect. Note that while connectivity is obvious in uniform grids of cubical cells, this may not be the case in other grid geometries. The second category of map addresses this issue. Dual to any grid, cubical or otherwise, is a simplicial complex—the *nerve* of the grid (see Appendix A). We call the map induced on this complex a *simplicial multivalued map* (SMM) in section 2.2. When the number of grid cells is finite, such maps are finitely representable: they can be stored precisely in a computer and used to perform exact computations.

The Conley index for an isolated invariant set (defined in Appendix B) can be computed using a corresponding CMM or SMM (techniques are recalled in Appendix C). Moreover, as described in section 2.2, if the cellular map is semicontinuous and *acyclic*, then the map that it induces on homology coincides with that induced by f , and thus it can be used to compute the Conley index of isolated invariant sets of f .

The associated computational cost of these computations depends on the geometry of the cells. To minimize this complexity while still preserving the essential features, our approach uses two constructions that play major roles in computational topology: the α -*diagram* [EM92, Ede95] and the *witness complex* [dSC04]. In section 2.3 we recall that the α -diagram of a data set is the intersection of its Voronoi diagram with the union of balls of radius α centered on the data points. The nerve of the α -diagram is the α -complex: it is generically a simplicial complex and is a subset of the Delaunay triangulation (see Appendix A), limiting to the latter as $\alpha \rightarrow \infty$ [Ede95]. Since the geometry of cells in the α -diagram is dictated by the data, rather than by rectilinear grid lines, the shape of the α -complex naturally follows that of the invariant set Λ .

While this flexible, data-driven representation has some appealing advantages, an α -complex constructed from a long time series $\Gamma = \{x_0, x_1, \dots, x_{T-1}\}$ would have at least one simplex for each point, and the complexity of algorithms that construct and manipulate these objects scales poorly with the number of simplices. It is useful, then, to represent these data using a global topological object that contains fewer simplices while preserving the Conley index. We use the *witness complex* [dSC04] for this purpose; see section 3. Instead of assigning a vertex to each point in Γ , we represent the data by a smaller set of vertices, a set of *landmarks*, $L \subset X$, and build a simplicial complex from those points. As described in section 2.3, there are a number of ways to choose landmarks. The computational complexity of this approach and its comparison to that of a cubical grid are discussed in Appendix D.

The witness complex is constructed from a relation on $\Gamma \times L$: each point in Γ may be a witness to one or more landmarks, and each landmark may have one or more witnesses. Of the many possible definitions of witness relation, we choose one in which a point $x \in \Gamma$ witnesses a set $\sigma \subset L$ if the distance between x and any landmark in σ is no more than ε greater than the minimum distance between x and the full set L of landmarks; see section 3.1. We use this witness relation to construct an abstract witness complex. The simplest implementation gives a clique or flag complex: it consists of simplices whose pairs of vertices have a common witness. As we show in section 3.2, if the landmarks are selected to be sufficiently uniform and the trajectory is sufficiently dense, then there are conditions under which the witness complex and the α -complex are the same.

The dynamics on the time series induces an SMM on the witness complex. This SMM also induces a corresponding CMM on a grid of α -cells based on the landmarks. These *witness maps*, which are a primary contribution of this paper, are described in section 3.3. The construction of the witness map is developed in several steps in section 3 in order to bring the well-developed theory of [KMM04] to bear on this new formulation and thereby establish the correspondence between the homology of the witness map and that of the true dynamics of the underlying system.

Finally, in section 4 we use data generated from the classic Hénon quadratic map to give a simple illustration of the ideas in this paper. We show that, under a verifiable set of assumptions, our techniques could be used to obtain rigorous results about the underlying dynamical system.

2. Multivalued maps. In this section we describe the concept of *multivalued maps* and obtain criteria for such maps to be an *enclosure* of a map f . A *cellular multivalued map* (CMM) is defined to be constant on each cell of a *grid*, generalizing the cubical case of [KMM04]. A cellular map gives rise to a *simplicial multivalued map* (SMM) on the nerve of the grid. In the final part of this section we show that one way to construct a grid is through an α -*diagram*. In this case, the nerve is a geometrical simplicial complex that is a deformation retract of the grid, and we will show that the cellular map and the simplicial map induce the same maps on homology.

Since we are interested in applications to data sets that correspond to real-valued measurements of continuous dynamical systems, we will assume that our time series Γ is obtained from a map f on a submanifold $X \subset \mathbb{R}^n$. We will use the Euclidean metric, $d(\cdot, \cdot)$, on \mathbb{R}^n . In the future, it might be useful to consider more general metrics on the submanifold X itself.

We begin by recalling some standard definitions for multivalued maps that approximate a dynamical system.¹

Definition (multivalued map). A multivalued map, $F : X \rightrightarrows X$, is a map from X to its power set. That is, for each $x \in X$, $F(x)$ is a subset of X .

We use multivalued maps to approximate continuous maps $f : X \rightarrow X$, and the approximation is taken to be “good” if the action on homology induced by F is equivalent to that induced by f . In order for this to be the case, the action of F must *enclose* the action of f

¹Following [DJM04, DFT08], we denote single-valued maps with lower-case letters, e.g., f , sets and set-valued maps with capital letters, e.g., F , and combinatorial objects and maps with calligraphic letters, e.g., \mathcal{F} .

and not introduce any extra homological structure. These requirements are spelled out in the following definitions.

Definition (outer approximation). A multivalued map $F : X \rightrightarrows X$ is an outer approximation of a continuous map $f : X \rightarrow X$ if $f(x) \in F(x)$ for each $x \in X$. In this case f is said to be a continuous selector for F .

The (weak) preimage of a multivalued map is itself a multivalued map defined as

$$F^{-1}(y) = \{x \in X : y \in F(x)\}.$$

Definition (semicontinuous). A multivalued map is (lower) semicontinuous if the preimage of each open set is open.

As usual, the n -dimensional homology group of a set X is denoted $H_n(X)$. We use, for simplicity, the homology over \mathbb{Z}_2 so that the torsion subgroups are ignored. Given this, a multivalued map preserves homology if the image of each point is homologous to a point. This is captured in the following definition.

Definition (acyclic). A multivalued map F is acyclic if for each $x \in X$,

$$H_n(F(x)) = \begin{cases} \mathbb{Z}_2, & n = 0, \\ \mathbf{0}, & n > 0. \end{cases}$$

The key point is that if F is semicontinuous and acyclic, then every continuous selector for F induces the same homomorphism in homology—a consequence of the acyclic carrier theorem [Mun84, Thm. 13.3].

Definition (enclosure). A semicontinuous, acyclic, multivalued map F is an enclosure of any continuous selector f .

Therefore, one can define the homology induced by an enclosure to be that of any of its continuous selectors.

In order that this homology be computable, however, it is necessary to obtain a finitely representable approximation of the map f , and it is this to which we turn next.

2.1. Grids. A grid allows one to construct finitely representable maps that can be outer approximations of a map f [KMV05, Mro99]. We will consider generalizations of the cubical cells of [KMM04] to a grid constructed from a collection of cells $\mathcal{A} = \{A_1, A_2, \dots\}$ in X . Associated with any such collection is its geometrical realization—the union of these cells as subsets of X —denoted by

$$(1) \quad |\mathcal{A}| := \bigcup_{A \in \mathcal{A}} A.$$

Since the shape and number of neighbors of each cell can vary, such a grid may permit more efficient computational algorithms than those for cubes of fixed size and shape.² Nevertheless, many of the results that appear here are easily adapted from the cubical case.

There are four basic properties for the cells of a grid.

Definition (grid [AKK⁺09]). A family of nonempty compact sets \mathcal{A} is a grid on X if

(a) $X = |\mathcal{A}|$;

²See Appendix D for more discussion of algorithms and associated computational complexity issues.

- (b) for all $A \in \mathcal{A}$, $A = cl(int(A))$;
- (c) for all $A, B \in \mathcal{A}$, if $A \neq B$, then $A \cap int(B) = \emptyset$; and
- (d) a finite subset of \mathcal{A} covers each compact $S \subset X$.

A prototypical grid is a lattice of closed cubes; indeed, this is the example studied extensively in [KMM04]. An example of a more general grid is shown in Figure 1. The cells in this grid are α -cells; see section 2.3.

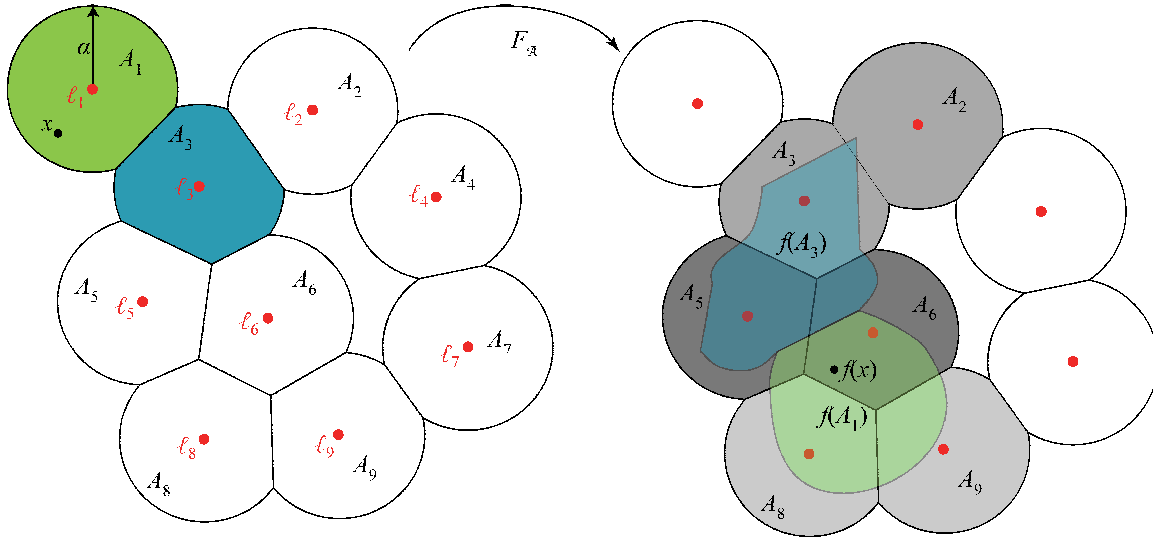


Figure 1. Sketch of a grid, $\mathcal{A}_\alpha(L)$, of α -cells (see section 2.3) based on a set of nine landmarks $\{l_i\}$ (red points) and the action of a cellular map $F_{\mathcal{A}}$ (2) on two of these cells. Here the image of any $x \in int(A_1)$ is $A_5 \cup A_6 \cup A_8 \cup A_9$, and that of any $x \in int(A_3)$ is $A_2 \cup A_3 \cup A_5 \cup A_6$. If $x \in A_1 \cap A_3$, then $F_{\mathcal{A}}(x) = A_5 \cup A_6$.

2.2. Cellular and simplicial maps. A cellular multivalued map (CMM) is a map on the geometrical realization, (1), of a grid \mathcal{A} that is constant on the interior of each cell of \mathcal{A} . It generalizes the cubical multivalued map of [KMM04] to a situation in which the cell boundaries need not be rectilinear.

Definition (cellular multivalued map (CMM)). A multivalued map $F_{\mathcal{A}} : |\mathcal{A}| \rightrightarrows |\mathcal{A}|$ on the geometrical realization of a grid \mathcal{A} is a cellular multivalued map (CMM) if it is the outer approximation of f defined by

$$(2) \quad F_{\mathcal{A}}(x) := \bigcap_{B \in \mathcal{A} : x \in B} \{A \in \mathcal{A} : A \cap f(B) \neq \emptyset\}.$$

The map $F_{\mathcal{A}}$ takes the interior of each cell to the union of the cells that intersect its image and the boundary shared by multiple cells to the cells that contain the intersection of their images. An illustration is shown in Figure 1. An implication is that CMMs are semicontinuous because they map the boundary of each cell to a subset of the image of the cell itself—this is a straightforward generalization of [KMM04, Prop. 6.17] for the related cubical case.

This construction is not easy to implement on a computer for two reasons: it is a map on a continuum $|\mathcal{A}|$, and to construct it we must know $f(x)$ for each point $x \in |\mathcal{A}|$. The second

problem is addressed by the *witness map* introduced in section 3.3. The first problem can be mitigated by defining a finite map, $\mathcal{F} : \mathcal{K} \rightrightarrows \mathcal{K}$, on a complex \mathcal{K} related to the grid \mathcal{A} . One such complex is the CW-complex [Hat02] formed from \mathcal{A} , that is, the collection of cells $\{A_i\}$, together with their faces $\{A_i \cap A_j\}$, their edges, and so forth. An example is the cubical CW-complex used in the approach of [KMM04]. Since the CMM $F_{\mathcal{A}}$ is constant on each cell of the CW-complex, it naturally gives rise to an associated combinatorial map. For example, in Figure 1 the two-cell A_1 would be mapped to $\{A_5, A_6, A_8, A_9\}$, and the one-cell represented by $A_1 \cap A_3$ would be mapped to $\{A_5, A_6\}$.

Our goal is to use a more easily described complex that is also naturally suited to homology calculations—in particular, a simplicial complex. A natural simplicial complex to use is the *nerve* $N(\mathcal{A})$; recall Appendix A. Let L denote a set of labels for the elements of \mathcal{A} , and let $\sigma = \langle l_0, l_2, \dots, l_k \rangle$ be any finite subset of L . The intersection of the cells labeled by the elements of σ is denoted

$$(3) \quad A_{\sigma} := \bigcap_{l \in \sigma} A_l.$$

A set σ is a simplex in the nerve when $A_{\sigma} \neq \emptyset$; thus

$$(4) \quad \mathcal{K} := N(\mathcal{A}) = \{\sigma = \langle l_0, l_1, \dots, l_k \rangle : A_{\sigma} \neq \emptyset\}.$$

For example, the nerve of the grid in Figure 1 is shown in Figure 2.

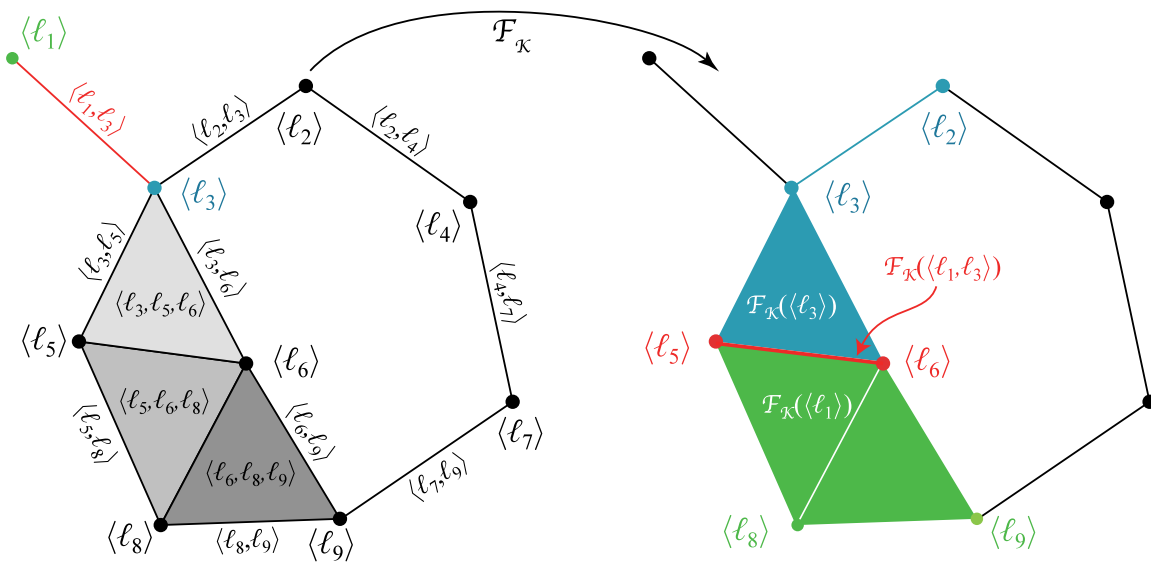


Figure 2. Simplicial complex $\mathcal{K}_{\alpha}(L)$ that is the nerve of the α -cells of Figure 1, and the action of the simplicial map $\mathcal{F}_{\mathcal{K}}$ (5) on the simplices $\langle l_1 \rangle$ (green), $\langle l_3 \rangle$ (blue), and $\langle l_1, l_3 \rangle$ (red).

A CMM $F_{\mathcal{A}}$ induces a combinatorial map $\mathcal{F}_{\mathcal{K}}$ on the nerve that is defined to commute with the correspondence between the grid and its nerve.

Definition (simplicial multivalued map (SMM)). If $F_{\mathcal{A}}$ is a CMM (2) on a grid \mathcal{A} , the map $\mathcal{F}_{\mathcal{K}} : \mathcal{K} \rightrightarrows \mathcal{K}$, defined by

$$(5) \quad \mathcal{F}_{\mathcal{K}}(\sigma) := \{\tau : A_{\tau} \subset F_{\mathcal{A}}(A_{\sigma})\},$$

is a simplicial multivalued map (SMM).

Note that $\mathcal{F}_\mathcal{K}$ is a combinatorial multivalued map. Its domain consists of simplices and its range of sets of simplices; moreover, $\mathcal{F}_\mathcal{K}(\sigma)$ is a subcomplex of \mathcal{K} (recall Appendix A). As an example, the simplicial map induced by the cellular map of Figure 1 is shown in Figure 2. In this case, since $A_{\langle l_1 \rangle} = A_1$, then $\{\langle l_5, l_6, l_8 \rangle, \langle l_6, l_8, l_9 \rangle\} \subset \mathcal{F}_\mathcal{K}(\langle l_1 \rangle)$ as are the nine faces of these two 2-simplices. Similarly, since $A_{\langle l_1, l_3 \rangle} = A_1 \cap A_3$, then $\mathcal{F}_\mathcal{K}(\langle l_1, l_3 \rangle) = \{\langle l_5, l_6 \rangle, \langle l_5 \rangle, \langle l_6 \rangle\}$.

The definition (5) satisfies the closed graph condition.

Lemma 1 (closed graph condition). *If $\mathcal{F}_\mathcal{K}$ is an SMM and $\tau \leq \sigma \in \mathcal{K}$ (i.e., τ is a face of σ), then $\mathcal{F}_\mathcal{K}(\tau) \supseteq \mathcal{F}_\mathcal{K}(\sigma)$.*

Proof. Since $\tau \leq \sigma$, then by (3) $A_\tau \supseteq A_\sigma$. Equation (2) implies that $F_\mathcal{A}(A_\tau) \supseteq F_\mathcal{A}(A_\sigma)$, and so by (5) $\mathcal{F}_\mathcal{K}(\tau) \supseteq \mathcal{F}_\mathcal{K}(\sigma)$. ■

When the nerve \mathcal{K} of a grid \mathcal{A} is a geometrical simplicial complex (again, recall Appendix A), it induces a natural CMM on the geometrical realization $|\mathcal{K}|$, the union of the convex hull $|\sigma| \in \mathbb{R}^n$ of each of its geometrical simplices. The map $F_\mathcal{K} : |\mathcal{K}| \rightrightarrows |\mathcal{K}|$ is the multivalued map induced by $\mathcal{F}_\mathcal{K}$; i.e., by analogy with (2),

$$(6) \quad F_\mathcal{K}(x) := \bigcap_{\sigma \in \mathcal{K} : x \in |\sigma|} \{|\mathcal{F}_\mathcal{K}(\sigma)|\}.$$

In certain cases, $F_\mathcal{A}$ and $F_\mathcal{K}$ contain the same information about homology. We can show this when the geometric realization of the complex is a (strong) deformation retract of the geometric realization of the grid, i.e., when $|\mathcal{K}| \subset |\mathcal{A}|$ and there exists a continuous map $r : |\mathcal{A}| \times [0, 1] \rightarrow |\mathcal{A}|$, such that

$$(7) \quad \begin{aligned} r(x, 0) &= x, \\ r(x, t) &= x \quad \text{if } x \in |\mathcal{K}|, \text{ and} \\ r(x, 1) &:= \rho(x) \in |\mathcal{K}|. \end{aligned}$$

As we will see in section 2.3, this assumption can be verified for an α -grid and its nerve. We begin with the following ‘‘partial commutativity’’ lemma.

Lemma 2. *Suppose that $\mathcal{K} = N(\mathcal{A})$ is a geometrical simplicial complex. Let $F_\mathcal{A}$ and $F_\mathcal{K}$ be CMMs as in (2) and (6). If $|\mathcal{K}|$ is a deformation retract of $|\mathcal{A}|$ and $\rho(A_i) = r(A_i, 1) \subset A_i$, then for any $x \in |\mathcal{A}|$, $(F_\mathcal{K} \circ \rho)(x) \subseteq (\rho \circ F_\mathcal{A})(x)$.*

Proof. Note that since ρ is onto $|\mathcal{K}|$ and is the identity on $|\mathcal{K}|$, $\rho(A_i) = |\mathcal{K}| \cap A_i$ and it suffices to show that for any $x \in |\mathcal{A}|$, $F_\mathcal{K}(\rho(x)) \subseteq F_\mathcal{A}(x)$. Furthermore, we note that it follows directly from (5) that for any simplex $\sigma \in \mathcal{K}$, $|\mathcal{F}_\mathcal{K}(\sigma)| \subset F_\mathcal{A}(A_\sigma)$. For an $x \in A_\sigma$, $\rho(x)$ is a point in a geometric simplex $|\sigma|$. Therefore, by (6),

$$F_\mathcal{K}(\rho(x)) \subseteq |\mathcal{F}_\mathcal{K}(\sigma)| \subseteq F_\mathcal{A}(A_\sigma),$$

as required. ■

This result is exactly what is needed to show that the maps on homology induced by $F_\mathcal{K}$ and $F_\mathcal{A}$ are isomorphic. In particular, we can prove the following theorem.

Theorem 3. *Under the hypotheses of Lemma 2, whenever $F_\mathcal{A}$ is an acyclic multivalued map, $F_\mathcal{K}$ induces the same map on homology as $F_\mathcal{A}$.*

Proof. Whenever $F_{\mathcal{A}}$ is an acyclic multivalued map, then the definition (5) implies that $F_{\mathcal{K}}$ is as well. Moreover, since there exists a deformation retract (7), ρ_* is the identity; thus, the maps $\rho \circ F_{\mathcal{A}}$ and $F_{\mathcal{K}} \circ \rho$ are also acyclic. Now, by Lemma 2, $F_{\mathcal{K}} \circ \rho$ is a submap of $\rho \circ F_{\mathcal{A}}$; therefore, it follows that there is a continuous map $u : |\mathcal{A}| \rightarrow |\mathcal{K}|$ that is a continuous selector for both $F_{\mathcal{K}} \circ \rho$ and $\rho \circ F_{\mathcal{A}}$. Thus, if v and w are continuous selectors for $F_{\mathcal{K}}$ and $F_{\mathcal{A}}$, respectively, then $(v \circ \rho)$ and $(\rho \circ w)$ are continuous selectors carried by $(F_{\mathcal{K}} \circ \rho)$ and $(\rho \circ F_{\mathcal{A}})$. The acyclic carrier theorem [Mun84, Thm. 13.3] then implies that

$$(v \circ \rho)_* = u_* = (\rho \circ w)_* \quad \Rightarrow \quad v_* \circ \rho_* = \rho_* \circ w_* \quad \Rightarrow \quad v_* = w_* \quad \blacksquare$$

By (2), $F_{\mathcal{A}}$ is semicontinuous and an outer approximation of f ; therefore, if $F_{\mathcal{A}}$ is acyclic, it induces a well-defined map on the homology groups such that $(F_{\mathcal{A}})_* = f_*$. That is, in order to determine the Conley index (recall Appendix B) for an isolated invariant set of f , we need only to know the map on homology induced by $F_{\mathcal{A}}$ and check acyclicity. In Appendix C, we recall the theoretical framework and algorithms from [DFT08] that can be used to compute the discrete Conley index for a multivalued map.

Note that our construction of the CMM $F_{\mathcal{A}}$ is still purely theoretical: since we only know the points in a time series, we do not know the image of every point in the state space and thus cannot compute the outer approximation. Without this, there is no direct method to compute the associated simplicial map $\mathcal{F}_{\mathcal{K}}$ on the nerve. In section 3 we define a new map, the witness map, that can be computed algorithmically from time-series data. We then provide a set of conditions under which the witness map and the CMM $F_{\mathcal{A}}$ contain the same information about homology. These results allow us to calculate the Conley index using the witness map rather than $F_{\mathcal{A}}$.

2.3. α -diagrams and complexes. The concepts in the previous sections can be specialized to the case of a grid based on an α -diagram³ for a finite set of *landmarks*, $L = \{l_1, \dots, l_\ell\} \subset \mathbb{R}^n$. Recalling that $d(x, y)$ is the Euclidean metric, we denote the closed ball of radius α centered at l by

$$(8) \quad B_\alpha(l) = \{x \in \mathbb{R}^n : d(x, l) \leq \alpha\}$$

and the distance from x to a set $S \subset \mathbb{R}^n$ by $d(x, S) = \inf_{y \in S} d(x, y)$. For a given $\alpha > 0$, the α -cell centered at a landmark point l_i is the intersection of $B_\alpha(l_i)$ with the Voronoi cell of l_i :

$$(9) \quad A_i(\alpha) := B_\alpha(l_i) \cap \{x \in \mathbb{R}^n : d(x, l_i) \leq d(x, L)\}.$$

We denote the collection of α -cells for a set of landmarks L by $\mathcal{A}_\alpha(L)$. An example of a set of landmarks and their associated α -cells is shown in Figure 1.

In our application it will be important to choose the landmarks to cover the time series Γ with some accuracy; in particular, there should be a landmark within some prescribed distance from each element of Γ . The landmarks should also be relatively sparse: the minimal distance between any two should not be too small. These two requirements are linked, as will be discussed in section 3. One selection method, advocated by [dSC04], is to choose

³See Appendix A for a discussion of α -diagrams and related simplicial complexes.

the landmarks from Γ itself by a “max-min” algorithm: choose L recursively by selecting the farthest point in Γ from the previous selection until the operative density and sparseness requirements are satisfied, if possible. For the simple example in section 4, we do not choose the landmarks from the time series, but uniformly in \mathbb{R}^n . The general problem of finding the most appropriate landmarks in an efficient way is an interesting question for future investigation.

It is straightforward to show that a collection of α -cells is a grid.

Lemma 4. *When $\alpha > 0$, the set of α -cells $\mathcal{A}_\alpha(L)$ is a grid on $|\mathcal{A}_\alpha(L)|$.*

This lemma allows us to define a CMM (2) for any α -grid with $\alpha > 0$.

The nerve of an α -diagram is the α -complex denoted

$$(10) \quad \mathcal{K}_\alpha(L) = N(\mathcal{A}_\alpha(L)).$$

Since the α -complex is a nerve, it is always an abstract simplicial complex, and, as for the general case of section 2.2, there is an associated simplicial map $\mathcal{F}_\mathcal{K}$ on (10). Moreover, whenever the vertices l_i are in “general position” (recall Appendix A), $\mathcal{K}_\alpha(L)$ is a geometrical complex so that its simplices have dimension at most n and their intersections are faces. We will always assume that the landmarks are selected to be in general position.

For this case, Edelsbrunner proved that the geometrical realization of the nerve $|\mathcal{K}_\alpha(L)|$ is a deformation retract, (7), of $|\mathcal{A}_\alpha(L)|$ [Ede95].⁴ He also showed that, since the α -cells are convex, $\rho(\cdot) = r(\cdot, 1)$ can be chosen to preserve inclusion in each specific cell: $\rho(A_i) \subset A_i$. Thus for an α -grid, the hypotheses of Lemma 2 hold. Consequently, Theorem 3 implies that whenever the cellular map $F_\mathcal{A}$ on an α -grid is an enclosure of a dynamical system f , then the Conley index of an isolated invariant set can be computed from the map induced by $F_\mathcal{K}$ on homology.

3. The witness complex and map. In this section, we define a simplicial multivalued *witness map* $\mathcal{F}_\mathcal{W}$ on a complex \mathcal{W} that is a variant of the *witness complexes* introduced by [dSC04]. The goal is to obtain an outer approximation of a continuous map $f : X \rightarrow X$ when the only data that we have is a time series $\Gamma = \{x_0, \dots, x_{T-1}\}$ near an invariant set $\Lambda \subset X$. To construct a multivalued map that approximates f , we view the data as “witnesses” to a set of ℓ nearby landmarks, $L = \{l_1, \dots, l_\ell\}$. There are many possible methods to choose appropriate landmarks. One strategy, as mentioned in section 2.3, is the max-min procedure of [dSC04]; but there are many other possible methods, and indeed it is not necessary that L be a subset of Γ . The landmarks can be viewed either as the centers of the cells of an α -grid $\mathcal{A}_\alpha(L)$ or as the vertices of a witness complex $\mathcal{W}(\Gamma, L)$. We show below that if the data satisfy certain density criteria and the landmarks are (more or less) uniformly spaced, then there is an α for which the witness complex is identical to the α -complex $\mathcal{K}_\alpha(L) = N(\mathcal{A}_\alpha(L))$.

The temporal ordering of Γ and the witness relation give rise to both an SMM $\mathcal{F}_\mathcal{W}$ and an associated CMM, $F_\mathcal{W}$, on the α -grid. We show that when the map f is Lipschitz and the trajectory Γ is dense enough, there is an α such that $F_\mathcal{W}$ is an outer approximation of f .

3.1. Witness complex. A witness complex is a simplicial complex based on a finite data set Γ that is intended to be more parsimonious than more traditional Rips or Čech complexes (recall Appendix A) [dSC04, dS08]. The vertices of the complex are taken from a set of

⁴See Lemma 10 and the discussion in Appendix A.

landmarks L whose cardinality is much smaller than that of Γ . The complex $\mathcal{W}(\Gamma, L)$ consists of those subsets of L that have a *witness* in Γ . For example, de Silva and Carlsson say a point $x \in \mathbb{R}^n$ is a *weak* witness to a k -simplex if the $k + 1$ nearest landmarks to x are the vertices of the simplex. If, in addition, x is equidistant from each of the vertices, then it is a *strong* witness to the simplex.

Another way to define witness complexes is through a general construction of Dowker [Dow52] that associates abstract simplicial complexes with a relation, that is, by the selection of a subset

$$(11) \quad R \subset \Gamma \times L.$$

For example, the strong witness complex corresponds to the relation $R_0 = \{(x, l) \in \Gamma \times L : d(x, l) = d(x, L)\}$. We say that a point x is a witness to a point l if $(x, l) \in R$. Thus,

$$W_R(\Gamma, l) = \{x \in \Gamma : (x, l) \in R\}$$

is the set of witnesses to the landmark l . Following Dowker, a relation gives rise to two abstract simplicial complexes, by vertical and horizontal slices, respectively. In our notation, the witness complex is the former: the vertices of each simplex in the complex share a witness: $\bigcap_{l \in \sigma} W_R(\Gamma, l) \neq \emptyset$.

We will use a more easily computed version of the witness complex, a *clique* complex, which is the maximal complex with a given set of edges (recall Appendix A). The clique complex for a given relation (11) is

$$(12) \quad \mathcal{W}_R(\Gamma, L) = \{\sigma : W_R(\Gamma, l) \cap W_R(\Gamma, l') \neq \emptyset \forall l, l' \in \sigma\}.$$

Note that, though the vertices of each edge in σ must share a witness, there need not be a common witness to every vertex in σ .

There are many possible choices for the witness relation R . We choose to use a fuzzy version of the strong witness relation:

$$(13) \quad R_\varepsilon = \{(x, l) \in \Gamma \times L : d(x, l) \leq d(x, L) + \varepsilon\}.$$

That is, x witnesses all landmarks no more than ε farther from x than its nearest landmark.⁵ The parameter ε represents the *fuzziness* of the boundary between cells. The definition (13) becomes the strong witness relation for $\varepsilon = 0$. We will denote the set of witnesses to a landmark using (13) by $W_\varepsilon(\Gamma, l)$, and the resulting clique complex (12) by $\mathcal{W}_\varepsilon(\Gamma, L)$.

An example is shown in Figure 3 for an orbit of the logistic map on $[0, 1]$ for a parameter value just above the first period-doubling accumulation point. Here, the landmarks were selected to be every 30th point in the sorted data, an orbit of length $T = 300$. The relation (13) for $\varepsilon = 0.01$ is the set of (blue) points near the diagonal. For the case shown, each point in Γ is a witness to at most two landmarks, and so the maximum dimension of a simplex in the complex (12) is one.

⁵In the notation of [dSC04], this relation corresponds to the complex $W(D, \varepsilon, 1)$, where D denotes the matrix of distances between landmarks and witnesses.

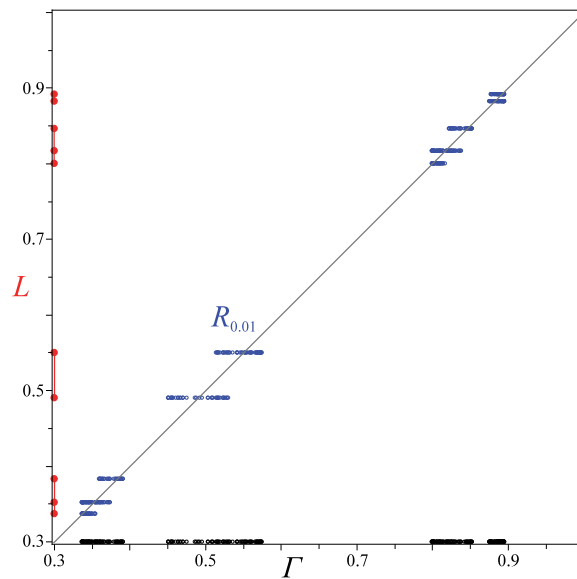


Figure 3. Witnesses (blue points) defined by the relation (13) for the logistic map, $f(x) = 3.58x(1 - x)$ with $\varepsilon = 0.01$. The orbit Γ , shown along the horizontal axis (black points), has $T = 300$ points, and there are $\ell = 10$ landmarks, shown along the vertical axes (red points). The witness relation defines six one-dimensional simplices (the line segments along the vertical axis), giving a complex with Betti number $\beta_0 = 4$. These correspond to the four major bands in the chaotic attractor of f .

One way to compute the relation (13) is to sort the rows of the $T \times \ell$ distance matrix $D_{tj} = d(x_t, l_j)$ in order of increasing size; thus, for the sorted matrix $D_{t,1}^s = d(x_t, L)$. Then x_t is a witness to all of the landmarks in the first few columns of the t th row of D^s , namely those for which $D_{t,j}^s \leq D_{t,1}^s + \varepsilon$. The main computational expense here—the distance calculations—can be reduced using a kd -tree [FBF77]. In addition, most implementations of efficient k -nearest-neighbor algorithms return their results sorted in size order. See Appendix D for more discussion of algorithms and complexity.

The following section describes how the complex $\mathcal{W}_\varepsilon(\Gamma, L)$ using the witness relation (13) can be related to an α -complex using the same landmark set, under some conditions on the selection of the landmarks and α .

3.2. Equivalence conditions for $\mathcal{K}_\alpha(L)$ and $\mathcal{W}_\varepsilon(\Gamma, L)$. The witness complex is based on the set of landmarks L in \mathbb{R}^n that can also be viewed as the centers of an α -grid $\mathcal{A}_\alpha(L)$. Since the α -diagram limits to the Voronoi diagram, it is clear that for large enough α , $|\mathcal{W}_\varepsilon(\Gamma, L)| \subset |\mathcal{A}_\alpha(L)|$. Moreover, for large enough α , the associated α -complex $\mathcal{K}_\alpha(L)$ is a clique complex—just as we have assumed for the fuzzy witness complex using (12). We will show here that when this is the case—and if the landmarks are not too closely spaced—then $\mathcal{W}_\varepsilon(\Gamma, L) \subset \mathcal{K}_\alpha(L)$. Conversely, when the data Γ are dense enough on $|\mathcal{A}_\alpha(L)|$, we will see that $\mathcal{K}_\alpha(L) \subset \mathcal{W}_\varepsilon(\Gamma, L)$. Consequently, when both sets of conditions are satisfied, the complexes are the same. As the hypotheses to obtain these results are independent, we state these two results separately.

Theorem 5. For a set of landmarks L , a time series Γ , and $\alpha, \varepsilon > 0$, let $\mathcal{K}_\alpha(L)$ be the α -complex (4) and $\mathcal{W}_\varepsilon(\Gamma, L)$ be the fuzzy witness complex (12) using the relation (13). Suppose

that there is a $\delta > 0$ such that Γ is δ -dense on $|\mathcal{A}_\alpha(L)|$ and $\delta \leq \varepsilon/2$. Then $\mathcal{K}_\alpha(L) \subseteq \mathcal{W}_\varepsilon(\Gamma, L)$.

Proof. Suppose $\sigma \in \mathcal{K}_\alpha(L)$, i.e., there is a $y \in |\mathcal{A}_\alpha(L)|$ such that $\Delta = d(y, L) = d(y, l_i) \leq \alpha$ for all $l_i \in \sigma$. We will show that there is an $x \in \Gamma$ that witnesses all the vertices in σ , i.e.,

$$x \in \bigcap_{l \in \sigma} W_\varepsilon(\Gamma, l).$$

Since Γ is δ -dense, for any $y \in |\mathcal{A}_\alpha(L)|$, there is at least one point $x \in \Gamma \cap B_\delta(y)$. Since $d(y, L) = \Delta$ and $d(x, y) \leq \delta$, it follows that

$$d(x, L) \geq \Delta - \delta.$$

Since $x \in B_\delta(y)$, for any $l_i \in \sigma$,

$$d(x, l_i) \leq \Delta + \delta \leq d(x, L) + 2\delta \leq d(x, L) + \varepsilon$$

since $\delta \leq \varepsilon/2$. Hence, $x \in W_\varepsilon(\Gamma, l_i)$ for each vertex of σ and, therefore, $\sigma \in \mathcal{W}_\varepsilon$. ■

Note that Theorem 5 applies even when the witness complex is not defined as a clique complex. However, to show the converse—as we do next—requires the clique assumption and also relies on the use of the Euclidean metric.

Theorem 6. *Suppose $\mathcal{K}_\alpha(L)$ and $\mathcal{W}_\varepsilon(\Gamma, L)$ are as in Theorem 5, and define*

$$M = \max_{x \in \Gamma} d(x, L) \quad \text{and} \quad \beta = \min_{i \neq j} d(l_i, l_j).$$

If α is chosen so that $\mathcal{K}_\alpha(L)$ is a clique complex and

$$(14) \quad M + \varepsilon \leq \alpha \leq \frac{\beta}{\sqrt{2}},$$

then $\mathcal{W}_\varepsilon(\Gamma, L) \subseteq \mathcal{K}_\alpha(L)$.

Proof. Note that $\mathcal{K}_\alpha(L)$ and $\mathcal{W}_\varepsilon(\Gamma, L)$ have the same vertex set, and by assumption each complex is a clique complex. This means that \mathcal{K}_α and \mathcal{W}_ε are each determined completely by their edges. It follows that we only need to verify that every edge in \mathcal{W}_ε is also an edge in \mathcal{K}_α .

Suppose that $\langle l_1, l_2 \rangle \in \mathcal{W}_\varepsilon$; then these landmarks share a witness; i.e., there is an $x \in \Gamma$ such that $d(x, l_i) \leq d(x, L) + \varepsilon$ for $i \in \{1, 2\}$. We want to show that there is a point $y \in |\mathcal{A}_\alpha(L)|$ such that $d(y, l_1) = d(y, l_2) = \Delta = d(y, L) \leq \alpha$. In the Euclidean metric there is always a point y equidistant from the two landmarks such that $d(y, l_i) = \frac{1}{2}d(l_1, l_2)$. Therefore, since $d(l_1, l_2) \leq d(x, l_1) + d(x, l_2) \leq 2(d(x, L) + \varepsilon)$, then

$$\Delta \leq d(x, L) + \varepsilon \leq M + \varepsilon \leq \alpha,$$

by (14). Let $l_3 \in L$ be the next closest landmark to y , besides l_1 and l_2 , and define $\beta_1 = d(l_3, l_1)$ and $\beta_2 = d(l_3, l_2)$. As illustrated in Figure 4, $\Delta' = d(l_3, y)$ is minimized when $\beta_1 = \beta_2 = \beta$. In this case, the segment from l_3 to y is the perpendicular bisector of the segment from l_1 to l_2 , and so we have $\Delta \leq \Delta'$ only if $\Delta \leq \frac{\beta}{\sqrt{2}}$. Since $\Delta \leq \alpha$, this condition is ensured by (14), and it follows that $\langle l_1, l_2 \rangle \in \mathcal{K}_\alpha$. Since \mathcal{K}_α is a clique complex, we have shown that $\mathcal{W}_\varepsilon \subseteq \mathcal{K}_\alpha$. ■

In order to apply Theorem 6, $\mathcal{K}_\alpha(L)$ must be a clique complex, which is not always true. However, since the Delaunay complex $\mathcal{D}(L)$ (recall Appendix A) is a clique complex (the

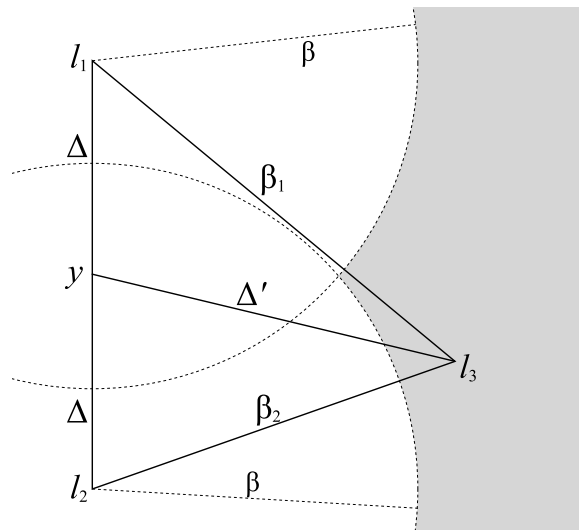


Figure 4. An illustration of the spacing between three landmarks $l_1, l_2,$ and $l_3,$ as in the proof of Theorem 6. The point y is the midpoint between the two landmarks l_1 and $l_2.$ By assumption, the distance from l_3 to l_1 or l_2 is at least $\beta,$ and it is thus minimized when $\beta_1 = \beta_2 = \beta.$

Voronoi cells cover $\mathbb{R}^d,$ then $\mathcal{K}_\infty(L) = \mathcal{D}(L)$ is a clique complex as well. Indeed, whenever α is larger than the radius of the biggest circumsphere that defines an n -dimensional simplex in $\mathcal{D}(L),$ then $\mathcal{K}_\alpha(L) = \mathcal{D}(L).$ For the simple case of a hexagonal array of landmarks in $\mathbb{R}^2,$ these circumspheres all have radius $\beta/\sqrt{3},$ so it is easy to determine when $\mathcal{K}_\alpha(L)$ is clique. For the trivial case when $\alpha < \beta/2,$ the α -balls about each landmark are disjoint, so the α -complex is trivial and also a clique complex.

When Theorems 5 and 6 both hold, then $\mathcal{W}_\varepsilon(\Gamma, L) = \mathcal{K}_\alpha(L).$ In this case, a map defined using the witness relation may have the same homology as a map on $\mathcal{A}_\alpha(L).$ It is to this issue that we turn next.

3.3. Witness map. Abstractly, we can define a CMM on a grid $\mathcal{A}(L)$ that contains the orbit Γ using any witness relation: whenever $x \in A_i$ and there is a witness $x_t \in \Gamma$ to the landmark $l_i,$ then the image of x should include the cells that x_{t+1} witnesses. The appropriate map is defined similarly to the cellular map $F_{\mathcal{A}},$ (2), but only using the data Γ and the witness relation $R.$

To obtain a cellular map that induces a simplicial map on the witness complex, we assume that the hypotheses of Theorems 5 and 6 are satisfied. In this case, there are values of α and ε such that $\mathcal{K}_\alpha(L) = \mathcal{W}_\varepsilon(\Gamma, L).$

Definition (cellular witness map). Suppose that α and ε are selected as in Theorems 5 and 6. The witness map $F_{\mathcal{W}} : |\mathcal{A}_\alpha(L)| \rightrightarrows |\mathcal{A}_\alpha(L)|$ for the fuzzy witness complex $\mathcal{W}_\varepsilon(\Gamma, L)$ is the CMM

$$(15) \quad F_{\mathcal{W}}(x) := \bigcap_{A_i \in \mathcal{A}_\alpha(L) : x \in A_i} \{A_j \in \mathcal{A}_\alpha(L) : \exists x_t \in W_\varepsilon(\Gamma, l_i) \text{ s.t. } x_{t+1} \in W_\varepsilon(\Gamma, l_j)\}.$$

Since by hypothesis the nerve $\mathcal{K}_\alpha(L) = N(\mathcal{A}_\alpha(L))$ is also the witness complex, the CMM

$F_{\mathcal{W}}$ induces an SMM

$$\mathcal{F}_{\mathcal{W}} : \mathcal{W}_{\varepsilon}(\Gamma, L) \rightrightarrows \mathcal{W}_{\varepsilon}(\Gamma, L)$$

in precisely the same way that $\mathcal{F}_{\mathcal{K}}$ was induced by $F_{\mathcal{A}}$, namely by (5). In other words, a simplex $\tau \in \mathcal{F}_{\mathcal{W}}(\sigma)$ whenever there are witnesses to σ that have images, under the temporal ordering of Γ , that are witnesses to τ . Indeed, the hypotheses of Theorem 5 imply that each nonempty image simplex $\tau \in \mathcal{F}_{\mathcal{W}}(\sigma)$, which is automatically in $\mathcal{K}_{\alpha}(L)$ since $A_{\tau} \neq \emptyset$, is also in the witness complex, since $\mathcal{K}_{\alpha}(L) \subset \mathcal{W}_{\varepsilon}(\Gamma, L)$. Thus, to guarantee that the image is in $\mathcal{W}_{\varepsilon}(\Gamma, L)$, we need Γ to be sufficiently dense on the α -shape ($\delta \leq \varepsilon/2$). Under the additional conditions of Theorem 6, the complexes $\mathcal{K}_{\alpha}(L)$ and $\mathcal{W}_{\varepsilon}(\Gamma, L)$ coincide, and we can view the witness map as having the domain $\mathcal{W}_{\varepsilon}(\Gamma, L)$ as well. Thus, to guarantee that the domain is well defined, we need the landmarks to be more or less uniformly spaced (β not too small), and each point in Γ to be not too far from a landmark (M not too large).

The fact that the α - and witness complexes coincide gives us hope that $F_{\mathcal{W}}$ will carry the same information about homology as the outer approximation $F_{\mathcal{A}}$. The following theorem ensures that this indeed is the case when the original map f satisfies a Lipschitz condition on the grid.

Theorem 7. *Suppose that $Y = |\mathcal{A}_{\alpha}(L)|$ is compact and f is Lipschitz on Y with constant c . Then if Γ is δ -dense on Y and $\delta \leq \frac{1}{2}\varepsilon \min\{1, \frac{1}{c}\}$, $F_{\mathcal{W}}$ is an outer approximation of f .*

Proof. We need to show that for any $y \in Y$, $f(y) \in F_{\mathcal{W}}(y)$. Note that any such $y \in A_i$ for some α -cell A_i and $f(y) \in A_j$ for some other α -cell A_j . We need to show that $A_j \subset F_{\mathcal{W}}(A_i)$ or, specifically, that there is an $x_t \in \Gamma$ such that $x_t \in W_{\varepsilon}(\Gamma, l_i)$ and $x_{t+1} = f(x_t) \in W_{\varepsilon}(\Gamma, l_j)$, where l_i and l_j are the landmarks associated with the α -cells A_i and A_j , respectively. Since Γ is δ -dense, it follows that there is $x \in \Gamma$ with $d(x, y) \leq \delta$. Thus, x is at most δ closer to any landmark than y (whose closest landmark is l_i),

$$(16) \quad d(x, L) \geq d(y, l_i) - \delta,$$

and consequently

$$(17) \quad d(x, l_i) \leq d(x, y) + d(y, l_i) \leq d(x, L) + 2\delta.$$

Since $2\delta \leq \varepsilon$, it follows that $x \in W_{\varepsilon}(\Gamma, l_i)$. In addition, since $d(f(x), f(y)) \leq cd(x, y)$ and $2c\delta \leq \varepsilon$, the same reasoning as in (17) leads to $f(x) \in W_{\varepsilon}(\Gamma, l_j)$.

Note that the points y and $f(y)$ may be in multiple α -cells, but the construction above applies to each cell, and so the conclusion is unaffected. ■

We have shown that, under the conditions of Theorems 5–7,

- the witness complex computed from data has the same homology as the union of a set of α -cells that cover the data, and
- when viewed as a multivalued map on \mathbb{R}^n , $F_{\mathcal{W}}$ is an outer approximation of the dynamical system f .

Since the cellular map $F_{\mathcal{W}}$ is semicontinuous (recall section 2.2), we know that whenever it is acyclic, then it is an enclosure of f . In this case, the acyclic carrier theorem implies that the induced map on homology can be computed from any continuous selector to the witness map [Mun84, Thm. 13.3]. However, note that acyclicity cannot be guaranteed; it must be checked when the map is numerically constructed.

3.4. Computing the map on homology. In our approach, all of the information about the topology of the invariant set $\Lambda \subset X$ is contained in the simplicial complex $\mathcal{W} = \mathcal{W}_\varepsilon(\Gamma, L)$, so our computation of the map f_* , the action induced by f on the homology groups, relies heavily on this simplicial complex. We begin by recalling the notion of a chain map.

A chain map from one simplicial complex to another consists of a homomorphism between the vertex sets, a homomorphism between the edge sets, etc., each of which commutes with the boundary operator. Commutation implies, for example, that the boundary of the image of a k -simplex is mapped to the image of the boundary of the k -simplex. An important feature of a chain map, φ , is that it induces a well-defined map in homology, φ_* [Mun84].

Our strategy in calculating f_* is to pick an appropriate chain map, φ , so that f_* coincides with φ_* . That is, we select $\varphi : \mathcal{W} \rightarrow \mathcal{W}$ to be a *chain selector*, so that $\varphi(\sigma) \in \mathcal{F}_\mathcal{W}(\sigma)$ for each $\sigma \in \mathcal{K}$. Such a selector can easily be constructed using a piecewise linear map between the topological realizations of the simplicial complexes. That is, we consider $|\varphi| : |\mathcal{W}| \rightarrow |\mathcal{W}|$. If φ is a chain selector for the SMM $\mathcal{F}_\mathcal{W}$, then it follows that $|\varphi|$ is a continuous selector for $F_\mathcal{W}$, and hence $\varphi_* = f_*$.

In summary, the strategy is as follows. To compute the Conley index of an isolated invariant set, it is sufficient to construct a CMM $F_\mathcal{A}$ that encloses f . Yet computing the α -grid on a given data set and its associated CMM $F_\mathcal{A}$ can be computationally expensive. In this section, we have shown how to construct a sparser simplicial complex—the witness complex $\mathcal{W}_\varepsilon(\Gamma, L)$. In addition, we define two associated multivalued witness maps—the cellular map, $F_\mathcal{W}$, defined implicitly on $|\mathcal{A}_\alpha(L)|$, and the combinatorial simplicial map, $\mathcal{F}_\mathcal{W}$, defined on the finitely determined complex $\mathcal{W}_\varepsilon(\Gamma, L)$. Moreover, when the hypotheses of the theorems in this section can be verified, $\mathcal{W}_\varepsilon(\Gamma, L) = \mathcal{K}_\alpha(L)$ and $F_\mathcal{W}$ encloses the dynamical system f . Then any continuous selector for this map will capture the homomorphism on homology f_* . We get this continuous selector of $F_\mathcal{W}$ by constructing a continuous selector of $\mathcal{F}_\mathcal{W}$ and taking its geometric realization.

4. An example: The Hénon map. The procedure for putting the mathematics of the previous sections into practice on time-series data from a dynamical system is as follows:

1. Given a time series $\Gamma = \{x_0, \dots, x_{T-1}\} \subset X$, which we assume lies near an invariant set $\Lambda \subseteq X$, select a set of landmarks, $L = \{l_0, \dots, l_{\ell-1}\}$, that are evenly distributed across Λ (cf. page 1286 and Appendix C). These landmarks will be the vertices of a simplicial complex.
2. Choose a value for the ε parameter that satisfies the requirement (14) for Theorem 6 and use witness/landmark relationships to simultaneously define a simplicial complex $\mathcal{W}_\varepsilon(\Gamma, L)$ and an SMM $\mathcal{F}_\mathcal{W}$. Note that, even though we never need to construct an α -complex, (14) implies that there is an α that has the same homology as \mathcal{W} .
3. Pick a subset of L as a starting guess for an isolated invariant set and use Algorithms 11, 12, and 14 of Appendix C to find an index pair $(|N|, |E|)$ for f . There are many different strategies for choosing the initial guess; if one is attempting to find periodic orbits, for instance, it makes sense to search for recurrent points in the time series and use the nearest landmark as the starting point for the algorithms. The important property is that the guess should be a subset of its period-image under the SMM.

4. Use a chain selector for \mathcal{F}_W to calculate $f_* : H_*(|N|, |E|) \rightarrow H_*(|N|, |E|)$.
 In the rest of this section, we illustrate this procedure on Hénon's classic map [Hén76]:

$$(18) \quad f(x, y) = (y + 1 - 1.4x^2, 0.3x).$$

This map has an invariant set Λ that is an attractor, and we generate a trajectory Γ of length $T = 10^5$ that starts from the initial condition $z_0 = (-0.4, 0.3)$, near Λ . The trajectory is shown in Figure 5. As a simple test of the witness map technique presented in the previous section, we use this trajectory to verify the trivial fact that f has a fixed point. We will assume that Γ is an exact trajectory of (18), making no claim that our computation is rigorous. The latter could, at least in principle, be done using interval arithmetic. Given (18), of course, a simple calculation shows that this system has *two* fixed points. Our goal is to find one of those fixed points using only the time series Γ . Note that this is a proof-of-concept example and not an exhaustive exploration of the parameter space of the algorithm. Moreover, it is simple enough that the homology calculations can be carried out by hand.

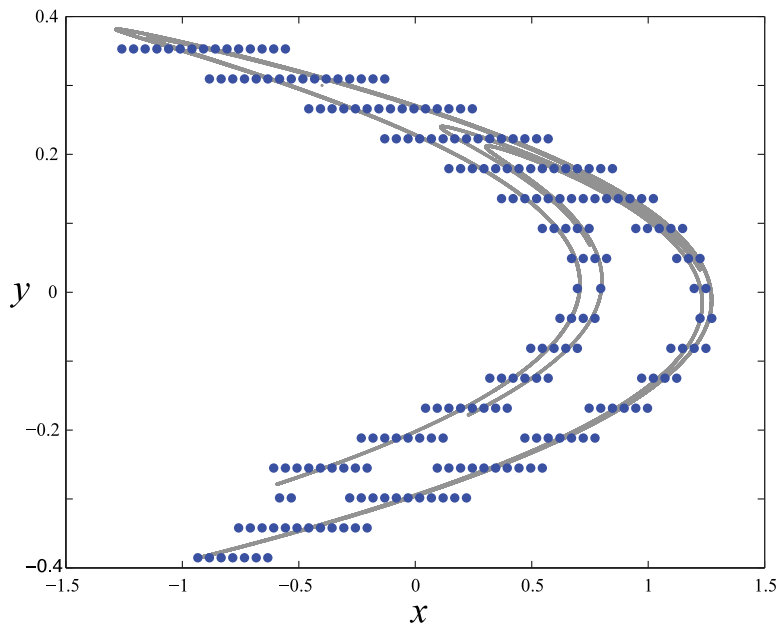


Figure 5. Time-series data, Γ , comprising 10^5 iterates of the Hénon map (18) (gray points) from an initial condition near the attractor, and the set L of 216 evenly spaced landmarks (blue circles) that approximate the attractor from which Γ was sampled. Note that the vertical and horizontal scales are different.

We begin by selecting a set of landmarks to approximate the attractor Λ . Again, as a proof of principle, we simply space these landmarks evenly within the bounding box of the orbit $[-1.5, 1.5] \times [-0.4, 0.4]$. With the goal of reflecting the structure of the attractor and yet having significantly fewer landmarks than points on the orbit ($\ell \ll T$), we use a hexagonal grid with spacing $\beta = 0.05$. Indeed, retaining only those landmarks that are within β of a time-series point gives $\ell = 216$ landmarks (so $\ell \sim \sqrt{T}$), as shown in Figure 5. This has the effect of distributing the landmarks across the attractor with enough resolution to detect some of its fractal structure.

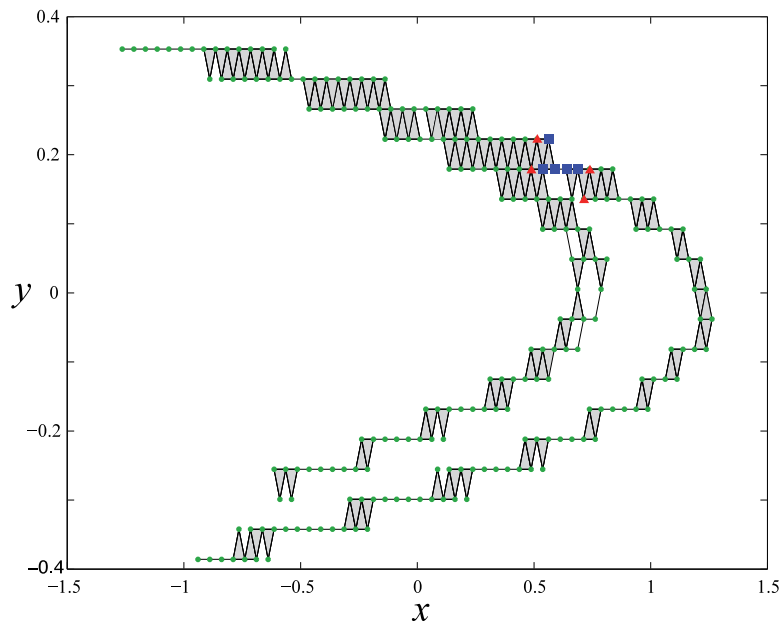


Figure 6. Witness complex and an index pair (N, E) for the fixed point of the Hénon map. The points correspond to the landmarks, L , the blue squares represent the isolating neighborhood, $N \setminus E$, and the red triangles are the exit set E . The witness complex corresponds to the landmarks, the lines, and the gray triangles.

The next step in the process is to define the witness-landmark relation R_ε of (13). Although we do not need to build an α -complex, it is possible to do so when the requirement (14) is satisfied. For example, given the hexagonal geometry, the α -complex will be a clique complex when $\alpha < \beta/2$, i.e., when it is trivially totally disconnected, or when $\alpha \geq \beta/\sqrt{3}$, the distance from a vertex to the center of the equilateral triangle of side β . Similarly, by construction, every data point is in one of the equilateral triangles formed from the landmarks, i.e., $M \leq \frac{\beta}{\sqrt{3}}$. The requirement (14) is then satisfied if we choose $\alpha = \beta/\sqrt{2}$, and

$$\varepsilon \leq \left(\frac{1}{\sqrt{2}} - \frac{1}{\sqrt{3}} \right) \beta.$$

Recall that, given an $\varepsilon \geq 0$, a pair $(x_t, l_j) \in R_\varepsilon \subset \Gamma \times L$ ($x_t \in W_\varepsilon(\Gamma, l_j)$), according to (13), if and only if $d(x_t, l_j) \leq d(x_t, L) + \varepsilon$. This witness relationship serves to define a clique complex $\mathcal{W}_\varepsilon(\Gamma, L)$ by (12): the edge $\langle l_i, l_j \rangle \in \mathcal{W}_\varepsilon$ if and only if $W_\varepsilon(\Gamma, l_i) \cap W_\varepsilon(\Gamma, l_j) \neq \emptyset$. By varying ε up to the bound above and looking at the resulting complexes, we finally selected $\varepsilon = 0.005$; this is large enough so that the complex is connected but small enough that the shape of the complex still reflects the primary fold in the attractor. The resulting complex is shown in Figure 6.

Having built the witness complex, we then construct the CMM $F_{\mathcal{W}} : |\mathcal{A}_\alpha(L)| \rightrightarrows |\mathcal{A}_\alpha(L)|$ using the witness-landmark relationships, as described by (15). Next we use the time series to search for an isolating neighborhood for $F_{\mathcal{W}}$. To apply Algorithm 12, we need a guess for an isolating neighborhood. For a periodic orbit, this can be found by looking for nearly recurrent points in the time series [LK89]. Consequently, to choose an initial guess for the

purpose of finding a fixed point, we can simply search for a time t that minimizes $d(x_t, x_{t+1})$. Following this approach, we find that $x_{39,436} = (0.6313, 0.1894)$ is a good candidate—indeed, it is close to the analytical fixed point $(x^*, y^*) \approx (0.6313544771, 0.1894063431)$. The cell of the landmark nearest to this point gives a useful initial guess for the isolating neighborhood for Algorithm 12. This isolating neighborhood is then used as the input of Algorithm 14 to obtain an index pair (N, E) for the fixed point of f . The result is shown in Figure 6.

The final step is to calculate the Conley index of the isolated invariant set that is the invariant part of $N \setminus E$. From this index, we can then infer the existence of a fixed point. We begin by taking a close look at the map \mathcal{F}_W restricted to the index pair (N, E) . The index pair is shown in Figure 6. Recall that F_W is a map that is constant on α -cells. Though we do not need to compute these α -cells, visualizing them—as in the sketch shown in Figure 7—helps in understanding the various multivalued maps involved in this process. In Figure 7, $N = \{A_1, \dots, A_9\}$ and $E = \{A_1, A_2, A_7, A_8\}$. The blue and red landmarks are the nexuses of the α -cells that make up $N \setminus E$ and E , respectively.

In this example, the map F_W restricted to the index pair can be described by the following transition matrix:

$$(19) \quad S = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \end{pmatrix},$$

where $S_{ij} = 1$ if and only if $A_j \subset F_W(\text{int}(A_i))$; i.e., there is a witness of the landmark associated with A_i whose image under the shift map is a witness of the landmark associated with A_j . For example, $F_W(A_6) = A_2 \cup A_3 \cup A_8 \cup A_9$. Geometrically, this image is a disk and is thus acyclic. Indeed, it is straightforward, if tedious, to verify from (19) that each cell maps to an acyclic set under the witness map $F_W|_N$. Moreover, by (15), whenever $x \in A_i \cap A_j$, it has an image that is the intersection of the images of the individual cells. In this way, one can compute the images under \mathcal{F}_W of the 12 one-dimensional simplices in the index pair. These images, inferred from (19), can be verified to be acyclic when restricted to N . Finally only one of the two-simplices in N remains in N under the map, namely $\mathcal{F}_W(\langle l_3, l_4, l_9 \rangle) = \langle l_6, l_7 \rangle$; this image has no homology. Consequently, the map F_W is acyclic on the index pair. This condition is sufficient for our limited purpose here of verifying that f has a fixed point in N . More generally, the acyclicity of the witness maps on the entire complex could easily be automated, as is done for cubical complexes in the software package “CHomP” [Mis14].

Now, we want to represent this index pair with a corresponding simplicial complex $\mathcal{W}(N, E)$, specifically, the witness complex associated with the landmarks $\{l_1, \dots, l_9\}$, corresponding to α -cells $\{A_1, \dots, A_9\}$. The witness complex $\mathcal{W}(N, E)$, which is also the α -complex in this case, is pictured in Figure 7. In order to compute the Conley index, we need a simplicial complex that represents the quotient space N/E . Since the α -cells A_1, A_2, A_7 , and A_8 make up the

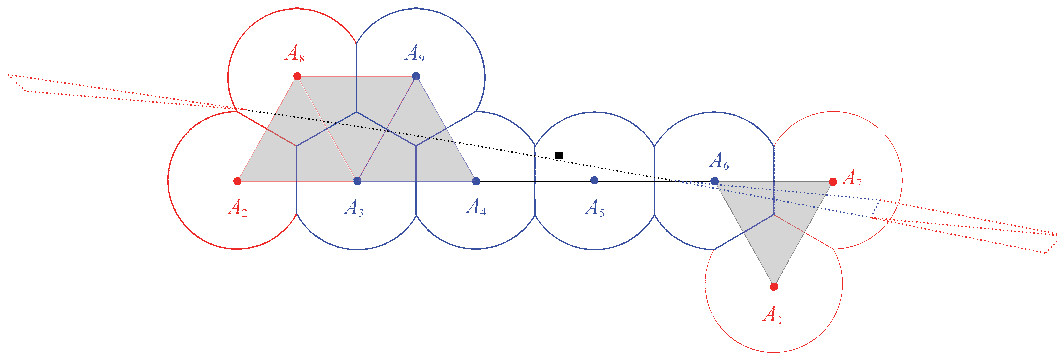


Figure 7. The witness complex (gray triangles, black lines, and dots) for the index pair (N, E) for the landmarks from Figure 6. In this case, the witness complex is also the nerve of the α -cells (shown as truncated red and blue spheres in the figure) since (14) is satisfied. Dashed lines show the analytical image of the simplicial complex under the Hénon map (18). The black square is the fixed point.

exit set E , we make the identification

$$l_1 \cong l_7 \cong l_2 \cong l_8 := E.$$

The resulting simplicial complex is shown in Figure 8. The homology of the simplicial complex $\mathcal{K}(N, E)$ can easily be computed by hand in this example. In particular, the quotient space N/E consists of a single connected component, so $H_0(N, E) = \mathbb{Z}_2$. The quotient space has a single, nonbounding cycle,

$$\sigma = \langle E, l_3 \rangle + \langle l_3, l_4 \rangle + \langle l_4, l_5 \rangle + \langle l_5, l_6 \rangle + \langle l_6, E \rangle,$$

so $H_1(N, E) = \mathbb{Z}_2$.

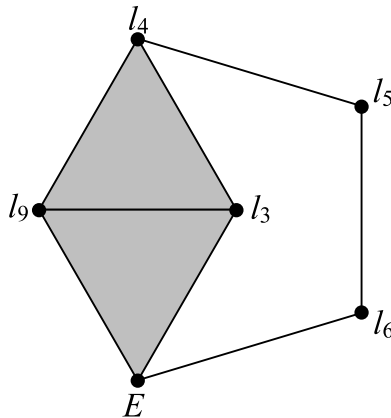


Figure 8. The quotient simplicial complex $N \setminus E$, where N is the simplicial complex shown in Figure 7 and $E = \{l_1, l_2, l_7, l_8\}$.

We are now ready to compute the Conley index of the isolating neighborhood $N \setminus E$. Specifically, we compute $f_* : H_*(N, E) \rightarrow H_*(N, E)$. As described in the previous section,

the homology of f is equivalent to the homology of φ , where φ is a chain selector for \mathcal{F}_W . Therefore, in order to compute the Conley index, we need to find $\varphi_*([\sigma]) := [\varphi(\sigma)]$.

The chain selector φ is defined inductively by first determining the image of each vertex in $\mathcal{K}(N, E)$ (to be enclosed by \mathcal{F}_W) and then determining the image of each edge so that φ commutes with the boundary operator. In addition, recall that \mathcal{F}_W on the quotient space must be an enclosure of the index map f_N . We begin with the initial assignment of vertices:

$$\begin{array}{c|cccccc} \langle i \rangle & \langle E \rangle & \langle l_3 \rangle & \langle l_4 \rangle & \langle l_5 \rangle & \langle l_6 \rangle & \langle l_9 \rangle \\ \hline \varphi_0(\langle l_i \rangle) & \langle E \rangle & \langle l_6 \rangle & \langle l_5 \rangle & \langle l_3 \rangle & \langle l_9 \rangle & \langle l_6 \rangle \end{array}$$

To compute $\varphi(\sigma)$, we need to find the images of the edges in σ as well. In order for φ to be a chain selector for \mathcal{F}_W , the image of each edge, τ , must be a subset of $\mathcal{F}_W(\tau)$. Furthermore, φ must commute with the boundary operator, so we need $\varphi_0 \circ \partial_1 = \partial_1 \circ \varphi_1$. Those two conditions yield the following edge assignments:

$$\begin{array}{c|cccccc} \tau & \langle E, l_3 \rangle & \langle l_3, l_4 \rangle & \langle l_4, l_5 \rangle & \langle l_5, l_6 \rangle & \langle l_6, E \rangle \\ \hline \varphi_1(\tau) & \langle E, l_6 \rangle & \langle l_6, l_5 \rangle & \langle l_5, l_4 \rangle + \langle l_4, l_3 \rangle & \langle l_3, l_9 \rangle & \langle l_9, E \rangle \end{array}$$

It follows that $\varphi_*([\sigma]) = [\sigma]$ (in \mathbb{Z}_2 homology). Since this map is not nilpotent, then it is not in the shift equivalence class of $[0]$ and thus the invariant set $S = \text{inv}(N \setminus E, f) \neq \emptyset$ [KMM04, Thm. 10.91].

With this very simple example, we have illustrated that the witness complex and the associated witness map can be used to compute the Conley index for a simple isolating neighborhood. We plan in the future to apply this method to more complex and higher-dimensional dynamics.

5. Conclusions and future work. Computational topology is a powerful way to analyze time-series data from dynamical systems. Existing approaches to the approximation of a dynamical system on algebraic objects construct multivalued maps from the time series using cubical discretizations and then use those maps to compute the Conley indices for isolated invariant sets of cubes. The approach described in this paper, by contrast, discretizes the dynamics using a simplicial complex that is constructed from a witness-landmark relationship. A natural discretization like this, whose cell geometry is derived from the data, is more parsimonious and thus potentially more computationally efficient than a cubical complex. We then use the temporal ordering of the data to construct a map on this simplicial witness complex that we call the *witness map*. Under the conditions established in section 3.3, this witness map gives an outer approximation of the dynamics and thus can be used to compute the Conley index of isolated invariant sets in the data.

As a proof of concept, we applied our methods to data from the classic Hénon map and located an isolating neighborhood for a fixed point of this dynamical system. There are many other potential applications in the study of dynamical systems. Our approach could also be used to find periodic orbits as well as connecting orbits between them—a strategy that ultimately leads to rigorous verification of chaotic dynamics.

An important question that we leave open is “Can one develop rigorous computational methods for multivalued maps based on simplicial complexes?” While interval arithmetic has

been done for cubical complexes [DFT08], an approach based on the selection of an appropriate value of ε to account for finite precision arithmetic might be appropriate.

In the future we also hope to explore the application of these techniques to scalar time-series data sampled from a dynamical system. In this case, delay-coordinate embedding [PCFS80, Tak81, SYC91] can be used to create a trajectory Γ of the form required by our methods. Of course, noise becomes an issue in any consideration of experimental data. In the case of bounded noise, it may be possible to turn our ε parameter to advantage—in the same spirit as in [MMRS99], where the size of the cells in the multivalued map is chosen to account for the experimental error.

Our techniques may also have significant impact in the numerical simulation of differential equations. In [MM95], numerical integration, while keeping track of the magnitude of round-off error, is used to prove that there is chaos present in the Lorenz equations [Lor63]. A key step in this proof is showing that one can construct a multivalued map that is truly an outer approximation of a given function f . Theorem 7 indicates that our techniques are appropriate for these types of proofs. The computational efficiency that the data-driven discretization confers upon the witness-map construction process should allow this approach to scale well with dimension, so it is likely that constructions based on this map could be used to generate computer-based proofs about high-dimensional differential equations. This would be a significant advance in the field.

A large body of research in the field of computational topology has revolved around the concept of *persistence* [DE95, EH08, Ghr08, Rob99]. The idea behind topological persistence is that many computations in this field depend upon simple scale parameters. For example, the α -complex of a point cloud depends upon the parameter α , the fuzzy witness complex depends upon the parameter ε , etc. It makes sense, then, to perform these calculations over a range of parameter values and to search for intervals where the topological properties remain constant. This was the rationale for the choice of ε in section 4. A major area for future research is the development of a theory of persistence in the context of Conley index theory. We believe that contribution described in this paper is a significant step in this direction.

Appendix A. Simplicial complexes for discrete data. An *abstract simplicial complex* \mathcal{K} is a collection of finite, ordered subsets $\sigma = \langle l_{i_0}, \dots, l_{i_k} \rangle$ in the power set, 2^L , of a set of vertices L such that if $\tau \leq \sigma$ is a “face” of σ (it contains only vertices also in σ), then it is also a simplex in \mathcal{K} . The vertices, $\langle l_i \rangle$, are zero-simplices; edges, $\langle l_i, l_j \rangle$, are one-simplices, etc. The empty set is a face of every simplex. For points $l_i \in \mathbb{R}^n$, the geometrical realization of a simplex, $|\sigma|$, is the convex hull of its vertices. A *geometrical simplicial complex* is an abstract complex \mathcal{K} such that each intersection $|\sigma| \cap |\tau|$ of simplices in \mathcal{K} is a face of both [Ede95].

A simplicial complex \mathcal{L} is a subcomplex of \mathcal{K} if every simplex in \mathcal{L} is in \mathcal{K} . The k -skeleton of \mathcal{K} is the subcomplex containing all simplices of dimension k or less. Thus a one-skeleton is the graph formed from the vertices and edges.

A *clique* (or *flag*) complex is the maximal complex with a given set of edges [Zom10] (called a “lazy” complex by [dSC04]). Thus a clique complex is determined by its one-skeleton.

The *nerve* $N(\mathcal{A})$ of a collection of sets \mathcal{A} is an abstract simplicial complex constructed from finite intersections. The vertices, $l_i \in L$, are labels for the elements $A_i \in \mathcal{A}$, and a simplex

$\sigma = \langle l_0, \dots, l_k \rangle$ is in the nerve if the $k + 1$ corresponding sets have nonempty intersection:

$$N(\mathcal{A}) = \left\{ \sigma : \bigcap_{l_i \in \sigma} A_i \neq \emptyset \right\}.$$

Thus each vertex in the nerve corresponds to the label l_i of the cell A_i , and each edge $\langle l_i, l_j \rangle$ to a nonempty intersection $A_i \cap A_j$, etc. It was shown by Borsuk and Weil that, in certain cases, the nerve has the same homology or homotopy type as the geometrical realization of the collection \mathcal{A} .

Lemma 8 (nerve lemma [Bor48, Wei52, BT82, Hat02]). *Let \mathcal{A} be a collection of closed sets such that every finite intersection between its members is either empty or contractible. Then $N(\mathcal{A})$ has the same homotopy type as $|\mathcal{A}|$.*

There are many natural ways of defining simplicial complexes for a finite point set $L = \{l_1, l_2, \dots, l_\ell\} \subset \mathbb{R}^n$. If $\mathcal{A} = \{B_r(l) : l \in L\}$ is the collection of closed radius- r balls (8) around the set of landmarks, then the Čech complex, $\mathcal{C}_r(L)$, is the nerve of \mathcal{A} . Since the balls are convex subsets of \mathbb{R}^n , the nerve lemma implies that \mathcal{C} has the same homotopy type as $|\mathcal{A}|$. The sequence of Čech complexes is nested: $\mathcal{C}_r(L) \subseteq \mathcal{C}_{r'}(L)$ when $r < r'$. A similar complex, the Rips (or Vietoris–Rips) complex, $\mathcal{R}_r(L)$, consists of all simplices whose vertices are pairwise within a distance r of each other:

$$\mathcal{R}_r(L) = \{ \sigma : d(l, l') \leq r \forall l, l' \in \sigma \}.$$

Since this complex is determined by its edges, it is a clique complex. Rips complexes are also nested as r grows, and, moreover, they are interleaved with Čech complexes:

$$\mathcal{R}_{r'}(L) \subset \mathcal{C}_r(L) \subset \mathcal{R}_{2r}(L)$$

whenever $r' < r\sqrt{2(n+1)/n}$ [Ghr08, dSG07]. This gives a relation between the persistent homologies of the family of Rips complexes and the family of Čech complexes.

The Voronoi diagram $\mathcal{V}(L) = \{V_l : l \in L\}$ is the covering of \mathbb{R}^n by the cells

$$V_l = \{x \in \mathbb{R}^n : d(x, l) \leq d(x, l') \forall l' \in L\}.$$

Note that each Voronoi cell is convex, since it is the intersection of half-spaces, and two such cells either are disjoint or meet on a portion of their boundaries. The associated simplicial complex is the Delaunay complex $\mathcal{D}(L) = N(\mathcal{V}(L))$, the nerve of the Voronoi diagram. When the points L are in “general position” (no more than $n + 1$ points lie on any $(n - 1)$ -sphere), $\mathcal{D}(L)$ is a geometrical complex [Ede95]. Since this is generically true, general position can be achieved by almost any, arbitrarily small, perturbation of the points in L . Thus it is common to assume L is in general position.

The cells in the α -diagram are the intersection of the Voronoi cells with a closed ball of radius α about a vertex:

$$\mathcal{A}_\alpha(L) = \{V_l \cap B_\alpha(l) : l \in L\}.$$

The corresponding nerve is the α -complex, $\mathcal{K}_\alpha(L) = N(\mathcal{A}_\alpha(L))$. An alternative characterization is that $\sigma \in \mathcal{K}_\alpha(L)$ if there exists a ball $B_r(x_0)$ with $r \leq \alpha$ that contains no vertices

in its interior, $\text{int}(B_r(x_0)) \cap L = \emptyset$, but for which $\sigma \subset \partial B_r(x_0)$. The boundary of such a ball is called a circumsphere for σ . For the Euclidean metric, each α -cell is convex, since it is the intersection of convex sets. Thus the nerve lemma implies that $|\mathcal{K}_\alpha(L)|$ is homotopy equivalent to $|\mathcal{A}_\alpha(L)|$. Note that the α -complex is the intersection of the Čech complex and the Delaunay complex,

$$\mathcal{K}_\alpha(L) = C_\alpha(L) \cap \mathcal{D}(L),$$

so it is a subcomplex of each. Moreover, as $\alpha \rightarrow \infty$, $\mathcal{K}_\alpha(L) \rightarrow \mathcal{D}(L)$.

A subset $A \subset X$, is a *deformation retract* of X if there exists a continuous map $r : X \times I \rightarrow X$ satisfying (7). The restriction $\rho : X \rightarrow A$ defined by $\rho(\cdot) = r(\cdot, 1)$ is then a *retraction* of X onto A . If A is a deformation retract of X , then it is homotopy equivalent to X . More generally, two spaces A and B are homotopy equivalent if and only if there are a space X and embeddings $a : A \rightarrow X$ and $b : B \rightarrow X$ such that both $a(A)$ and $b(B)$ are deformation retracts of X [Hat02, Cor. O.21].

In fact the α -shape, $|\mathcal{K}_\alpha(L)|$, is a deformation retraction of the α -grid $|\mathcal{A}_\alpha(L)|$ [Ede95]. There are two parts to Edelsbrunner's result, and we give only a brief discussion of the ideas in his paper.

Lemma 9. *For any $\alpha \geq 0$ and any finite set of landmarks $L \subset \mathbb{R}^n$ in general position, $|\mathcal{K}_\alpha(L)| \subset |\mathcal{A}_\alpha(L)|$.*

This follows because when a collection of α -cells mutually intersect, i.e., when $A_\sigma \neq \emptyset$, they must do so at a point x in the interior of $|\sigma|$ (viewed as a subset of its spanning k -plane). The proof proceeds by induction (it is easy for 0-simplices) and uses the facts that the union $S(\sigma) = \cup_{l \in \sigma} A_l$ is star-convex, relative to x , and that $|\sigma|$ is itself convex. The implication is that $|\sigma| \subset S(\sigma)$ for each simplex in $\mathcal{K}_\alpha(L)$. The result is used to construct the deformation retract.

Lemma 10. *For any $\alpha \geq 0$ and any finite set of landmarks $L \subset \mathbb{R}^n$ in general position, $|\mathcal{K}_\alpha(L)|$ is a deformation retract of $|\mathcal{A}_\alpha(L)|$.*

The construction of a deformation retract is based on planes that are orthogonal to points on simplices that are on the boundary of $|\mathcal{K}_\alpha(L)|$ (for points in the interior, the deformation is the identity map). If σ is a k -simplex, then each point in its interior is the intersection of the k -plane containing $|\sigma|$ and the $(n-k)$ -plane that is its orthogonal complement. Convexity implies that these families of orthogonal $(n-k)$ -planes cover $|\mathcal{A}_\alpha(L)| \setminus |\mathcal{K}_\alpha(L)|$, and each point in this set lies in exactly one such plane. The deformation is defined as linear flow from the boundary of $|\mathcal{A}_\alpha(L)|$ to the boundary of $|\mathcal{K}_\alpha(L)|$. A consequence of this construction is that the deformation maintains membership in each α -cell: $r(A_i, t) \subset A_i$. This last property is a hypothesis for Lemma 2.

Appendix B. Discrete Conley index. A key tool in computational topology is the Conley index [Con78], which can be expressed in terms of the algebraic topology of a pair of sets that are acted upon by a map f . Given the Conley index of such a pair, one can sometimes prove the existence of fixed points, periodic orbits, and equivalence to shift dynamics for the dynamics of an invariant set. In this appendix we briefly recall the definition of the index and some related concepts; for more details see [Eas98, MM02, KMM04].

Given a homeomorphism $f : X \rightarrow X$, a set Λ is *invariant* if $f(\Lambda) = \Lambda$. The *maximal*

invariant set contained in a set K is

$$\text{inv}(K) = \{x \in K : f^t(x) \in K \forall t \in \mathbb{Z}\},$$

which, of course, could be empty. A compact set K is an *isolating neighborhood* if the subset that remains in K for all time is contained in its interior: $\text{inv}(K) \subset \text{int}(K)$. Similarly, a set S is an *isolated invariant set* of f when it is the maximal invariant set in the interior of some isolating neighborhood K : $S = \text{inv}(K) \subset \text{int}(K)$.

The computation of the Conley index relies on the construction of an *index pair* that gives rise to an isolating neighborhood [MM02].

Definition (index pair). A pair of compact sets $P = (N, E)$ with $E \subset N \subset X$ is called an index pair for $S = \text{inv}(N \setminus E)$ relative to f if it satisfies the following three properties:

- $\text{cl}(N \setminus E)$ is an isolating neighborhood.
- $f(E) \cap N \subseteq E$.
- $f(N \setminus E) \subset N$.

These three properties are illustrated in Figure 9. The first states that the set $K = \text{cl}(N \setminus E)$ is an isolating neighborhood that isolates some (possibly empty) invariant set S . The second property implies that once a trajectory enters E , it will not return to K before leaving the index pair entirely. The third property states that E contains the exit set of N ; that is, the images of points not in E must remain in N . It is possible to show that every isolated invariant set S has an index pair [FR00].

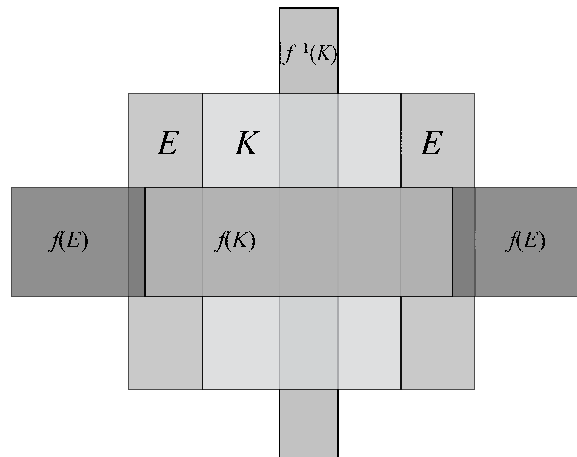


Figure 9. An index pair $N = K \cup E$ (a square) and E (two rectangles). The image of E either leaves N or remains in E . For this picture the isolating neighborhood K is guaranteed to contain a fixed point. A simple map with this index pair is $f(x, y) = (ax, by)$ with $a > 1 > b > 0$.

For any index pair $P = (N, E)$ there is a quotient space N/E with the equivalence relation $[x] = [y]$ if $x, y \in E$. A continuous *index map* f_P can be defined on N/E by

$$f_P([x]) = \begin{cases} f(x) & \text{if } x, f(x) \in N \setminus E, \\ [E] & \text{otherwise.} \end{cases}$$

A complication is that two index pairs P, P' for the same isolated invariant set can have topologically distinct quotient spaces and index maps that are not homotopic. However, any

such index maps induce *shift equivalent* maps on the homology group $H_*(N, E)$ of N relative to E . Shift equivalence is a less rigid version of conjugacy for noninvertible maps that was introduced by Robert Williams and used in the definition of the discrete Conley index by Franks and Richeson [FR00].

Definition (shift equivalence). *A pair of endomorphisms $f : X \rightarrow X$ and $g : Y \rightarrow Y$ are shift equivalent if there exist continuous maps $h : X \rightarrow Y$ and $k : Y \rightarrow X$ such that $h \circ f = g \circ h$ and $f \circ k = k \circ g$, and there exists an $m \in \mathbb{N}$ such that $h \circ k = g^m$ and $k \circ h = f^m$.*

Note that if f and g are homeomorphisms, they are shift equivalent if and only if they are conjugate, with the conjugacy defined by $c = h \circ f^{-m} = g^{-m} \circ h$, and $c^{-1} = k$.

The point is that for any two index pairs (N, E) , (N', E') that isolate the same invariant set, the maps on homology f_{N^*} and $f_{N'^*}$ are shift equivalent [FR00]. This shift equivalence class, $[f_{P^*}]_s$, is the *discrete homology Conley index* $\text{Con}(S, f)$ of the invariant set S isolated by K .

One of the fundamental advantages of the Conley index is its structural stability; for example, if K is an isolating neighborhood for f , then there is $\varepsilon > 0$ such that K is also an isolating neighborhood for \tilde{f} , whenever $\|f - \tilde{f}\|_\infty < \varepsilon$. Moreover, so long as an invariant set remains isolated by K , its Conley index does not change [MM02].

The simplest implication of a nontrivial Conley index is the *Wazewski property*: whenever $\text{Con}(S, f) \neq [0]$, then $S \neq \emptyset$ [KMM04, Thm. 10.91]. In addition, periodic orbits are guaranteed when the “Lefschetz number” is nonzero [KMM04, Thm. 10.46], and (for C^∞ maps) the topological entropy is positive whenever the shift equivalence class of f_P has spectral radius greater than one [Bak02].

Appendix C. Computing the Conley index. In order to use the Conley index to obtain information about a map f , we start by using a multivalued map $F_{\mathcal{A}}$ or $F_{\mathcal{W}}$ to locate isolating neighborhoods for f . Since the construction of the CMM is analogous to the cubical map of [KMM04], we borrow the presentation as well as relevant theorems and algorithms from that work and from [DFT08]. In most cases, the proofs of the theorems in this section are identical to those in the original citations if one simply substitutes the concept of a grid-cell for that of a cube. A thorough treatment of these results—with respect to any grid satisfying the first definition in section 2.1—can be found in [Mro99]. Our goal is to move beyond the cubical complexes used in previous work and devise a method to *efficiently* build an SMM that contains the same information as the CMM.

We begin by defining trajectories and invariant sets for the multivalued map, following [DJM04, DFT08].

Definition (combinatorial trajectory). *A combinatorial trajectory of $F_{\mathcal{A}}$ through $A \in \mathcal{A}$ is a bi-infinite sequence of cells, $\Gamma_A = (\dots, A^{(-1)}, A^{(0)}, A^{(1)}, \dots)$, such that $A^{(0)} = A$ and $A^{(n+1)} \subseteq F_{\mathcal{A}}(A^{(n)})$ for all $n \in \mathbb{Z}$.*

Definition (combinatorial invariance). *Given a CMM $F_{\mathcal{A}} : |\mathcal{A}| \rightrightarrows |\mathcal{A}|$, the combinatorial invariant part of $N \subset \mathcal{A}$ is defined by*

$$\text{inv}(N, F_{\mathcal{A}}) := \{A \in \mathcal{A} : \exists \text{ a trajectory } \Gamma_A \text{ for which } A^{(n)} \subset N \text{ for all } n \in \mathbb{Z}\}.$$

The following algorithm can be used to locate the combinatorial invariant part of a compact

set N .

Algorithm 11. *invariantPart*($N, F_{\mathcal{A}}$)

$S \leftarrow N$

repeat

$S' \leftarrow S$

$S \leftarrow F_{\mathcal{A}}(S) \cap S \cap F_{\mathcal{A}}^{-1}(S)$

until $S = S'$

return S

It is proved in [KMM04, Thm. 10.83]—in the context of cubical sets—that if N is finite, this algorithm terminates and returns $\text{inv}(N, F_{\mathcal{A}})$ (which could be empty). The extension to the cellular case is straightforward.

Associated with this notion of invariance, there is a property of isolation, which is defined using the following definition.

Definition (combinatorial neighborhood). *The combinatorial neighborhood of a set $S \subset \mathcal{A}$ is*

$$o(S) := \{B \in \mathcal{A} : B \cap S \neq \emptyset\}.$$

More plainly, the combinatorial neighborhood consists of S and all of the cells that touch its boundary. In order for a combinatorial invariant set to be isolated, it must be the invariant set of some neighborhood.

Definition (combinatorial isolating neighborhood). *A set $K \subset \mathcal{A}$ is a combinatorial isolating neighborhood if*

$$o(\text{inv}(K, F_{\mathcal{A}})) \subseteq K.$$

Given a guess, K , for such a neighborhood, we might be able to find an isolating neighborhood by *growing* it: if $K' = \text{inv}(o(K), F_{\mathcal{A}}) \subset K$, then K is isolating; otherwise we replace K by K' and repeat. For example, in section 4, where we are looking for a fixed point, we use the cell containing a nearly recurrent point as the initial guess. This leads to the algorithm of [DJM04, DFT08].

Algorithm 12. *growIsolating*($K, F_{\mathcal{A}}$)

while $\text{inv}(o(K), F_{\mathcal{A}}) \not\subseteq K$ **do**

$K \leftarrow \text{inv}(o(K), F_{\mathcal{A}})$

if $K \cap \partial|\mathcal{A}| \neq \emptyset$ **then**

return \emptyset

end if

end while

return K

If *growIsolating* is called with a combinatorial set $K \subset \mathcal{A}$ and a CMM $F_{\mathcal{A}}$, then it returns a combinatorial isolating neighborhood for $F_{\mathcal{A}}$ —or else it fails when K intersects the boundary of the grid \mathcal{A} . A sufficient condition for this not to occur is that $|\mathcal{A}|$ is itself an isolating neighborhood because then each cell that touches the boundary of $|\mathcal{A}|$ has a neighborhood whose invariant part is contained in $|\mathcal{A}|$.

An important point is that when K is isolating for $F_{\mathcal{A}}$, then under certain conditions, $|K|$

is isolating for any continuous selector f of $F_{\mathcal{A}}$.

Theorem 13. *Let $F_{\mathcal{A}} : |\mathcal{A}| \rightrightarrows |\mathcal{A}|$ be a CMM for f . Then if $K \subset \mathcal{A}$ is a combinatorial isolating neighborhood for $F_{\mathcal{A}}$, $|K|$ is an isolating neighborhood for f .*

This is essentially [KMM04, Thm. 10.87], generalized to the cellular case.

The computation of the Conley index begins with an isolating neighborhood K of a CMM, with the goal of finding a pair of sets (N, E) that satisfy the definition of an index pair. We compute these using the following algorithm.

Algorithm 14. *indexPair($K, F_{\mathcal{A}}$)*

```

 $S \leftarrow \text{inv}(K, F_{\mathcal{A}})$ 
 $C \leftarrow o(S) \setminus S$ 
 $E \leftarrow F_{\mathcal{A}}(S) \cap C$ 
repeat
   $E' \leftarrow E$ 
   $E \leftarrow F_{\mathcal{A}}(E) \cap C \cap E'$ 
until  $E = E'$ 
 $N \leftarrow S \cup E$ 
return  $(N, E)$ 

```

This is similar to Algorithm 10.86 in [KMM04] which was stated for cubical sets. It was proven there that if this algorithm is called with a combinatorial isolating neighborhood K and an outer approximation $F_{\mathcal{A}}$ of f , then the geometric realization of the pair it returns is an index pair for f . This proof can be adapted to the cellular map situation.

Given an index pair $(|N|, |E|)$ for f , the computation of the discrete Conley index reduces to finding a representative of the shift equivalence class $[f_{P*}]_s$ and its action on the relative homology groups, $H_*(|N|, |E|)$.

Appendix D. Computational complexity. To analyze the computational complexity of the approach proposed in this paper and compare it to that of the cubical grid version of [KMM04], one must consider both run time and memory use.

In the cubical grid case, all of the cells that are occupied by data points must be processed in order to compute the homology. The run time costs of this have two components. Determining whether an individual data point is in a particular cell in a d -dimensional cubical grid is a matter of evaluating $2d$ inequalities: the computational cost is $O(d)$. Constructing the multivalued map requires checking the images of the 2^d grid squares that touch the corner points of the occupied cells and iteratively expanding that set until there are no empty intersections [MMRS99]. This iterative expansion step can be a significant computational expense. Finding an isolated invariant set can require iteratively checking the forward and backward images of the neighboring cells in the grid (cf. Algorithm 12). This too can require significant computational effort.

The witness-complex approach sidesteps all of this complexity in two ways: first by using a *subset* of the data, and second by building a *simplicial* complex from those landmarks. The computational costs of this approach are balanced differently than in the cubical grid case: building the complex is harder, but using it is easier. In particular, constructing a witness complex involves calculating the distances between every point and every landmark, which has cost $O(\ell \log T)$ if there are T points and ℓ landmarks (using, e.g., a kd -tree algorithm).

But this computation parallelizes beautifully; moreover, $\ell \ll T$ in practice—indeed, that is the point of the “coarsening” inherent in the witness complex. Moreover, the dimension of each simplex is only high enough to cover the corresponding part of the invariant set, whereas all of the grid elements in the cubical case necessarily have the dimension of the ambient space. This means that Algorithm 12 has not only fewer cells to process in the simplicial case but also far fewer neighbors to check. For all of these reasons, the overall complexity in computing the homology of a witness complex is substantially lower than that of the cubical grid case. Note, too, that the cellular witness map is automatically an outer approximation if the conditions of Theorem 7 are satisfied.

The memory costs of the two approaches also arise in different ways. Informally speaking, in order to use a cubical grid to capture the dynamics with the same fidelity as a witness complex constructed from landmarks whose minimum spacing is β , one would need to use grid elements of size β/\sqrt{d} , where d is the dimension of the ambient space. The number of cells in this grid would be larger than the number of d -dimensional simplices in the corresponding witness complex. This effect, which holds even if one disregards empty grid cells, may not be significant in low dimensions and small data sets but can become an issue if the data are large and/or high-dimensional. Moreover, if the landmarks are spaced uniformly in time along the trajectory, that spacing—and the geometry of the witness complex—naturally adapts to the dynamics. Cubical grids do not share this advantageous property.

Another important difference arises in storing the complex in the computer’s memory. There are a number of extremely efficient ways to store information about which cells of a cubical grid are occupied by data points. The free-form nature of simplices would appear to make storing information about them (points, edges, faces, etc.) more of a challenge, but that cost can be mitigated by using creative algorithms. Note, for instance, that if one stores the results of the witness-landmark calculations mentioned above in the form of a linked list whose t th element contains a list of the landmarks that are witnessed by the t th data point, sorted in increasing order by distance, that data structure *contains all of the information one needs to describe the witness complex*. Algorithmic creativity can lower the expense of working with that data structure; we are currently investigating an approach that stores the witness relationships in bitmap data structures and uses them as “masks” (together with logical operations) to find landmarks that are shared between different sets of witnesses. And the clique assumption made here can be used to further streamline this search, since all one needs to consider is the edges.

While we have not provided a test of these claims about computational efficiency on a large set of high-dimensional data in this paper, we plan to do so in future work.

Finally, we would like to note that while building α -complexes is a computationally demanding task in high dimensions, we never actually *construct* an α -complex. The only roles of that construct in this work are as a vehicle for extending the proofs of [KMM04] to the simplicial case.

Acknowledgments. The authors would like to thank the referees for extensive comments that improved the paper. They also thank Robert Ghrist for pointing out the Dowker reference and Konstantin Mischaikow and Robert Easton for useful conversations.

REFERENCES

- [AKK⁺09] Z. ARAI, W. KALIES, H. KOKUBU, K. MISCHAIKOW, H. OKA, AND P. PILARCZYK, *A database schema for the analysis of global dynamics of multiparameter systems*, SIAM J. Appl. Dyn. Syst., 8 (2009), pp. 757–789.
- [Bak02] A.W. BAKER, *Lower bounds on entropy via the Conley index with applications to time series*, Topol. App., 120 (2002), pp. 333–354.
- [Bor48] K. BORSUK, *On the imbedding of systems of compacta in simplicial complexes*, Fund. Math., 35 (1948), pp. 217–234; also available online from <http://matwbn.icm.edu.pl/ksiazki/fm/fm35/fm35119.pdf>.
- [BT82] R. BOTT AND L.W. TU, *Differential Forms in Algebraic Topology*, Grad. Texts in Math. 82, Springer, Berlin, 1982.
- [Con78] C. CONLEY, *Isolated Invariant Sets and the Morse Index*, CBMS Reg. Conf. Ser. Math. 38, AMS, Providence, RI, 1978.
- [DE95] C.J.A. DELFINADO AND H. EDELSBRUNNER, *An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere*, Comput. Aided Geom. Design, 12 (1995), pp. 771–784.
- [DFT08] S. DAY, R. FRONGILLO, AND R. TREVIÑO, *Algorithms for rigorous entropy bounds and symbolic dynamics*, SIAM J. Appl. Dyn. Syst., 7 (2008), pp. 1477–1506.
- [DJM04] S. DAY, O. JUNGE, AND K. MISCHAIKOW, *A rigorous numerical method for the global analysis of infinite-dimensional discrete dynamical systems*, SIAM J. Appl. Dyn. Syst., 3 (2004), pp. 117–160.
- [Dow52] C.H. DOWKER, *Homology groups of relations*, Ann. of Math. (2), 56 (1952), pp. 84–95.
- [dS08] V. DE SILVA, *A weak characterisation of the Delaunay triangulation*, Geom. Dedicata, 135 (2008), pp. 39–64.
- [dSC04] V. DE SILVA AND E. CARLSSON, *Topological estimation using witness complexes*, in Proceedings of the 1st Eurographics Symposium on Point-Based Graphics, M. Alexa and S. Rusinkiewicz, eds., The Eurographics Association, Zürich, Switzerland, 2004, pp. 157–166.
- [dSG07] V. DE SILVA AND R. GHRIST, *Coverage in sensor networks via persistent homology*, Algebr. Geom. Topol., 7 (2007), pp. 339–358.
- [Eas98] R.W. EASTON, *Geometric Methods for Discrete Dynamical Systems*, Oxford University Press, Oxford, UK, 1998.
- [Ede95] H. EDELSBRUNNER, *The union of balls and its dual shape*, Discrete Comput. Geom., 13 (1995), pp. 415–440.
- [EH08] H. EDELSBRUNNER AND J. HARER, *Persistent homology: A survey*, in Surveys on Discrete and Computational Geometry: Twenty Years Later, Contemp. Math. 453, J.E. Goodman, J. Pach, and R. Pollack, eds., AMS, Providence, RI, 2008, pp. 257–282.
- [EM92] H. EDELSBRUNNER AND E.P. MÜCKE, *Three-dimensional alpha shapes*, in Proceedings of the 1992 Workshop on Volume Visualization, ACM, New York, 1992, pp. 75–82.
- [FBF77] J.H. FRIEDMAN, J.L. BENTLEY, AND R.A. FINKEL, *An algorithm for finding best matches in logarithmic expected time*, ACM Trans. Math. Software, 3 (1977), pp. 209–226.
- [FR00] J. FRANKS AND D. RICHESON, *Shift equivalence and the Conley index*, Trans. Amer. Math. Soc., 352 (2000), pp. 3305–3322.
- [Ghr08] R. GHRIST, *Barcodes: The persistent topology of data*, Bull. Amer. Math. Soc. (N.S.), 45 (2008), pp. 61–75.
- [Hat02] A. HATCHER, *Algebraic Topology*, Cambridge University Press, Cambridge, UK, 2002.
- [Hén76] M. HÉNON, *A two-dimensional mapping with a strange attractor*, Comm. Math. Phys., 50 (1976), pp. 69–77.
- [KMM04] T. KACZYNSKI, K. MISCHAIKOW, AND M. MROZEK, *Computational Homology*, Appl. Math. Sci. 157, Springer-Verlag, New York, 2004.
- [KMV05] W.D. KALIES, K. MISCHAIKOW, AND R. VANDERVORST, *An algorithmic approach to chain recurrence*, Found. Comput. Math., 5 (2005), pp. 409–449.
- [LK89] D.P. LATHROP AND E.J. KOSTELICH, *Characterization of an experimental strange attractor by periodic orbits*, Phys. Rev. A, 40 (1989), pp. 4028–4031.
- [Lor63] E.N. LORENZ, *Deterministic nonperiodic flow*, J. Atmospheric Sci., 20 (1963), pp. 130–141.

- [Mis14] K. MISCHAIKOW, *CHomP: Computational Homology Project*, <http://chomp.rutgers.edu> (2014).
- [MM95] K. MISCHAIKOW AND M. MROZEK, *Chaos in the Lorenz equations: A computer-assisted proof*, Bull. Amer. Math. Soc., 32 (1995), pp. 66–72.
- [MM02] K. MISCHAIKOW AND M. MROZEK, *Conley index theory*, in Handbook of Dynamical Systems, Volume 2, North-Holland, Amsterdam, 2002, pp. 393–460.
- [MMRS99] K. MISCHAIKOW, M. MROZEK, J. REISS, AND A. SZYMCZAK, *Construction of symbolic dynamics from experimental time series*, Phys. Rev. Lett., 82 (1999), pp. 1144–1147; also available online from <http://link.aps.org/doi/10.1103/PhysRevLett.82.1144>.
- [MMSR97] K. MISCHAIKOW, M. MROZEK, A. SZYMCZAK, AND J. REISS, *From Time Series to Symbolic Dynamics: An Algebraic Topological Approach*, Technical report, Georgia Institute of Technology, Atlanta, GA, 1997, <http://www.math.rutgers.edu/%7Emischaik/papers/tseries.ps>.
- [Mro99] M. MROZEK, *An algorithmic approach to the Conley index theory*, J. Dynam. Differential Equations, 11 (1999), pp. 711–734.
- [Mun84] J.R. MUNKRES, *Elements of Algebraic Topology*, Benjamin/Cummings, Menlo Park, CA, 1984.
- [PCFS80] N.H. PACKARD, J.P. CRUTCHFIELD, J.D. FARMER, AND R.S. SHAW, *Geometry from a time series*, Phys. Rev. Lett., 45 (1980), pp. 712–716.
- [Rob99] V. ROBINS, *Towards computing homology from finite approximations*, Topology Proc., 24 (1999), pp. 503–532; also available online from <http://at.yorku.ca/b/a/a/k/28.htm>.
- [SYC91] T. SAUER, J.A. YORKE, AND M. CASDAGLI, *Embedology*, J. Stat. Phys., 65 (1991), pp. 579–616.
- [Tak81] F. TAKENS, *Detecting strange attractors in turbulence*, in Dynamical Systems and Turbulence (Warwick, 1980), Lecture Notes in Math. 898, D. Rand and L.-S. Young, eds., Springer, Berlin, Heidelberg, 1981, pp. 366–381.
- [Wei52] A. WEIL, *Sur les théorèmes de Rham*, Comment. Math. Helv., 26 (1952), pp. 119–145.
- [Zom10] A.J. ZOMORODIAN, *Fast construction of the Vietoris-Rips complex*, Computers & Graphics-UK, 34 (2010), pp. 263–271.