

Exploring the Topology of Dynamical Reconstructions

Joshua Garland^{a,*}, Elizabeth Bradley^{a,b}, James D. Meiss^c

^a*Department of Computer Science, University of Colorado, Boulder CO USA*

^b*Santa Fe Institute, Santa Fe NM USA*

^c*Department of Applied Mathematics, University of Colorado, Boulder CO USA*

Abstract

Computing the state-space topology of a dynamical system from scalar data requires accurate reconstruction of those dynamics and construction of an appropriate simplicial complex from the results. The reconstruction process involves a number of free parameters and the computation of homology for a large number of simplices can be expensive. This paper is a study of how to compute the homology efficiently and effectively without a full (diffeomorphic) reconstruction. Using trajectories from the classic Lorenz system, we reconstruct the dynamics using the method of delays, then build a simplicial complex whose vertices are a small subset of the data: the “witness complex.” Surprisingly, we find that the witness complex correctly resolves the homology of the underlying invariant set from noisy samples of that set even if the reconstruction dimension is well below the thresholds for assuring topological conjugacy between the true and reconstructed dynamics that are specified in the embedding theorems. We conjecture that this is because the requirements for reconstructing *homology* are less stringent: a homeomorphism is sufficient—as opposed to a diffeomorphism, as is necessary for the full dynamics. We provide preliminary evidence that a homeomorphism, in the form of a delay-coordinate reconstruction map, may exist at a lower dimension than that required to achieve an embedding.

Keywords: Topology, Delay-Coordinate Embedding, Nonlinear Time-Series Analysis, Computational Homology, Witness Complex

1. Introduction

Topology is of particular interest in dynamics, since many properties—the existence of periodic orbits, transitivity, recurrence, entropy, etc.—depend only upon topology. This idea is commonly exploited in the computational topology community, often using the Conley index of isolating neighborhoods, to study dynamical invariants [1]. However, computing topology from time series can be a real challenge. First, one typically has only *scalar* data, not the full trajectory, and hence one must begin by reconstructing the full dynamics from that data—e.g., via delay-coordinate reconstruction. Success of this reconstruction procedure depends on several

*Corresponding author

Email address: joshua.garland@colorado.edu (Joshua Garland)

free parameters. In practice, the embedding theorems provide little guidance regarding how to choose these parameters. A number of creative strategies have been developed for doing so, but these methods require good data and input from a human expert. Moreover, the delay-coordinate reconstruction machinery (both theorems and heuristics) targets the computation of dynamical invariants like the correlation dimension and the Lyapunov exponent. If one just wants to extract the topological structure of an invariant set, as we show in this paper, a scaled-back version of that machinery may be sufficient. Nevertheless, there are issues of parameter choice here, as in the standard approach. Moreover, real-world data sets have finite length, nonzero sampling interval, limited precision, and may be contaminated by noise. In the face of these issues, one obviously cannot compute the topology to arbitrary precision, but computations can still be useful to extract information about the large-scale features.

Coarse-graining the topological analysis of data also addresses another issue: the associated computations are expensive, and that expense grows with the number of simplices in the complexes that one constructs during that process. The pioneering work in this area used cubical complexes and multivalued maps for this purpose [2], and these results can be computationally rigorous even in the face of noise. For more efficiency, one can use a simplicial complex that follows the natural geometry of the data—e.g., the *witness complex* of [3]. To construct a witness complex, one chooses a set of “landmarks,” typically a subset of the data, that become the vertices of the complex. The connections between the landmarks are determined by their nearness to the rest of the data—the “witnesses.” Two landmarks in the complex are joined by an edge, for instance, if they share at least one witness. As described in §2, there are many possible definitions for a witness “relation.” The one that we use includes a scale parameter, ε , that is intended to provide a measure of noise immunity. The ideas of persistent homology [4, 5] can be used to choose ε , build the complex, and then explore the changes in its topology with changing reconstruction dimension.

Our initial work on this approach suggests that *the witness complex correctly resolves the homology of the underlying invariant set even if the reconstruction dimension is well below the thresholds for which the embedding theorems assure smooth conjugacy between the true and reconstructed dynamics*. This paper reports upon an exploration of that conjecture in the context of the classic Lorenz system and suggests some implications and applications. To set the stage for that discussion, the rest of this section gives a brief review of delay-coordinate reconstruction. The witness complex is covered in more depth in §2, which also describes the notion of persistence and demonstrates how that idea is used to choose scale parameters for a complex built from reconstructed time-series data. In §3, we explore how the homology of such a complex changes with reconstruction dimension.

Delay-coordinate reconstruction [6, 7] is arguably the most well-established technique for obtaining the dynamics of a system from scalar time-series data. Suppose that \vec{Y} is a point on a compact invariant set $M \subset \mathbb{R}^d$, and $\vec{Y}(t)$ represents its trajectory. A smooth *measurement function* $h : M \rightarrow \mathbb{R}$ gives rise to a scalar time series, $x(t) = h(\vec{Y}(t))$, from that trajectory. Then the delay-coordinate map, $F : M \rightarrow \mathbb{R}^m$

$$F(\vec{Y}(t); h, m, \tau) = (x(t), x(t - \tau), \dots, x(t - (m - 1)\tau)) , \quad (1)$$

is almost always a diffeomorphism whenever $\tau > 0$ and m is large enough, i.e., $m > 2d_{\text{box}}$, where d_{box} is the box-counting dimension of M [8]. When these conditions are met, the reconstructed

(a) (b) (c)

Figure 1: Classic Lorenz attractor ($r = 28$, $b = 8/3$, $\sigma = 10$): (a) A 10^5 -point trajectory in \mathbb{R}^3 generated using fourth-order Runge-Kutta with a time step of $T = 0.001$. (b) A time-series trace of the x coordinate of that trajectory. (c) A 3D projection of a delay-coordinate embedding with dimension $m = 5$ and delay $\tau = 174T$, following (1).

attractor and the true attractor are diffeomorphic, and thus certainly have the same topology. The left panel of Fig. 1 shows an example: a trajectory from the classic Lorenz system [9]. The middle panel shows the corresponding time series of the x coordinate of that trajectory (i.e., $h(x, y, z) = x$), and the right panel shows a delay-coordinate reconstruction using $\tau = 174T$, where T is the interval between points in the time series. Note that a reconstruction dimension of five ($m = 5$) is required in order to satisfy the $m > 2d_{box}$ requirement for this attractor, since $d_{box} \approx 2.06$. Of course, it is not easy to display the 5D picture; Fig. 1(c) shows a 3D projection of this reconstruction.

In practice, one is presented with a scalar time series so that the dimension, d , of the original state space is unknown, and one cannot compute d_{box} without first embedding the data. Thus, choosing the reconstruction dimension m is a challenge. There are a number of heuristics for doing so. Perhaps the most well-known is the family of false near-neighbor methods pioneered in [10]. The basic idea behind this class of methods is to increase the reconstruction dimension until the geometry of the neighbor relationships stabilizes; this is taken to indicate that any false crossings created by the measurement function h have been eliminated and the dynamics are properly unfolded. The choice of the delay τ also plays a role in this unfolding. Though the theorems only require $\tau > 0$, in practice one needs to ensure that τ is large enough to make the coordinates numerically independent, but not so large that the coordinates become causally unrelated [11]. The standard approach for this—which we used to select the τ value in Fig. 1(c)—is to calculate the time-delayed average mutual information of the time series and choose τ at the first minimum of that curve [12]. There are many other methods for estimating both m and τ ; see [13] for a deeper discussion. All of these methods are subtle and subjective. Invoking them and interpreting their results requires good data and expert knowledge; the false-near neighbor method, for instance, typically overestimates the embedding dimension when noise is present in the time series—something that is unavoidable in experimental data.

In this paper, we adopt the philosophy that one might only desire knowledge of the topology of the invariant set, and we conjecture that this might be possible with a lower reconstruction dimension than that needed to obtain a true “embedding.” That is, the reconstructed dynamics might be *homeomorphic* to the original dynamics at a lower dimension than that needed for a diffeomorphically correct embedding. We will return to this idea below.

2. Witness Complexes for Dynamical Systems

To compute the topology of data that sample an invariant set of a dynamical system, we need a complex that captures the shape of the data but is robust with respect to noise and other sampling issues. To do so *efficiently*, the complex should have as few simplices as possible while still accurately representing the topology, i.e., it should be parsimonious. A witness complex is an

Figure 2: Illustration of the fuzzy witness relation (2). The closest landmark to witness w_a is l_1 , and since $\|w_a - l_2\| < D_{a1} + \varepsilon$, the simplex $\langle l_1, l_2 \rangle$ is in the complex. Similarly w_b witnesses the edge $\langle l_2, l_3 \rangle$.

ideal choice for these purposes. Such a complex is determined by the reconstructed time-series data, $W \subset \mathbb{R}^m$ —the *witnesses*—and an associated set $L \subset \mathbb{R}^m$, the *landmarks*, which can (but need not) be chosen from among the witnesses. The landmarks form the vertex set of the complex; the connections between them are dictated by the geometric relationships between W and L . In a general sense, a witness complex can be defined through a relation $R(W, L) \subset W \times L$. As Dowker noted [14], any relation gives rise to a pair of simplicial complexes. We will use one: a point $w \in W$ is a witness to an abstract k -dimensional simplex $\sigma = \langle l_{i_1}, l_{i_2}, \dots, l_{i_{k+1}} \rangle \subset L$ whenever $\{w\} \times \sigma \subset R(W, L)$. The collection of simplices that have witnesses is a complex relative to the relation R . For example, two landmarks are connected if they have a common witness—this is a one-simplex. Similarly, if three landmarks have a common witness, they form a two-simplex, and so on.

There are many possible definitions for a witness relation R . One very natural construction is to use the matrix $D(W, L)$ of distances $D_{ij} = \|w_i - l_j\|$ to define R . Sorting each row of this matrix from smallest to largest determines the set of landmarks that are closest to the i^{th} witness. One relation corresponds to assigning a cut-off, which thereby determines the simplices witnessed by w_i . For example, one can choose a fixed number (viz., k -nearest neighbors), a strict size (neighbors within some distance), or an increment. The first concept gives the “weak witness complex” of de Silva and Carlsson [3], but suffers from the problem that there is no limit on the distance to the nearest neighbors and thus a simplex might be too spread out. The second notion seems too restrictive: a portion of the invariant set M that has a low density may not be covered enough to be represented in the complex. The third idea is a compromise and gives the notion of an ε -weak witness [15], or what we call a “fuzzy” witness [16]: a point witnesses a simplex if all the landmarks in that simplex are within ε of the closest landmark to the witness:

Definition (Fuzzy Witness). *The fuzzy witness set for a point $l \in L$ is the set of witnesses*

$$\mathcal{W}_\varepsilon(l) = \{w \in W : \|w - l\| \leq \min_{l' \in L} \|w - l'\| + \varepsilon\}. \quad (2)$$

In this case, the relation consists of the collections $R = \cup_{l \in L} (\mathcal{W}_\varepsilon(l) \times \{l\})$ and a simplex σ is in the complex whenever $\cap_{l \in \sigma} \mathcal{W}_\varepsilon(l) \neq \emptyset$ —that is, when all of its vertices share a witness. The relation (2) is illustrated in Fig. 2.

The fuzzy witness complex reduces to “strong witness complex” of de Silva and Carlsson [3] when $\varepsilon = 0$. In such a complex, an edge exists between two landmarks *iff* there exists a witness that is exactly equidistant from those landmarks. This is not a practical notion of shared closeness for finite data sets. A simpler implementation of the fuzzy witness complex gives a “clique” or “flag” complex, analogous to the Rips complex [17], that consists of simplices whose pairs of vertices have a common witness. This is called a “lazy” complex in [3] and instantiated as the LazyWitnessStream class in the javaPlex [18] software. In our notation, the complex is

$$\mathcal{K}_\varepsilon(W, L) = \{\sigma \subset L : \mathcal{W}_\varepsilon(l) \cap \mathcal{W}_\varepsilon(l') \neq \emptyset, \quad \forall l, l' \in \sigma\}. \quad (3)$$

Figure 3: Varying the fuzziness parameter ε : One-skeletons of clique complexes $\mathcal{K}_\varepsilon(W, L)$ constructed from the trajectory of Fig. 1(a) using 201 landmarks (red dots) and four values of ε .

Following [16], we will use this particular construction because it minimizes computational complexity.

Figure 3 shows four witness complexes built in this fashion from the 100,000-point trajectory of the Lorenz system that is shown in Fig. 1(a). The landmarks (red dots) consist of $\ell = 201$ points equally spaced along the trajectory, i.e., every $\Delta t = 500^{\text{th}}$ point of the time series. There are many ways to choose landmarks; this particular strategy distributes them according to the invariant measure of the attractor. One could also choose landmarks randomly from the trajectory or using the “max-min” selector of [3] (choose the first landmark at random, and given a set of landmarks, choose the next to be the data point farthest away from the current set); each of these gives results similar to those shown. When ε is small, very few witnesses fall in the thin regions required by (2), so the resulting complex does not have many edges and is thus not a good representation of the shape of the data. As ε grows, more witnesses fall in the “shared” regions and the complex fills in, revealing the basic homology of the attractor of which the trajectory is a sample. There is an obvious limit to this, however: when ε is very large, even the largest holes in the complex are obscured.

Studying the change in homology under changing scale parameters is a well-established notion in computational topology. The underlying idea of *persistence* [4, 5, 19] is that any topological property of physical interest should be (relatively) independent of parameter choices in the associated algorithms. One useful way to represent information about the changing topology of a complex is the *barcode persistence diagram* [17]. Fig. 4 shows barcodes of the first two Betti numbers for the witness complexes of Fig. 3. Each horizontal line in the barcode is the interval in ε for which there exists a particular non-bounding cycle, thus the number of such lines is the rank of the homology group—a Betti number. We computed these values for β_0 and β_1 using `javaPlex` [18] over the range $0.017 \leq \varepsilon \leq 1.7$, using the `ExplicitMetricSpace` to choose the equally spaced points and the `LazyWitnessStream` to obtain a clique complex from the $\ell = 201$ landmarks. There are no three-dimensional voids in the results, i.e., β_2 was always zero for this range of ε —a reasonable implication for the 2.06-dimensional attractor. When ε is very small, as in Fig. 3(a), the witness complex has many components and the β_0 barcode shows a large number of entries. As ε grows, the spurious gaps between these components disappear, leaving a single component that persists above $\varepsilon \approx 0.014$. That is, witness complexes constructed with $\varepsilon > 0.014$ correctly capture the connectedness of the underlying attractor. The β_1 barcode plot shows a similar pattern: there are many holes for small ε that are successively filled in as that parameter grows, leaving the two main holes (i.e., $\beta_1 = 2$) for $\varepsilon > 1.01$. Above $\varepsilon > 3.2$ (not shown in Fig. 4), one of those holes disappears; eventually, for $\varepsilon > 4.05$, the complex becomes topologically trivial. Above this value, the resulting complexes—recall Fig. 3(d)—have no non-contractible loops and are homologous to a point (acyclic).

This notion of persistence can be turned around and used to select good values for the parameters that play a role in topological data analysis—i.e., looking for the ε value at which the homology stabilizes. However, definitions of what constitutes stabilization are subjective and can

Figure 4: Persistence barcodes computed using `javaPlex` for a $\ell = 201$ witness complex of the trajectory of Fig. 1(a). Each plot tabulates the two lowest Betti numbers of the complex for 100 values of the scale parameter ε . The left panel shows the behavior when $0.001 \leq \varepsilon \leq 0.1$, and the right $0.017 \leq \varepsilon \leq 1.7$.

Figure 5: Varying the number of landmarks ℓ : One-skeletons of witness complexes (3) constructed from the trajectory of Fig. 1(a) with $\varepsilon = 1.2$ and six values of ℓ .

be problematic. This kind of issue turns up routinely in nonlinear time-series analysis [13]. Even so, persistence is a powerful technique and we make use of it in a number of ways here.

Another critical step in our approach is the selection of the landmarks that constitute the vertex set of the witness complex. For efficiency, the number of landmarks should be much smaller than the number of points in the time series, but for efficacy they should be distributed so as to capture the shape of the data. A simple method for this, advocated by [3], is to use a max-min algorithm that chooses $L \subset W$ by selecting the farthest point in W from the previous selection and iterating until desired density and sparseness requirements are satisfied, if possible. For data from a dynamical system, one can alternatively exploit the natural temporal ordering and select points that are equally spaced in time (Δt) along the trajectory. If the attractor is ergodic, this will distribute the landmarks evenly relative its invariant measure. **This strategy is advantageous if one wishes to construct a map to capture the flow of the dynamics on the witness complex, as described in §4. Needless to say, the choice of the time interval between landmarks invokes an tradeoff between accuracy and computational efficiency. The nature of this form of sampling raises other issues: since the invariant measure induced by the dynamics may not agree with the underlying geometry of the invariant set, the statistical results of [20] may not be applicable. However, the algorithmic simplicity that follows from an even sampling strategy for landmark choice, together with the potential advantages for building a map on the witness complex that captures the dynamics, caused us to adopt this approach here.**

In practice, we use persistence to select the number of landmarks: i.e., given a trajectory, we build the complex for different ℓ values, calculate the homology, and choose a value at which the results stabilize. In our experience, the homology of the witness complex appears to be highly robust to the landmark spacing, but that issue will require more exploration, as described in §4. Fig. 5 shows a series of examples: fuzzy witness complexes constructed from the trajectory of Fig. 1(a) increasing ℓ from 26 to 5001. As one would expect, the complexes fill in as the number of landmarks increases. Visually, the two main holes become apparent in the $\ell = 101$ complex, which begs the question: does one really need the extra structure of an $\ell \gg 100$ complex if the goal is to resolve the large-scale topology of the attractor?

The number of landmarks required for this will obviously depend, in a complex way, on the structure of the underlying invariant set. In the simplest case—if this set were a Riemannian submanifold—the results of [20] imply that, if the “feature size” of the manifold is not too small, a well-defined number of sample points are needed to reconstruct the topology of the manifold—with high probability—from a Čech complex with a given ball size. It is also possible to guarantee that the witness complex has the same topology as an alpha complex, under appropriate conditions on the density of the sets L and W and with appropriate selection of α and ε [16, 21].

For the case at hand, however, the invariant set is a fractal, and it is impossible to obtain all of its structure from a finite sample, even though we expect to resolve more of the complexity with more landmarks. The witness complex in Fig. 5(f), for instance, resolves some of the spiral-shaped gaps in the wings of the attractor. For a fixed data set, however, there is a fundamental limit: it does not make sense to think about resolving the fine-grained structure of an object beyond the limits that are inherent in a finite-length, finite-precision sampling of that object. Since all data have these limitations, and all real-world data are noisy, it makes sense to content ourselves with an approximate topology. In this case, the roughest topology corresponds to the Williams branched manifold model of the Lorenz attractor [22], which has the homology of a figure eight, i.e. $\beta_0 = 1$ and $\beta_1 = 2$. This topology corresponds to that implied by longest bars in the persistence diagram of Fig. 4(b).

The witness complex elegantly balances effectiveness and efficiency in topological data analysis. Its use of a small number of landmarks sidesteps the issues that arise when one builds a fine-grained complex that touches every data point, which is both computationally expensive and sensitive to noise. One could also use cubical complexes to address those issues, as in [2], but the witness complex is computationally less demanding—both because it naturally follows the data and because of the “tight” way that a simplicial complex covers a space. Among other things, the dimension of each simplex in the clique complex can be restricted to be just 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 computational homology algorithms like those instantiated in javaPlex not only have fewer cells to process in the simplicial case, but also far fewer neighbors to check—e.g., during computations of isolating neighborhoods¹. Witness complexes have begun to see some use in topological data analysis [15], but they are not completely immune to the foibles of real-world data. Complexes constructed using the standard witness relations, for example, can contain “false positive” edges due to distant witnesses (in the case of the weak relation) and “false negative” edges because of the strong relation’s very stringent requirement on shared witnesses, which is inappropriate in noisy, incompletely sampled data. The fuzzy witness relation used here is intended to mitigate these issues. Of course, this approach is not without disadvantages. Both the landmark spacing Δt that makes witness complexes computationally efficient and the ε that makes the *fuzzy* witness complex robust with respect to noise are parameters that one needs to choose, and choose well. Moreover, these choices interact, as described in the following section.

The examples presented so far involve a full trajectory from a dynamical system. In the real world, however, one is generally working with reconstructions of scalar time-series data—structures whose topology is guaranteed to be identical to that of the underlying dynamics if the reconstruction process is carried out properly. But what if the dimension m does *not* satisfy the requirements of the theorems? Can one obtain useful results about the topology of that underlying system using the ideas, even if those dynamics are not properly unfolded in the sense of [6, 7, 8]? It is to this issue that we turn next.

¹See Appendix C of [16] for a detailed analysis of the associated computational costs.

3. Topologies of reconstructions

A scalar time-series of a dynamical system is a projection of the d -dimensional dynamics onto \mathbb{R}^1 —an action that does not automatically preserve the topology of the object. The method of delays allows one to reconstruct the underlying dynamics, up to diffeomorphism, if the reconstruction dimension is large enough. There are a number of conditions for the successful execution of this procedure, as mentioned in §1. According to [7], seven dimensions ($m > 2d$) are almost always sufficient to reconstruct the structure of the Lorenz attractor from the time series of Fig. 1(b); the looser bounds of [8], however, suggest that $m = 5$ is sufficient, since the box-counting dimension of that attractor is 2.06. Since the state-space dimension is generally unknown and one needs an embedding to compute d_{box} , choosing m is a primary challenge in nonlinear time-series analysis. The question we wish to address in this section is: can one use the witness complex to obtain a useful, coarse-grained description of the topology from lower-dimensional reconstructions—say, the basic connectivity of the invariant set, or the number of holes in it that are larger than a certain scale?

The short answer to that question is yes. Figure 6 shows a side-by-side comparison of witness complexes and barcode diagrams for the Lorenz trajectory of Fig. 1(a) and a two-dimensional reconstruction ($m = 2$) using the x coordinate of that trajectory. For the full 3D trajectory on the left, javaPlex needed 6942 simplices to resolve the two main holes in the attractor, with $\varepsilon = 1.2$. For the 2D reconstruction on the right—constructed with $\tau = 174T$, the first minimum of the time-delayed mutual information—the witness complex with $\varepsilon = 0.2$ has only 1916 simplices but has the same homology as the 3D complex². *In other words, the correct large-scale homology is accessible from a witness complex of a 2D reconstruction, in a computationally efficient manner, even though (a) that complex does not involve all of the data points and (b) the reconstruction does not satisfy the conditions of the associated theorems.*

Figure 6: One-skeletons of the witness complexes (top row) and barcode diagrams for β_1 (bottom row) of the Lorenz system. The plots in the left-hand column were computed from the three-dimensional (x, y, z) trajectory of Fig. 1(a); those in the right-hand column were computed from a two-dimensional ($m = 2$) delay-coordinate reconstruction from the x coordinate of that trajectory with $\tau = 174T$. In both cases, $\ell = 201$ equally spaced landmarks (red \times s) were used. Both complexes have two persistent nonbounding cycles (green and blue edges) but the 2D reconstruction requires only ≈ 1900 simplices to resolve those cycles (at $\varepsilon = 0.2$), while the full 3D trajectory requires ≈ 7000 simplices (at $\varepsilon = 1.2$) to eliminate spurious loops.

And that leads to the central question of this paper: how does the homology of the witness complex change with the reconstruction dimension m ? Intuitively, one would think that it would change at first, then stabilize when m became large enough to correctly unfold the topology of the underlying attractor. In practice, however, if m is too large, the so-called “curse of dimensionality”—when a finite amount of data is spread over a large volume—will destroy the

²The difference between the ε values that yields a persistent result in these two cases makes sense because the diameters of the true and reconstructed attractors are different. This issue is discussed further below.

Figure 7: The effect of reconstruction dimension: One-skeletons of witness complexes of different reconstructions of the scalar time series of Fig. 1(b). Both reconstructions use $\tau = 174T$, the first minimum of the average time-delayed mutual information[12], $\ell = 198$ equally spaced landmarks (red dots), and $\xi = 0.54\%$, as defined in (4).

fidelity of the complex. Moreover, increasing m exacerbates both noise effects and computational expense. For all of these reasons, it would be a major advantage if one could obtain useful information about the homology of the underlying attractor from a low-dimensional delay-coordinate reconstruction of scalar time-series data.

Again, it appears that this is possible. Figure 7 shows witness complexes for $m = 2$ and $m = 3$ reconstructions of the Lorenz time series of Fig. 1(b). The barcodes for first two Betti numbers for these two complexes, as computed using `javaPlex`, have similar structure: the complexes become connected ($\beta_0 = 1$) at a small value of ε , and the dominant, persistent homology corresponds to the two primary holes ($\beta_1 = 2$) in the attractor. Note, by the way, that Fig. 7(a) is not simply a $2D$ projection of Fig. 7(b); the edges in each complex reflect the geometry of the witness relationships in different spaces, and so may differ. Higher-dimensional reconstructions—not easily displayed—have the same homology for suitable choices of ε , though for $m > 5$, it is necessary to increase the number of landmarks to obtain a persistent $\beta_1 = 2$.

That brings up an important point: if one wants to sensibly compare witness complexes constructed from different reconstructions of a single data set, one has to think carefully about the ℓ and ε parameters. Here, we used persistence to choose a good value of ℓ . We found that the results were robust with respect to changes in that value, across all reconstruction dimension values in this study, so we fix $\ell \approx 200$ for all the experiments reported in this section³. Since the number of data points required to properly sample an object should generally grow with dimension [23], this will require more exploration, as mentioned in §4.

In the experiments in the previous section, the scale parameter ε was given in absolute units. To generalize this approach across different examples, it makes sense to compare reconstructions with ε chosen to be a fixed fraction,

$$\varepsilon = \xi \operatorname{diam}(W) \tag{4}$$

of the diameter, $\operatorname{diam}(W)$, of the set W . For example, for the full $3D$ attractor in Fig. 1(a),

$$\operatorname{diam}(W_{xyz}) = \sqrt{(x_{\max} - x_{\min})^2 + (y_{\max} - y_{\min})^2 + (z_{\max} - z_{\min})^2} = 75.3,$$

so the ε values used in Fig. 4— $0.017 \leq \varepsilon \leq 1.7$ in absolute units—translate to $2.3 \times 10^{-4} \leq \xi \leq 0.023$ in this diameter-scaled measure.

The diameter of the reconstruction varies in a natural way with the dimension m . Since delay-coordinate reconstruction of scalar data unfolds the full range of those data along every added dimension, the diameter of an m -dimensional reconstruction will be

$$\operatorname{diam}(W_m) = \sqrt{m(x_{\max} - x_{\min})^2} = 37.0 \sqrt{m},$$

³The precise value varies slightly because the length of a trajectory reconstructed from a fixed-length data set decreases with increasing m (since one needs a full span of $m \times (\tau/T)$ data points to construct a point in the reconstruction space).

where x represents the scalar time-series data. Since this unfolding will change the geometry of the reconstruction, we need to scale ε accordingly. The witness complexes in Fig. 7 were constructed with a fixed value of $\xi = 0.54\%$. Thus, for Fig. 7(a), $\varepsilon = 37.0 \sqrt{2}(0.0054) = 0.283$ in absolute units, while for Fig. 1(b), $\text{diam}(W_3) = 37.0 \sqrt{3}$ and $\varepsilon = 0.346$. This scaling of ε —which is used throughout the rest of this section—should allow the witness complex to adapt appropriately to the effects of changing reconstruction dimension and finite data.

To formalize the exploration of the reconstruction homology and extend that study across multiple dimensions, one can use a variant of the classic barcode diagram that shows, for each simplex, the reconstruction dimension values at which it appears in and vanishes from the complex. Fig. 8(a) shows such a plot for edges that involve l_0 , the first landmark on the reconstructed trajectory. A number of interesting features are apparent in this image. Unsurprisingly, most of the one-simplices that exist in the $m = 1$ witness complex—many of which are likely due to the strong effects of the projection of the underlying \mathbb{R}^d trajectory onto \mathbb{R}^1 —vanish when one moves to $m = 2$. There are other short-lived edges in the complex as well: e.g., the edge from l_0 to l_{120} that is born at $m = 2$ and dies at $m = 3$. The sketch in Fig. 8(b) demonstrates how edges can be born as the dimension increases: in the $m = 2$ reconstruction, ℓ_1 and ℓ_3 share a witness (the green square); when one moves to $m = 3$, spreading all of the points out along the added dimension, that witness is moved far from ℓ_3 —and into the shared region between ℓ_1 and ℓ_2 . There are also long-lived edges in the complex of Fig. 8(a). The one between l_0 and l_{140} that persists from $m = 1$ to $m = 8$ is particularly interesting: this pair of landmarks has shared witnesses in the scalar data *and in all reconstructions*. Possible causes for this are explored in more depth below. All of these effects depend on ξ , of course; lowering ξ will decrease both the number and average length of the edge persistence bars.

Figure 8: (a) Dimension barcode for edges in the witness complex of the reconstructed scalar time series of Fig. 1(b) that involve l_0 , the first landmark, for reconstructions with $m = 1, \dots, 8$. The vertical axis is labeled with the indices of the remaining 197 landmarks in the complex; a circle at the $m - 1 \rightarrow m$ tickmark on the horizontal axis indicates the transition at which an edge between l_0 and l_i is born; a square indicates the transition at which that edge vanishes from the complex. An arrow at the right-hand edge of the plot indicates an edge that was still stable when the algorithm completed. For all reconstructions, $\tau = 174T$, $\ell = 198$, and $\xi = 0.54\%$. (b) Sketch of the birth and death of edges at the $m = 2 \rightarrow 3$ transition.

While this Δm barcode image is interesting, the amount of detail that it contains makes it somewhat unwieldy. To study the m -persistence of all of $\ell \times \ell$ edges in a witness complex, one would need to examine ℓ of these plots—or condense them into a single plot with ℓ^2 entries on the vertical axis. Instead, one can plot what we call an *edge lifespan diagram*: an $\ell \times \ell$ matrix whose $(i, j)^{\text{th}}$ pixel is colored according to the maximum range of m for which an edge exists in the complex between the i^{th} and j^{th} landmarks; see Fig. 9. If the edge $\{l_i, l_j\}$ existed in the complex for $2 \leq m < 3$ and $5 \leq m < 8$, for instance, Δm would be three and the i, j^{th} pixel would be coded in cyan. Edges that do not exist for any dimension are coded white.

A prominent feature of Fig. 9 is a large number (683) of edges with a lifespan 1 (blue). Of these edges, 463 exist for $m = 1$, but not for $m = 2$, and thus reflect the anomalous behavior of projecting a 2.06 dimensional object onto a line. This was also seen, as described above, in the barcode of Fig. 8.

Figure 9: Edge lifespan diagram: pixel i, j on this image is color-coded according to the maximum range Δm of dimension for which an edge exists between landmarks l_i and l_j in the witness complex of the reconstructed scalar time series of Fig. 1(b) for $m = 1, \dots, 8$. For all reconstructions, $\tau = 174T$, $\ell = 198$, and $\xi = 0.54\%$.

Another interesting set of features in the lifespan diagram is the diagonal line segments. Note that the *color* of the pixels in these segments varies, though most of them correspond to edges with longer lifespans. These segments indicate the existence of Δm -persistent edges $\{l_i, l_j\}, \{l_{i+1}, l_{j+1}\}, \{l_{i+2}, l_{j+2}\} \dots$. This is likely due to the continuity of the dynamics [24]. Recall that the landmarks are evenly spaced in time, so l_{i+1} is the Δt -forward image of l_i . Thus a diagonal segment may indicate that the Δt -forward images of (at least one) witness that is shared between l_i and l_j is shared between l_{i+1} and l_{j+1} , and so on. The lengths of the longer line segments suggest that that continuity fails after 5-10 Δt steps, probably because of the positive Lyapunov exponents on the attractor. As a simple check on this reasoning, one can compute an edge lifespan diagram for a dynamical system with a limit cycle. The structure of such a plot (not shown) is dominated by diagonal lines of high Δm -persistence, with a few other scattered one-persistent edges. One can capture the underlying dynamical information that gives rise to these effects more fully using what we call the *witness map*, as mentioned briefly in §4 and described at more length in [16]. Choosing landmarks according to the invariant measure, as mentioned previously, has an important advantage if one is ultimately interested in constructing such a map, since the speed, relative to the simplices in the complex, will be more even.

The rationale behind studying the *maximal* m -lifespan goes back to one of the basic premises of persistence: that features that persist for a wide range of parameter values are in some sense meaningful. To explore this, Fig. 10 shows the witness complex of Fig. 7(a), highlighting the $\Delta m \geq 2$ -persistent edges: those that exist at $m = 2$ and persist at least to $m = 4$. There exists a fundamental core to the complex that persists as the dimension grows and thus is robust to geometric distortion, but there are also short-lived edges that fill in the complex in accord with the local geometric structure of the reconstruction. Indeed, when $m = 2$, the projection artificially compresses near the origin; small simplicies fill in this region due to the landmark clustering there. However, in the transition to $m = 3$ —viz., Fig. 7(b)—this region stretches away from the origin, spreading the landmarks out. There is a similar cluster of fragile edges near the lower left corner of the complex.

This geometric evolution with increasing reconstruction dimension leads to the death of many local edges. Even so, the large-scale topology is correct in both complexes of Fig. 7, although the fine-scale topology is resolved differently by the dimension-dependent geometry. So while the edges with longer lifespan are indeed more important to the core structure, the short-lived edges are also important because they allow the complex to adapt to the geometric evolution of the attractor and fill in the details of the skeleton that are necessary and meaningful in that dimension.

In the spirit of the false near-neighbor method [10], one might be tempted to take the short-lived edges as an indication that the reconstruction dimension is inadequate. However, one computes homology *from the overall complex*. As the example above shows, homology is relatively robust with respect to individual edges. The moral of this story is that the lifespan of an edge is not necessarily an obvious indication of its importance to the homology of the complex; Δm -persistence

Figure 10: Witness complex of Fig. 7(a) with $\Delta m \geq 2$ -persistent edges shown as thick (black) lines, and the $\Delta m = 1$ edges as (red) dashed lines.

Figure 11: The effects of noise: One-skeletons and barcodes of witness complexes for $m = 2$, $\xi = 0.41\%$, $\tau = 174T$, $\ell = 201$ reconstructions of the scalar time series of Fig. 1(b) with added uniform noise of size $\nu = 1$ (a and c) and $\nu = 4$ (b and d)—respectively, 1.9% and 7.6% of the diameter of the reconstructed attractor.

plays a different role here than the abscissa of traditional barcode persistence plots.

A closely related issue is noise, which is always present in real-world data and can disturb the geometric relationships between points in the complex. To study the effect on the fuzzy witness complex, we add uniformly distributed noise on the interval $[-\nu/2, \nu/2]$ to each point of the trajectory of Fig. 1(b), and then perform a delay-coordinate reconstruction using $m = 2$. Visually, Fig. 11(a), where $\nu = 1$, shows a surprisingly similar complex to the noise-free case of Fig. 6(b); however, the structure is clearly different in Fig. 11(b), where $\nu = 4$. Comparing the barcodes, one sees that the fine-scale structure is, unsurprisingly, washed out for smaller values of ε than in the noise-free case. In particular, the persistence intervals of small-scale loops are decreased. It is encouraging, however, to see that the two major holes persist over a wide range of ε even when the noise level is close to 2% of the diameter of the attractor. The larger noise level (7.6% of that diameter) is enough to destroy the two large loops so that the complex becomes acyclic when $\varepsilon > 0.152$.

The relative immunity to noise is a general feature of persistent homology for point-cloud data [17]. The additional robustness of the witness complex—an important factor in our choice of this representation—has two sources: the fuzziness parameter and the fact that multiple points can witness a given simplex. Even if noise moves one witness out of the shared region that is defined by ε , it may not move *all* of the witnesses to a particular simplex out of that shared region. This suggests another potential noise mitigation technique: build a complex that only contains simplices that have at least n witnesses—and perhaps explore n -persistence of the associated homology. We plan to explore these ideas in future work.

Assessing the change in topology with changing reconstruction dimension is a new flavor of persistence—an idea that has traditionally been applied in the context of scale parameters like our fuzziness parameter ε . Recall that persistent homology is based on the idea of a filtration, i.e., a nested set of complexes. Most of the standard simplicial complex constructions for data have a natural parameter that gives rise to a filtration: for example, the radius for a sequence of Čech complexes. This filtration can be used to define persistent homology [19] and identify robust topological features—e.g., those that correspond to long-lived bars in a barcode diagram. Since edges of the witness complex are both created and destroyed with increasing reconstruction dimension, the idea of persistence would not seem to work for this case (except perhaps in the sense of zigzag persistence [25]). However, if we use Δm -persistent edges—edges that exist for some *range* of m —we can get a filtration. Specifically, if $L^s(\Delta m)$ is the set of edges that exist for a range Δm , then $L^s(\Delta m) \subset L^s(\Delta m - 1)$. The same would be true for a (clique) complex built from these edges. Note that this gives, somewhat surprisingly, an inclusion as Δm decreases, the reverse

of what one might think at first.

4. Conclusion & Future Work

We have shown that it is possible to compute the topology of an invariant set of a dynamical system using a coarse-grained simplicial complex—a *witness complex*—built from a low-dimensional reconstruction of a scalar time series. These results have a number of interesting implications. Among other things, they suggest that the traditional delay-coordinate reconstruction process may be excessive if one is only interested in topological structure. Indeed, this explains why it is possible to construct accurate predictions of the future state of a high-dimensional dynamical system using a two-dimensional delay-coordinate reconstruction [26]. The delay-coordinate machinery strives to obtain a diffeomorphism—not a homeomorphism—between the true and reconstructed attractors. However, many of the important properties of attractors (transitivity, continuity, recurrence, entropy, etc.) are topological, so requiring only a homeomorphism is natural and more efficient [2].

Given a finite set of data, and finite computational power, one can never compute the *full* topology of a fractal attractor. Nevertheless, it is useful to obtain a coarse approximation of the topological features—e.g., the two main holes in the canonical Lorenz attractor. The approach described in this paper produces such an approximation in a computationally efficient and robust manner. The scale of the exploration can be controlled by the spacing between the landmarks in the witness complex, so finer structure can be observed if needed. There is a computational cost to doing so, of course, even with the natural parsimony of the witness complex.

To choose the various free parameters in this procedure, we use ideas from *persistent homology* to investigate the dependence of the complex on the number of landmarks ℓ and on the fuzziness parameter ε . It might be possible to make this multi-parameter exploration of persistence rigorous using ideas related to those of [27]. To study persistence across reconstruction dimension, we introduce the maximum lifespan Δm . This parameter gives a filtration so that the standard persistent homology methods apply. We found that the long-lived edges determine the core of the complex and that the shorter-lived edges fill in the fine-grained structure. Of course, some edges are superfluous for the homology, and destroying them does not change the Betti numbers. This is part of the strength of the witness complex approach, as well as a part of the source of its noise immunity.

The approach described in this paper uses the fuzzy witness relation of (2) and builds a clique complex (3) from the resulting landmarks. Since the structure of such a complex is determined by its edges, that process is computationally efficient. In the future, we plan to investigate other choices for the witness relation, e.g. [3, 28], and to see if removing the clique assumption allows for greater fidelity without excess computational burden. It would also be interesting to investigate how the robustness with respect to noise varies with the choice of the witness relation.

Fuzzy witness complexes not only support useful strategies for selecting a reconstruction dimension that reveals the (approximate) homology of the attractor. This approach is also a step along the path to detecting and characterizing bifurcations [29]. Suppose that, for example, the first part of the data set corresponds trivially to an equilibrium—that is, to a set with $\beta_k = 0$ for all $k > 0$ —but that this equilibrium undergoes a bifurcation to an oscillatory regime midway through

the data set. In this case, a shift in β_1 signals a regime change. Regime changes in a nonstationary data set could be due to dynamical parameters changing with time, or to switching from one dynamical system to another, e.g., such as an iterated function system [24].

Further analysis of the dynamics, beyond the static homology of the points in the reconstructed trajectory, requires the construction of a map on the witness complex that is induced by the temporal shift on the time series. Following the ideas of [16], the time-ordering of the data gives rise to a simplicial multi-valued map, that we call the *witness map*, and under certain conditions this map can be shown to induce a map on homology that is the same as any continuous “selector”, allowing one to compute the Conley index, which can be used to prove the existence of various invariant sets. **In this context, it is useful to choose landmarks following the invariant measure of the attractor so that the speed is nearly constant relative to simplices in the complex. If the simplices were chosen to be spatially uniform, then the trajectory would spend extra time in a large simplex where the speed was small, and skip over simplices where the speed was large. Of course if one is only interested in the static topology of the attractor, this is not important. Indeed similar barcodes to those in Fig. 6 are obtained by using either random or max-min landmarks.**

This paper uses a single example—the classic Lorenz system—but we believe that the approach will work on other dynamical systems. We plan in the future to do a careful exploration of additional systems, both maps and flows, to explore the effects of the size and spacing of the data, to study the interaction with the fuzziness parameter ε and the number ℓ of landmarks. The edge-length distribution might be a useful aid in choosing good values for the former; for the latter, it may be useful to examine the distribution of witness-landmark distances. The underlying problem is not simple; the data sample the invariant set, the landmarks sample the data, and the complex reflects the geometric relationships between the landmarks and the rest of the data. One could get at some of this by changing the data length—say, using half of the time-series data—repeating the analysis, and looking to see if the “best” parameter values change. It will be important to determine guidelines for the optimal number and distribution of the landmarks, and how this should scale with m for a fixed trajectory.

Acknowledgments

This material is based upon 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. The authors thank the IMA for hosting the Algebraic Topology and Dynamics conference in February 2014, and K. Mischaikow for numerous useful discussions.

References

- [1] T. Kaczynski, K. Mischaikow, M. Mrozek, Computational Homology, Vol. 157 of Applied Mathematical Sciences, Springer-Verlag, New York, 2004.
- [2] K. Mischaikow, M. Mrozek, J. Reiss, A. Szymczak, Construction of symbolic dynamics from experimental time series, Phys. Rev. Lett. 82 (1999) 1144–1147, <http://link.aps.org/doi/10.1103/PhysRevLett.82.1144>.

- [3] V. de Silva, E. Carlsson, Topological estimation using witness complexes, in: M. Alexa, S. Rusinkiewicz (Eds.), Eurographics Symposium on Point-Based Graphics (2004), The Eurographics Association, Zurich, 2004, pp. 157–166, <http://pages.pomona.edu/~vds04747/public/papers/witness.pdf>.
- [4] H. Edelsbrunner, D. Letscher, A. Zomorodian, Topological persistence and simplification, in: IEEE Symposium on Foundations of Computer Science, 2000, pp. 454–463.
- [5] V. Robins, Computational topology for point data: Betti numbers of α -shapes, in: K. Mecke, D. Stoyan (Eds.), Morphology of Condensed Matter, Vol. 600 of Lect. Notes in Physics, Springer Berlin Heidelberg, 2002, pp. 261–274, http://dx.doi.org/10.1007/3-540-45782-8_11.
- [6] N. Packard, J. Crutchfield, J. Farmer, R. Shaw, Geometry from a time series, Phys. Rev. Lett 45 (9) (1980) 712–716, <http://dx.doi.org/10.1103/PhysRevLett.45.712>.
- [7] F. Takens, Detecting strange attractors in fluid turbulence, in: D. Rand, L.-S. Young (Eds.), Dynamical Systems and Turbulence, Springer, Berlin, 1981, pp. 366–381.
- [8] T. Sauer, J. Yorke, M. Casdagli, Embedology, J. Stat. Phys. 65 (1991) 95–116, <http://dx.doi.org/10.1007/BF01053745>.
- [9] E. Lorenz, Deterministic nonperiodic flow, J. Atmos. Sci. 20 (1963) 130–141, [http://dx.doi.org/10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](http://dx.doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).
- [10] M. Kennel, R. Brown, H. D. I. Abarbanel, Determining minimum embedding dimension using a geometrical construction, Phys. Rev. A 45 (1992) 3403–3411, <http://dx.doi.org/10.1103/PhysRevA.45.3403>.
- [11] M. Casdagli, S. Eubank, J. Farmer, J. Gibson, State space reconstruction in the presence of noise, Physica D 51 (1-3) (1991) 52–98, [http://dx.doi.org/10.1016/0167-2789\(91\)90222-U](http://dx.doi.org/10.1016/0167-2789(91)90222-U).
- [12] A. Fraser, H. Swinney, Independent coordinates for strange attractors from mutual information, Phys. Rev. A 33 (2) (1986) 1134–1140, <http://dx.doi.org/10.1103/PhysRevA.33.1134>.
- [13] E. Bradley, H. Kantz, Nonlinear time-series analysis revisited, Chaos 25 (2015) 097610, <http://arxiv.org/abs/1503.07493>.
- [14] C. Dowker, Homology groups of relations, Ann. of Math 56 (1) (1952) 84–95, <http://dx.doi.org/10.2307/1969768>.
- [15] G. Carlsson, Topological pattern recognition for point cloud data, Acta Numerica 23 (2014) 289–368, <http://dx.doi.org/10.1017/s0962492914000051>.
- [16] Z. Alexander, E. Bradley, J. Meiss, N. Sanderson, Simplicial multivalued maps and the witness complex for dynamical analysis of time series, SIAM J. Appl. Dyn. Sys. 14 (3) (2015) 1278–1307, <http://dx.doi.org/10.1137/140971415>.
- [17] R. Ghrist, Barcodes: The persistent topology of data, Bull. Amer. Math. Soc. 45 (1) (2008) 61–75, <http://dx.doi.org/10.1090/S0273-0979-07-01191-3>.
- [18] A. Tausz, M. Vejdemo-Johansson, H. Adams, javaPlex: a research platform for persistent homology, Book of Abstracts Minisymposium on Publicly Available Geometric/Topological Software (2012) 7.
- [19] A. Zomorodian, G. Carlsson, Computing persistent homology, Disc. Comp. Geom. 33 (2005) 249–274, <http://dx.doi.org/10.1007/s00454-004-1146-y>.
- [20] P. Niyogi, S. Smale, S. Weinberger, Finding the homology of submanifolds with high confidence from random samples, Discr. Comput. Geom. 39 (1-3) (2008) 419–441, <http://dx.doi.org/10.1007/s00454-008-9053-2>.
- [21] Z. Alexander, A topology-based approach for nonlinear time-series analysis with applications in computer performance analysis, Ph.D. thesis, University of Colorado (April 2012).
- [22] R. Williams, The structure of Lorenz attractors, Publications Mathematiques de l’IHES 50 (1979) 73–99, <http://dx.doi.org/10.1007/BF02684770>.
- [23] A. Tsonis, J. Elsner, K. Georgakakos, Estimating the dimension of weather and climate attractors: Important issues about the procedure and interpretation, J. Atmos. Sci. 50 (15) (1993) 2549–2555, [http://dx.doi.org/10.1175/1520-0469\(1993\)050<2549:ETDOWA>2.0.CO;2](http://dx.doi.org/10.1175/1520-0469(1993)050<2549:ETDOWA>2.0.CO;2).
- [24] Z. Alexander, J. Meiss, E. Bradley, J. Garland, Iterated function system models in data analysis: Detection and separation, Chaos 22 (2) (2012) 023103, <http://dx.doi.org/10.1063/1.3701728>.
- [25] G. Carlsson, V. de Silva, Zigzag persistence, Found. Comput. Math 10 (4) (2010) 367–405, <http://dx.doi.org/10.1007/s10208-010-9066-0>.

- [26] J. Garland, E. Bradley, Prediction in projection, *Chaos* 25 (2015) 123108, <http://dx.doi.org/10.1063/1.4936242>.
- [27] E. Carlsson, G. Singh, A. Zomorodian, Computing multidimensional persistence, *J. Comp. Geom.* 1 (1) (2010) 72–100, <http://jocg.org/index.php/jocg/article/view/19>.
- [28] D. Attali, H. Edelsbrunner, J. Harer, Y. Mileyko, Alpha-beta witness complexes, in: F. Dehne, J. Sack, N. Zeh (Eds.), *Lecture Notes in Computer Science: Algorithms and Data Structures 10th Intl. Workshop, Wads 2007*, Vol. 4619, Springer, 2007, pp. 386–397, http://link.springer.com/chapter/10.1007/978-3-540-73951-7_34.
- [29] J. Berwald, M. Gidea, M. Vejdemo-Johansson, Automatic recognition and tagging of topologically different regimes in dynamical systems, *Discontinuity, Non-linearity, and Complexity* 3 (4) (2015) 413–426.

Figure 1 (a)




Figure 1 (b)




Figure 1 (c)




Figure 2




Figure 3 (a)




Figure 3 (b)




Figure 3 (c)




Figure 3 (d)

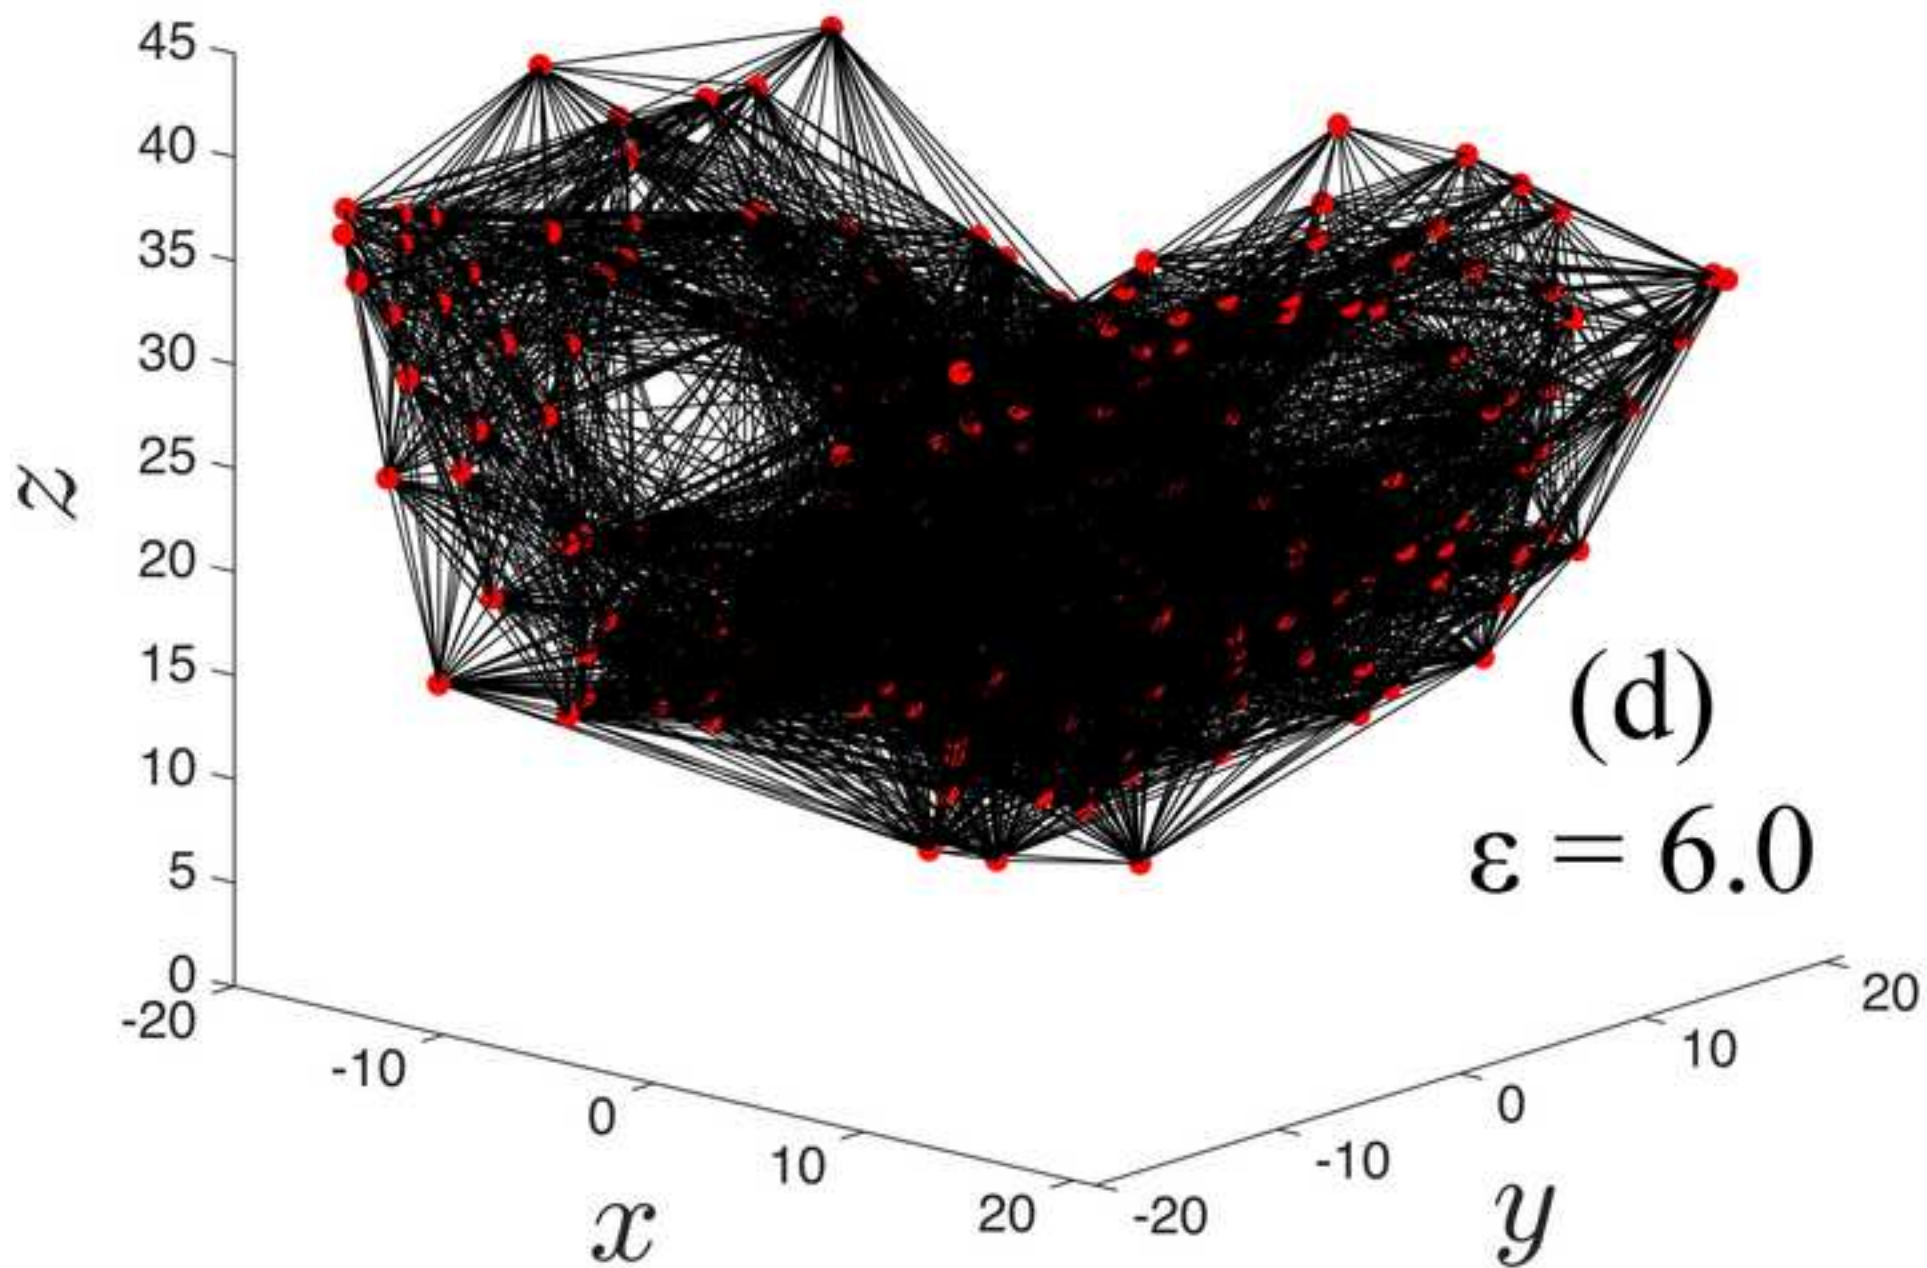


Figure 4 (a)




Figure 4 (b)




Figure 5 (a)




Figure 5 (b)




Figure 5 (c)




Figure 5 (d)




Figure 5 (e)




Figure 5 (f)

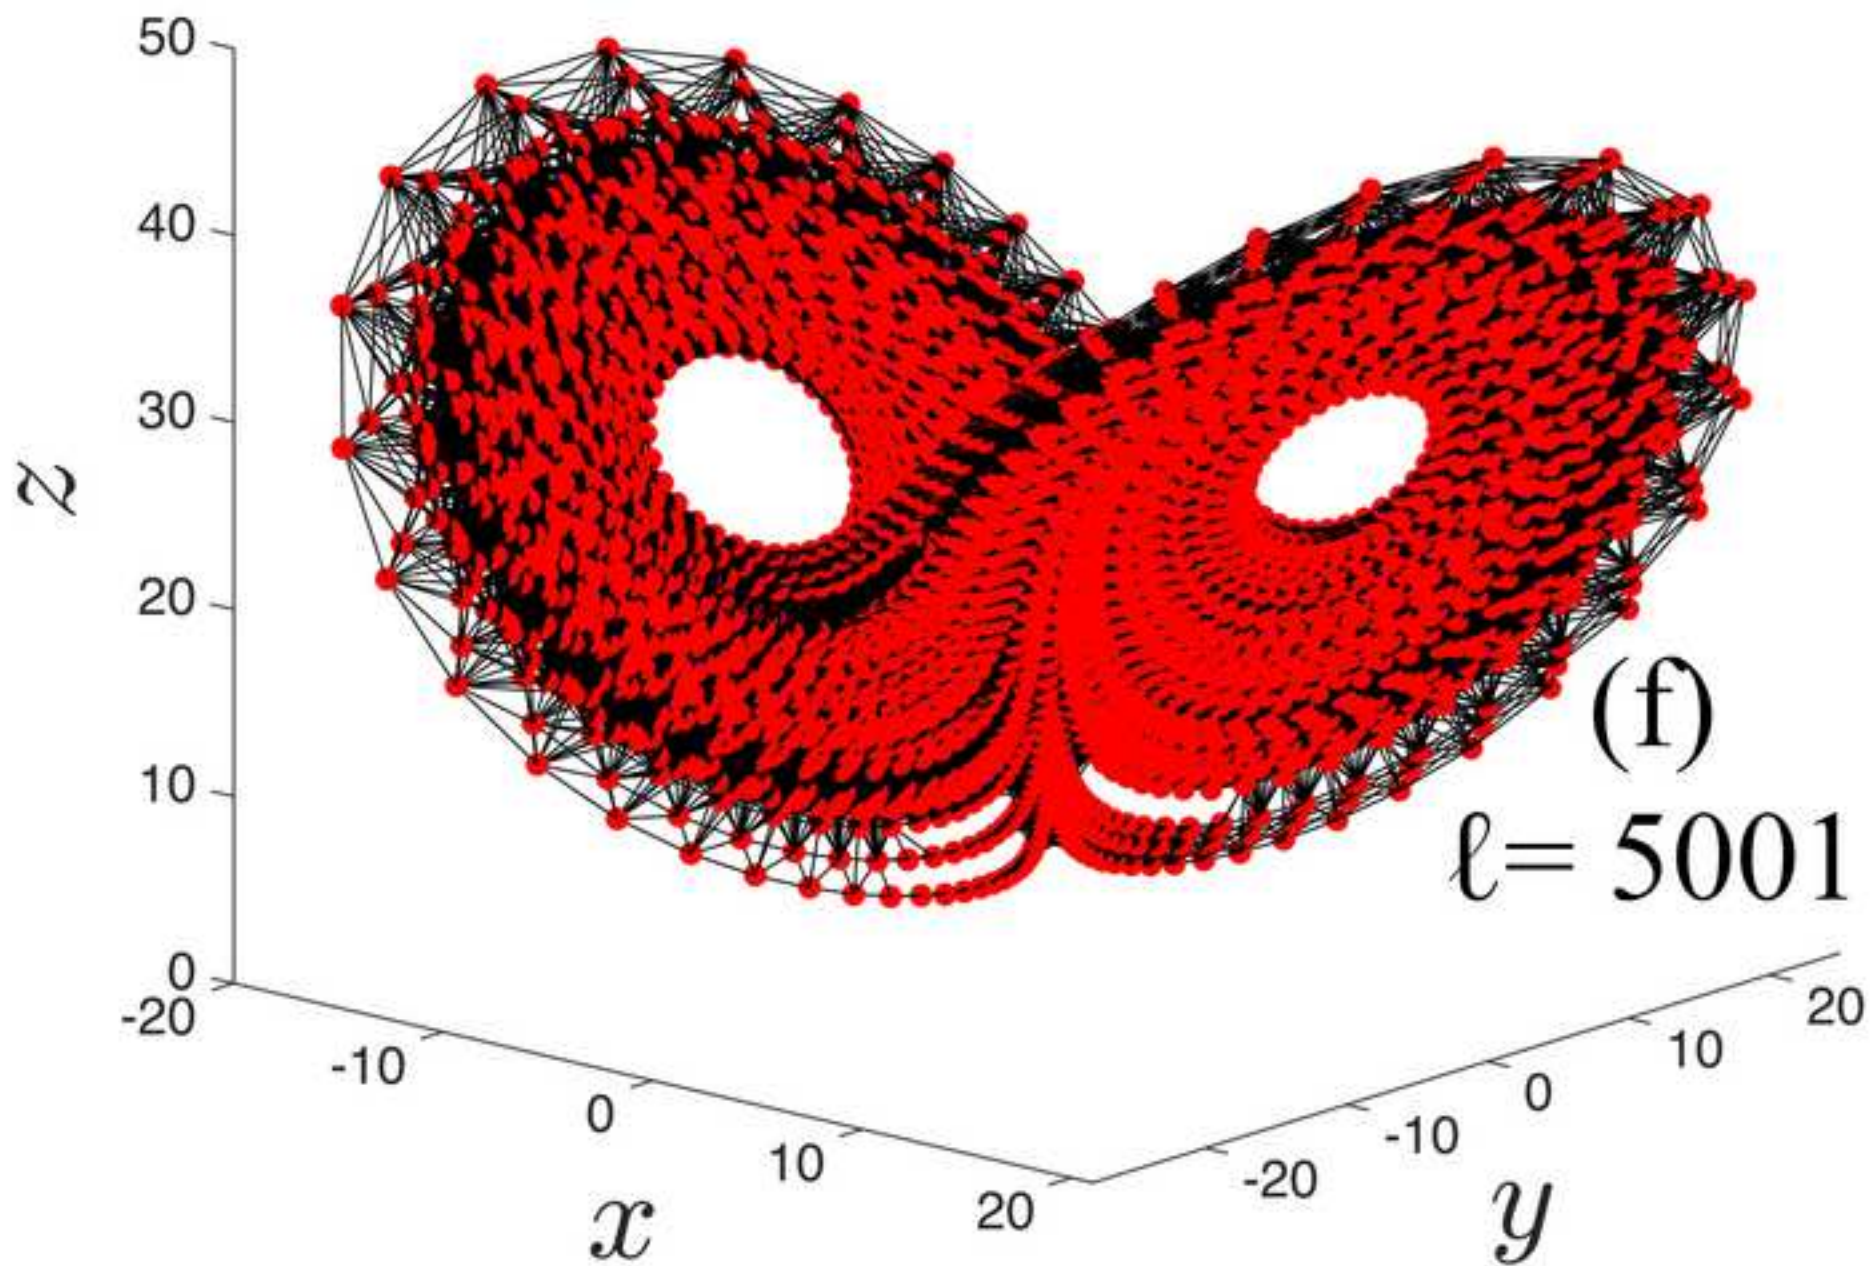


Figure 6 (a)




Figure 6 (b)

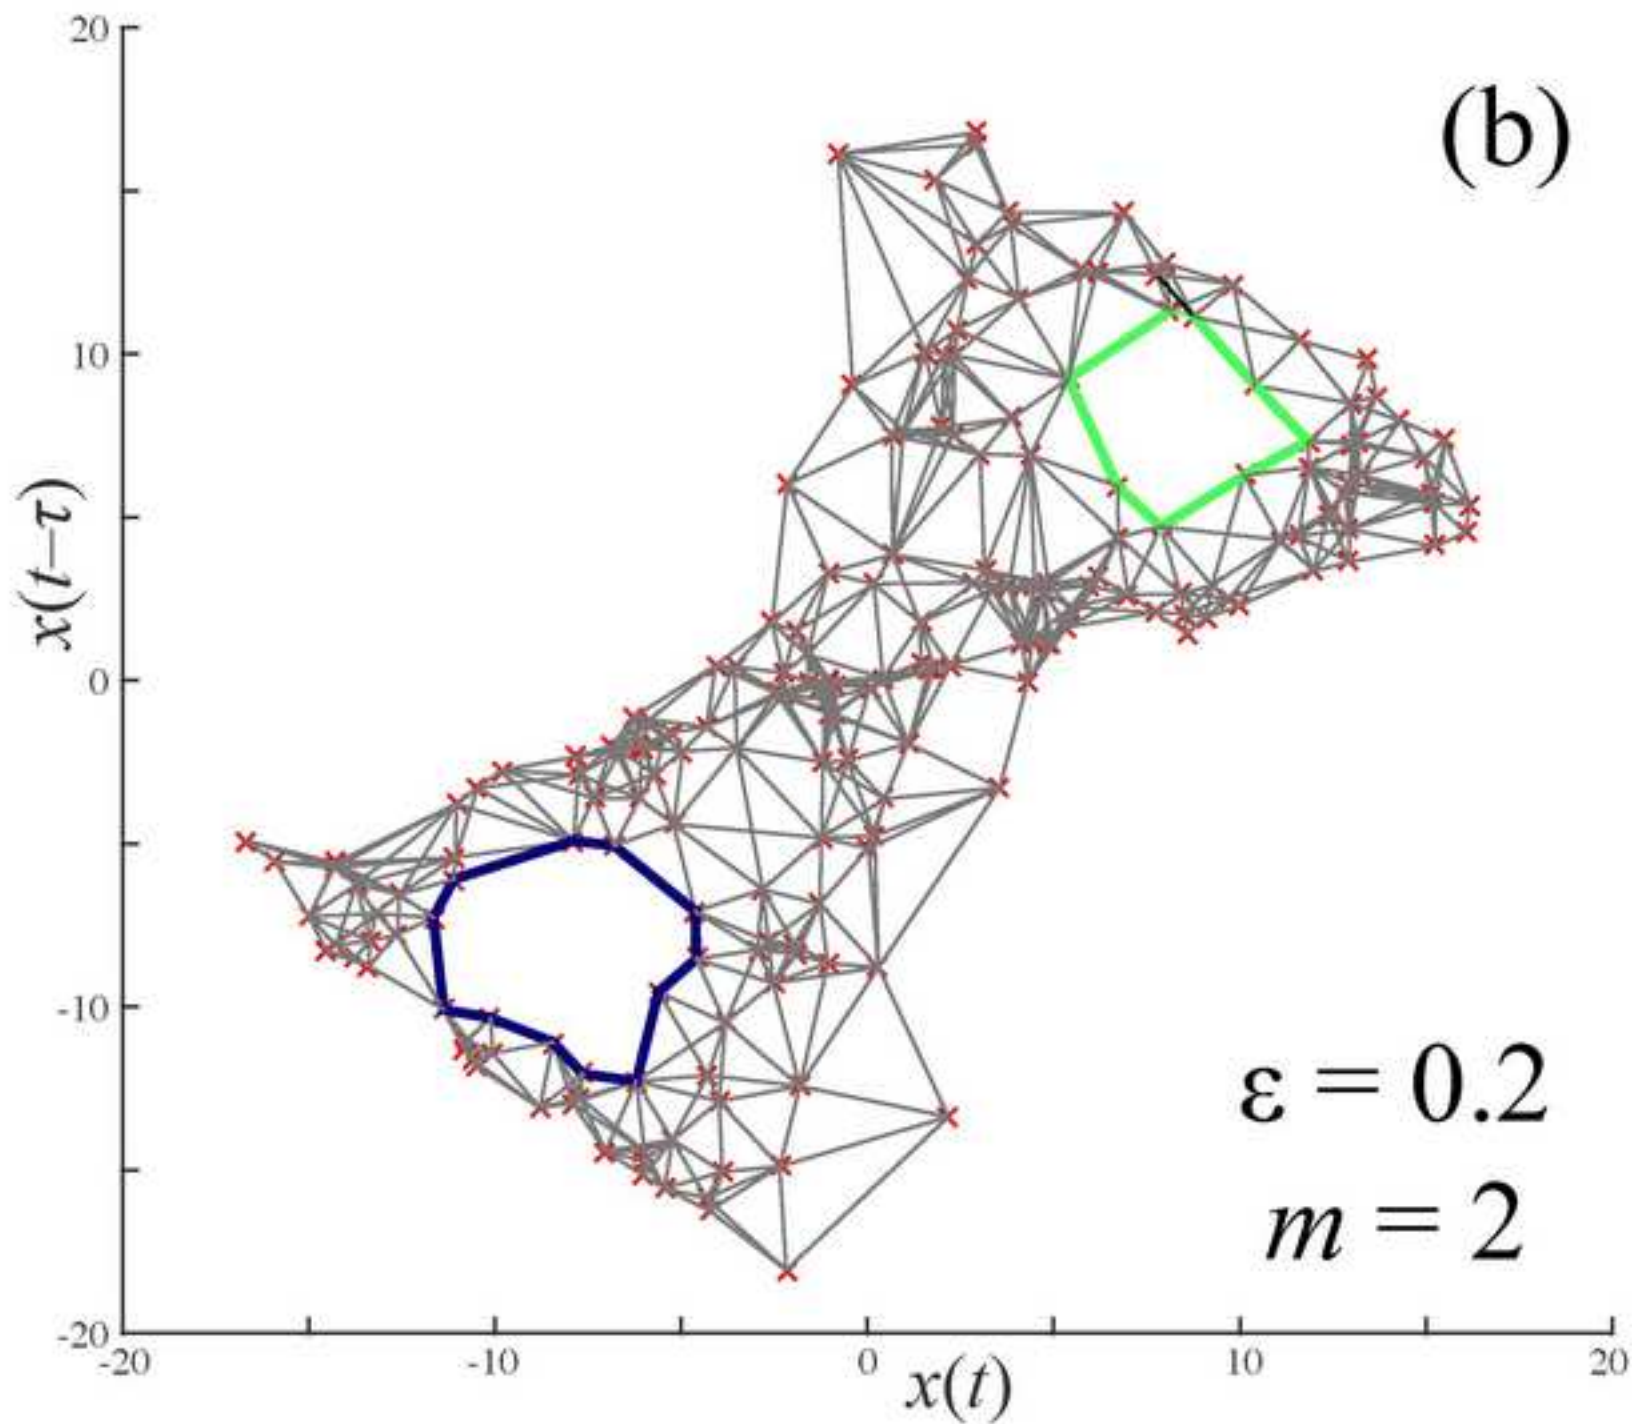


Figure 6 (c)




Figure 6 (d)




Figure 7 (a)




Figure 7 (b)

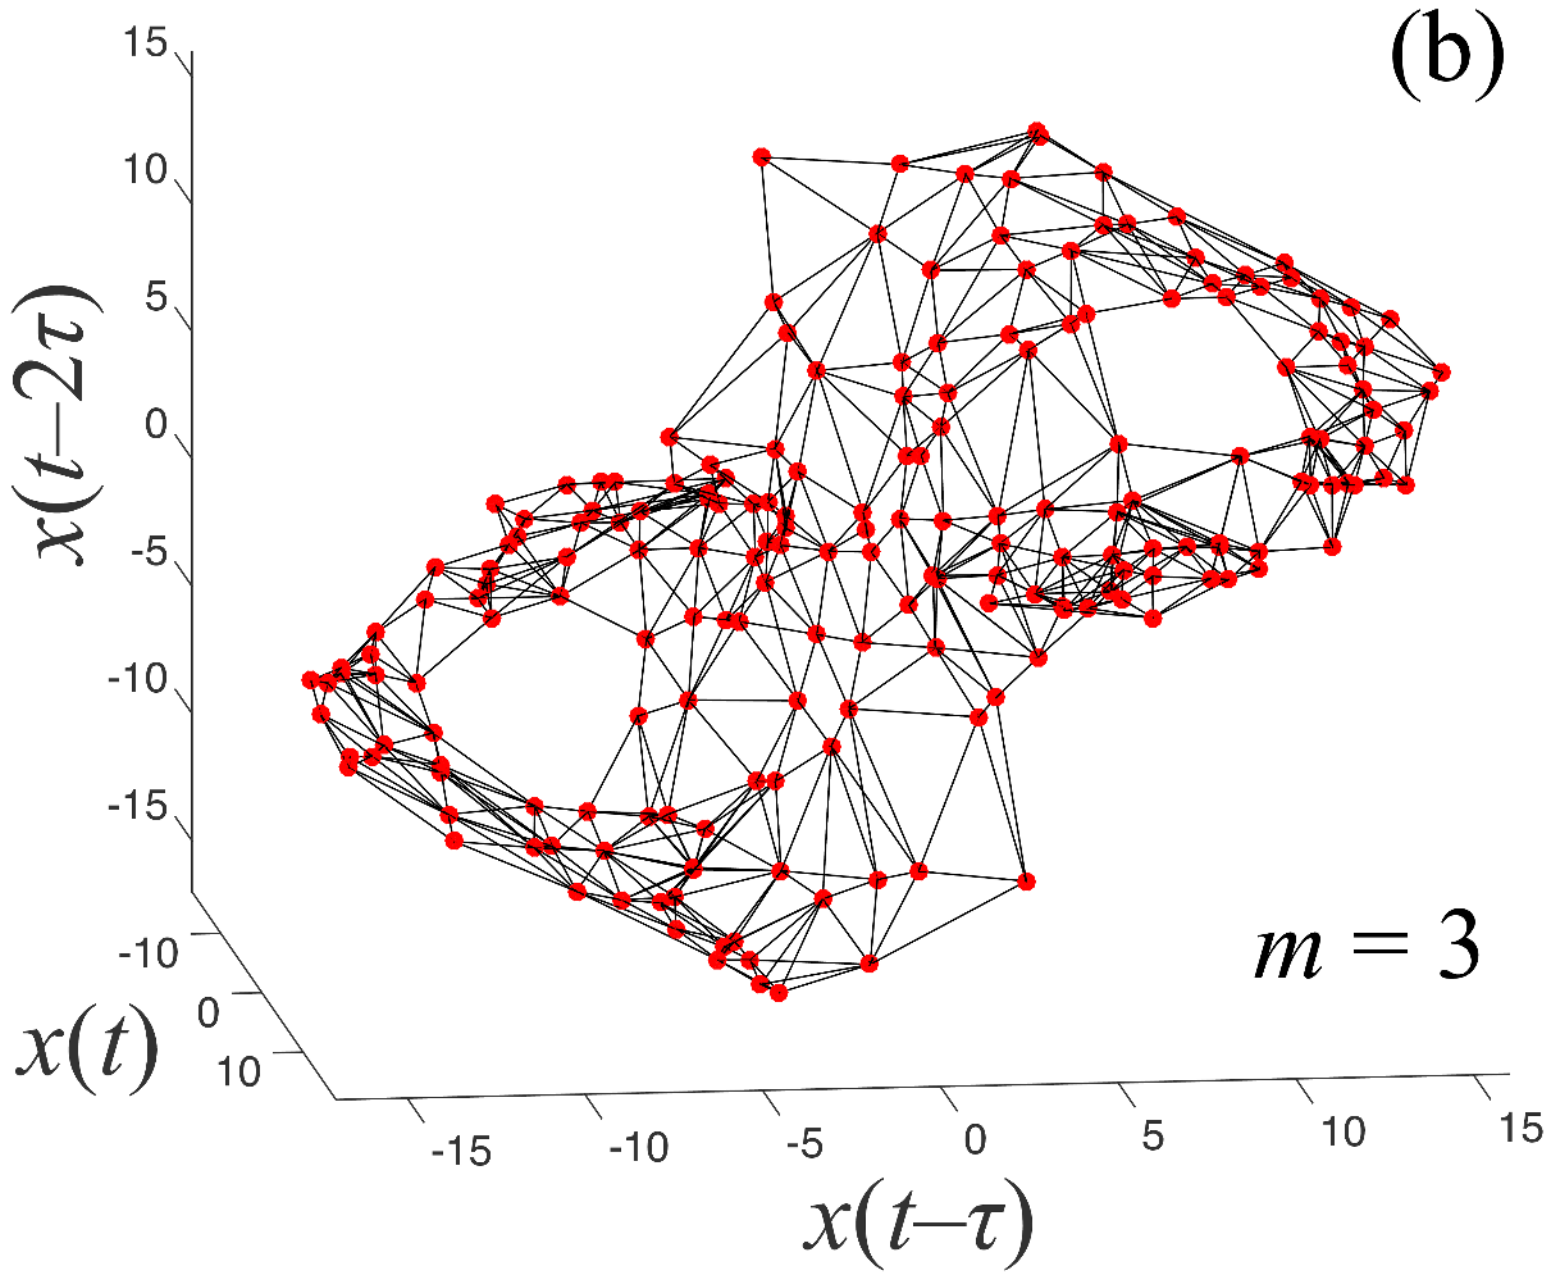


Figure 8 (a)




Figure 8 (b)




Figure 9




Figure 10




Figure 11 (a)




Figure 11 (b)




Figure 11 (c)




Figure 11 (d)

