

# FRactal Dimension Estimation with Persistent Homology: A Comparative Study

JONATHAN JAQUETTE AND BENJAMIN SCHWEINHART

ABSTRACT. We propose that the recently defined persistent homology dimensions are a practical tool for fractal dimension estimation of point samples. We implement an algorithm to estimate the persistent homology dimension, and compare its performance to classical methods to compute the correlation and box-counting dimensions in examples of self-similar fractals, chaotic attractors, and an empirical dataset. The performance of the 0-dimensional persistent homology dimension is comparable to that of the correlation dimension, and better than box-counting.

## 1. INTRODUCTION

Loosely speaking, fractal dimension measures how local properties of a set depend on the scale at which they are measured. The Hausdorff dimension was perhaps the first precisely defined notion of fractal dimension [23, 41]. It is difficult to estimate in practice, but several other more computationally practicable definitions have been proposed, including the box-counting [10] and correlation [34] dimensions. These notions are in-equivalent in general.

Popularized by Mandelbrot [57, 58], fractal dimension has applications in a variety of fields including materials science [17, 44, 92], biology [6, 52, 66], soil morphology [68], and the analysis of large data sets [7, 86]. It is also important in pure mathematics and mathematical physics, in disciplines ranging from dynamics [81] to probability [8]. In some applications, it is necessary to estimate the dimension of a set from a point sample. These include earthquake hypocenters and epicenters [39, 47], rain droplets [28, 54], galaxy locations [37], and chaotic attractors [34, 81].

We propose that the recently defined persistent homology dimensions [1, 77] (Definition 1 below) are a practical tool for dimension estimation of point samples. Persistent homology [20] quantifies the shape of a geometric object in terms of how its topology changes as it is thickened; roughly speaking,  $i$ -dimensional persistent homology ( $PH_i$ ) tracks  $i$ -dimensional holes that form and disappear in this process. Recently, it has found many applications in fields ranging from materials science [43, 73] to biology [31, 89]. There are further applications of persistent homology in machine learning, for which different methods of vectorizing persistent homology have been defined [2, 11]. In most applications, larger geometric features represented by persistent homology are of greatest interest. However, in the current context it is the smaller features — the “noise” — from which the dimension may be recovered.

For a finite subset of a metric space  $\{x_1, \dots, x_n\}$  and a positive real number  $\alpha$  define the  $\alpha$  weight  $E_\alpha^i(x_1, \dots, x_n)$  to be the sum of the lengths of the (finite)  $PH_i$  intervals to the  $\alpha$  power:

$$E_\alpha^i(x_1, \dots, x_n) = \sum_{I \in PH_i(x_1, \dots, x_n)} |I|^\alpha.$$

---

*Date:* July 2019.

Research of the first author was supported in part by NIH T32 NS007292 and NSF DMS-1440140 while the author was in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2018 semester.

Research of the second author was supported in part by a NSF Mathematical Sciences Postdoctoral Research Fellowship under award number DMS-1606259.

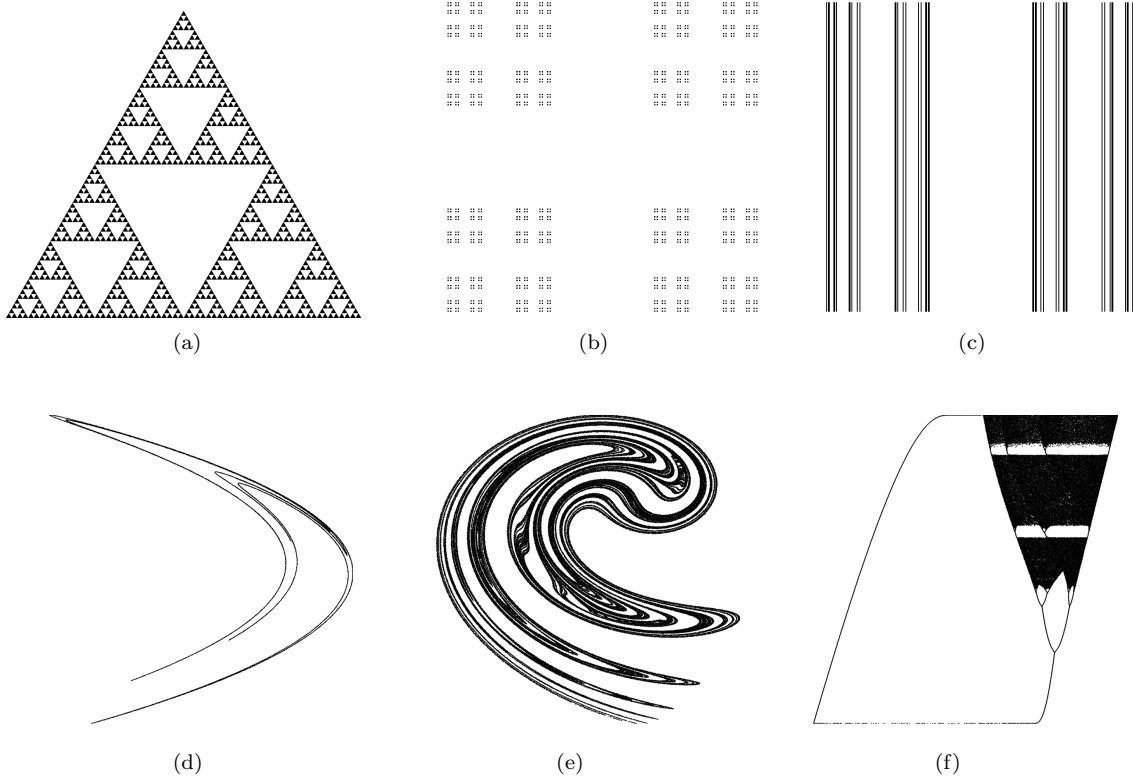


FIGURE 1. Three self-similar fractals: (a) the Sierpinski Triangle, (b) the Cantor dust, and (c) the Cantor set cross an interval; and three chaotic attractors arising from the: (d) Hénon map (e) Ikeda map and (f) Rulkov map.

If  $\mu$  is a measure on bounded metric space,  $\{x_j\}_{j \in \mathbb{N}}$  are i.i.d. samples from  $\mu$ , and  $i$  is a natural number the  $PH_i$ -dimension  $\dim_{PH_i}(\mu)$  measures how  $E_\alpha^i(x_1, \dots, x_n)$  scales as  $n \rightarrow \infty$ . In particular, if  $E_\alpha^i(x_1, \dots, x_n) \approx n^{\frac{d-\alpha}{d}}$ , then  $\dim_{PH_i}(\mu) = d$ . See below for a precise definition.

We estimate the persistent homology dimension of several examples, and compare its performance to classical methods to estimate the correlation and box-counting dimensions. We study the convergence of estimates as the sample size increases, and the variability of the estimate between different samples. The examples we consider are from three broad classes: self-similar fractals, chaotic attractors, and empirical data. The sets in the former class (shown in Figure 1(a)–1(c)) have known dimensions, and are regular in a sense that implies that the various notions of fractal dimension agree for them. That is, they have a single well-defined “fractal dimension.” This makes it easy to compare the performance of the different dimension estimation techniques. In those cases, our previous theoretical results [77] imply that at least the zero-dimensional version of persistent homology dimension will converge to the true dimension.

The second class of fractals we study are strange attractors arising from chaotic dynamical systems, see e.g. Figure 1(d)–1(f). Generically, these sets are not known to be regular, and the various definitions of fractal dimension may disagree for them. Finally, we apply the dimension estimation techniques to the Hauksson–Shearer Southern California earthquake catalog [40].

We also propose that a notion of persistent homology complexity due to MacPherson and Schweinhart [56] may be a good indicator of how difficult it is to estimate the correlation dimension or persistent homology dimension of a shape.

In summary, we conclude the following.

- **Effectiveness.** Based on our experiments, the  $PH_0$  and correlation dimensions perform comparably well. In cases where the true dimension is known, they approach it at about the same rate. In most cases, the box-counting,  $PH_1$ , and  $PH_2$  dimensions perform worse.
- **Efficiency.** Computation of the  $PH_0$  dimension is fast and comparable with the correlation and box-counting dimensions. The  $PH_1$  dimension is reasonably fast subsets of  $\mathbb{R}^2$ , but the  $PH_1$  and  $PH_2$  dimensions are quite slow for point clouds in  $\mathbb{R}^3$ , and computations for sets with higher ambient dimension are impractical.
- **Equivalence.** For a large class of regular fractals the  $PH$  dimension coincides with various classical definitions of fractal dimension. However, there are fractal sets for which these definitions do not agree, with sometimes surprisingly large differences (e.g. the Rulkov and Mackey-Glass attractors). Within the class of  $PH_i$  dimensions, there is variation between dimension estimates coming from different choices of the homological dimension or scaling weight. All this said, it can be difficult to determine whether dimensions truly disagree, or if the convergence is very slow.
- **Error** Error estimates, whether they are the “statistical test error” of the correlation dimension or the empirical standard deviation of estimates between trials, do not contain meaningful information about the difference between the dimension estimate and the true dimension. In general, it is difficult to tell whether a dimension estimate has approached its limiting value.
- **Ease-of-use** We found one simple rule for fitting a power law to estimate the  $PH_0$  which worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension.

In the following, we briefly survey previous work comparing different methods for the estimation of fractal dimension (Section 1.1), outline the different methods considered here (Section 2), and discuss the results for each example (Sections 3, 4, and 5).

**1.1. Background and Previous Work.** Many previous studies have compared different methods for fractal dimension estimation. These include surveys focusing on applications to chaotic attractors [83], medical image analysis [52], networks [70], and time series and spatial data [32]. Fractal dimension estimation techniques have also been applied for intrinsic dimension analysis in applications where an integer-valued dimension is assumed [13, 63].

Several previous studies have observed a relationship between fractal dimension and persistent homology. Estimators based on 0-dimensional persistent homology (minimum spanning trees) were proposed by Weygart et al. [87] and Martinez et al. [61]; we explain their relationship to the current work after Definition 1 below. The PhD thesis of Robins [69], arguably the first publication in the field of topological data analysis, studied persistent Betti numbers of fractals and proved results for the 0-dimensional homology of totally disconnected sets. In 2012, MacPherson and Schweinhart introduced an alternate definition of  $PH$  dimension, which we refer to here as “ $PH_i$  complexity” to avoid confusion (Section 2.1.3). The two previous notions measure the complexity of a shape rather than a classical dimension; they are trivial for  $\mathbb{R}^m$ , for example. A 2019 paper of Adams et al. [1], proposed a  $PH$  dimension that is a special case of the one we consider here, and performed computational experiments on self-similar fractals. However, they do not compare the estimates with those of other fractal dimension estimation techniques. In 2018, Schweinhart [78] proved a relationship between the upper box-counting dimension and the extremal properties of the persistent homology of a metric space, which was a stepping stone to the paper where the current definition was introduced [77]. The latter paper did not undertake a computational study of the persistent homology dimension; we mention relevant theoretical results in Section 2.1.1.

The second class of fractals we study arise in dynamical systems. A dynamical system describes the motion of trajectories  $x(t)$  within a phase space, with a rule for describing how  $x(t)$  changes as time  $t$

increases, and are referred to as either discrete or continuous depending on whether  $t \in \mathbb{Z}_+$  or  $t \in \mathbb{R}_+$ . An attractor  $A$  is an invariant set (that is, if  $x(t) \in A$ , then  $x(s) \in A$  for all  $s > t$ ) in the phase space such that if  $x(t)$  is close to  $A$ , then its distance to  $A$  will decrease to zero asymptotically as  $t \rightarrow \infty$ . Some attractors are simple, consisting of a collection of points or smooth periodic orbits, but some are complicated, such as those shown in Figure 1 (d-f). In dissipative systems *strange attractors* — attractors with a fractional dimension — can occur.

For self-similar fractals, dimension is easy to compute exactly and approximate computationally. The situation for strange attractors is quite different; in general, the fractal dimension of an attractor cannot be computed analytically. The measures on these sets may not be regular, and the different notions of dimension appear to disagree in many cases. Note that when estimates for different notions of fractal dimension disagree, it is difficult to tell whether this is because (1) one method converges much faster than the other (2) there is a significant systematic error in the estimates, or (3) the two definitions are genuinely different.

To offer a further cautionary tale, in [46] the authors study a family of quasi-periodically forced 1D maps with an attracting invariant curve. As a parameter changes, the curve becomes less stable and it becomes increasingly wrinkled; for any resolution, parameters exist for which the curve will appear to be a strange non-chaotic attractor [35]. However, until the invariant curve completely loses stability, the attractor remains a smooth 1-dimensional curve. Hence, no matter how fine a resolution we use to compute our dimension estimate, we can never be sure that our estimate is close to its convergent value.

The Lyapunov exponents of a dynamical system describe its exponential rates of expansion and contraction, and can be used to define the Lyapunov dimension of the attractor [91]. The Kaplan-Yorke conjecture claims that the Lyapunov dimension of a generic dynamical system will equal the information dimension of its attractor [27, 48]. For a class of strange attractors arising from 2D maps, Young showed this to be the case, and moreover that the Lyapunov, Hausdorff, box-counting, and Rényi dimensions all coincide [90]. Other attractors appear to exhibit multifractal properties in the sense that different notions of fractal dimension disagree. The multifractal properties of such an attractor can be studied with a one-parameter family of dimensions such as the generalized Hausdorff [33] or Rényi dimensions [5, 67]. We can also examine this with the  $PH_i$  dimensions by varying the weight parameter  $\alpha$  — see Figure 12.

The Takens embedding theorem also illustrates the importance of fractal dimension to the study of attractors. In an experimental setting, it is often impossible to record all of the relevant dynamic variables. However, one can reconstruct the entire attractor from the time series of a single observed quantity using a time delay embedding [81]. Namely, for a choice of time delay  $\tau$  and embedding dimension  $m$ , a 1-dimensional time series  $\{x_1, \dots, x_N\} \subseteq \mathbb{R}$  may be used to construct a  $m$ -dimensional time series consisting of points  $\{x_i, x_{i-\tau}, \dots, x_{i-(m-1)\tau}\} \in \mathbb{R}^m$ . If  $m$  is at least twice the box-counting dimension of the attractor, then generically this reconstruction will be diffeomorphic to the original attractor [75]. These techniques have been widely applied and we briefly mention some references which use this in the context of topological data analysis [29, 62, 64].

## 2. DIMENSION ESTIMATION METHODS

**2.1. Persistent Homology Dimension.** We give a brief, informal introduction to persistent homology for the special case of a subset of Euclidean space. For a more in depth survey, see [14, 21, 22, 30]. Note that the definition of the  $PH$  dimension also makes sense for subsets of an arbitrary metric space.

Persistent homology [20] quantifies the shape of a geometric object in terms of how the topology of the set changes as it is thickened. To be precise, if  $X$  is a subset of  $\mathbb{R}^m$ , define the family of  $\epsilon$ -neighborhoods  $\{X_\epsilon\}_{\epsilon \in \mathbb{R}^+}$  by

$$X_\epsilon = \{y \in \mathbb{R}^m : d(y, X) < \epsilon\} .$$

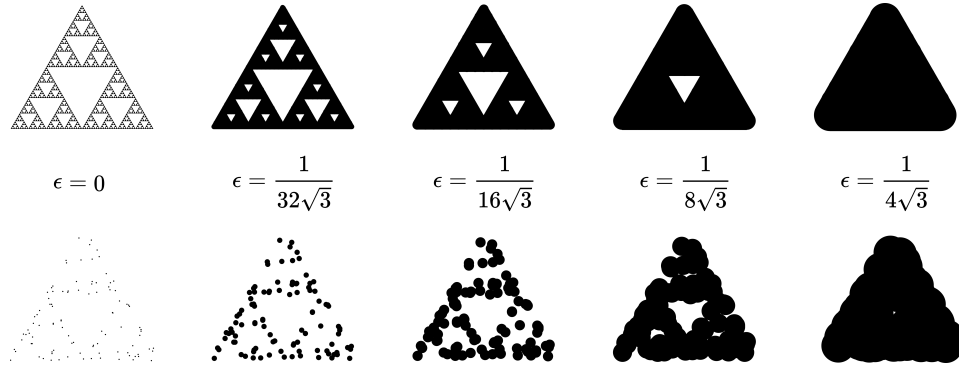


FIGURE 2.  $\epsilon$ -neighborhoods for the Sierpinski triangle (above) and a sample of 100 points from that set (below).

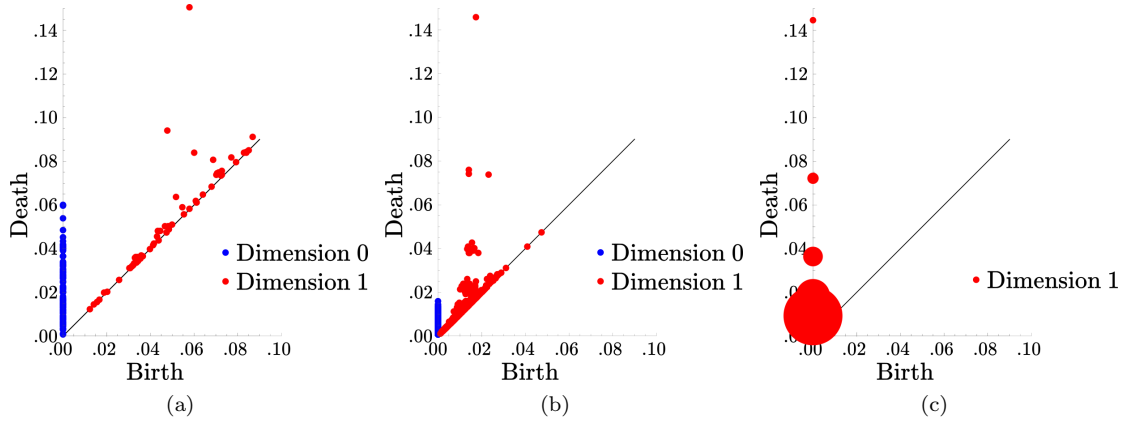


FIGURE 3. Persistence diagrams for (a) 100 point sample from the Sierpinski triangle, (b) 1000 points from the Sierpinski triangle and (c) the Sierpinski triangle itself (where the area of a dot is proportional to the number of persistent homology intervals with corresponding birth and death time).

Figure 2 shows  $\epsilon$ -neighborhoods of the Sierpinski triangle  $S$  and a sample of 100 points from that set. The Sierpinski triangle  $S$  contains infinitely many holes which disappear as we thicken it. The first homology group of  $S_\epsilon$  is an algebraic invariance which essentially counts the number of holes in  $S_\epsilon$ . The last hole to disappear is in the center of the triangle; it vanishes when we have thickened the triangle by  $\epsilon = \frac{1}{4\sqrt{3}}$ . Persistent homology represents this largest hole by the single interval  $(0, \frac{1}{4\sqrt{3}})$ . The next three largest holes disappear at  $\epsilon = \frac{1}{8\sqrt{3}}$  and correspond to three intervals  $(0, \frac{1}{8\sqrt{3}})$ . The following nine holes are represented by nine intervals  $(0, \frac{1}{16\sqrt{3}})$ , and so on. These are the first dimensional persistent homology intervals of the Sierpinski triangle, which we denote by  $PH_1(S)$ .

The zero-dimensional persistent homology of a metric space tracks the connected components that merge together as the geometric object is thickened. The Sierpinski triangle is already connected at  $\epsilon = 0$ , so  $PH_0(S)$  is trivial. However, the finite point sample  $\mathbf{x}$  shown in Figure 2 has 100 components. As we thicken  $\mathbf{x}$  by an amount  $\epsilon$  the first component disappears when  $\epsilon$  equals  $\delta/2$ , where  $\delta$  is the smallest pairwise distance between the points. This corresponds to the interval  $(0, \delta/2)$  in  $PH_0(\mathbf{x})$ . One can find all  $PH_0$  intervals of  $\mathbf{x}$  by increasing  $\epsilon$  and forming an interval whenever there are two points  $x_1, x_2 \in \mathbf{x}$  so that  $d(x_1, x_2) < \epsilon$  and  $x_1$  and  $x_2$  are in different components of  $\mathbf{x}_{\epsilon_0}$  for all  $\epsilon_0 < \epsilon$ .

This is essentially the same as Kruskal’s algorithm for computing the minimum spanning tree on  $\mathbf{x}$ , which leads to a proof that there is a bijection between the edges of that tree and the intervals of  $PH_0(\mathbf{x})$  where an interval corresponds to an edge of twice its length.

Roughly speaking, the higher dimensional homology groups  $H_i(X_\epsilon)$  count the number of higher dimensional “holes” in  $X_\epsilon$ . The higher dimensional persistent homology is defined in terms of how these groups  $H_i(X_\epsilon)$  change as  $\epsilon$  increases. The structure of the persistent homology is captured by a unique set of intervals that track the birth and death of homology generators as  $\epsilon$  changes [12, 14]. We denote this set of intervals by  $PH_i(X)$ .

Traditionally, the information contained in persistent homology is plotted in a persistence diagram showing the scatter plot (birth, death) for each interval. The persistence diagrams of the Sierpinski triangle, and point samples from it with 100 and 1000 points are shown in Figure 3. Note that the 1-dimensional persistent homology of the point sample with 1000 roughly approximates that of the Sierpinski triangle, with 1 interval that dies around  $\epsilon = \frac{1}{4\sqrt{3}}$ , three that die around  $\epsilon = \frac{1}{8\sqrt{3}}$ , and 9 that die around  $\epsilon = \frac{1}{16\sqrt{3}}$ . This is a consequence of the “bottleneck stability” of persistent homology [15].

It is a remarkable fact that if  $\mathbf{x}$  is a finite subset of Euclidean space, then the intervals  $PH_i(\mathbf{x})$  can be computed exactly and efficiently. This is done by replacing the infinite family of  $\epsilon$ -neighborhoods  $\{\mathbf{x}_\epsilon\}_{\epsilon \in \mathbb{R}^+}$  with a finite sequence of finite simplicial complexes, called the Alpha complex of  $\mathbf{x}$ ; these complexes are subcomplexes of the Delaunay triangulation on  $\mathbf{x}$  [20].

Note that there is more than one way to define the persistent homology of a metric space. Here, we use the persistent homology of the Čech complex persistent homology of  $X$ . If  $X$  is a subset of Euclidean space, this is equivalent to the persistent homology of the  $\epsilon$ -neighborhood filtration of  $X$  described above (as well as the persistent homology of the Alpha complex, if  $X$  is finite). Another common (in-equivalent) notion is the persistent homology of the Vietoris–Rips Complex on  $X$  [18, 88]. We use the Čech complex because there are efficient algorithms to compute its persistent homology for finite subsets of  $\mathbb{R}^2$  and  $\mathbb{R}^3$ , as discussed below. A persistent homology dimension defined in terms of the Vietoris–Rips complex also has nice properties [77], and may work better in examples without an embedding into a small-dimensional Euclidean space.

**2.1.1. Definition of the PH dimension.** Consider a sample  $\{x_1, \dots, x_n\}$  of independent points from the natural measure on the Sierpinski triangle. As  $n$  increases, the persistence diagram of the point sample will converge (in the bottleneck distance) to the persistence diagram of the Sierpinski triangle itself. As such, it might seem strange that we can recover the dimension of that set from the 0-dimensional persistent homology of the samples ( $PH_0(S)$  is trivial). However, one may notice in Figure 3 that there is a large cluster of points, both 0- and 1-dimensional, along the plotted diagonal line. These points are generally considered to be “noise” and do not significantly contribute to the larger features of interest in other applications of persistent homology. However, the rate at which this “noise” decays is linked to the dimension of the underlying object.

To track the growth of this “noise”, we define the power-weighted sum for  $\alpha > 0$  by:

$$E_\alpha^i(X) = \sum_{I \in PH_i(X)} |I|^\alpha,$$

where the sum is taken over all finite intervals and  $|I|$  denotes the length of an interval. The scaling properties of random variables of the form  $E_\alpha^0(x_1, \dots, x_n)$  as  $n \rightarrow \infty$  has been studied extensively in probabilistic combinatorics [3, 49, 80, 93], and the case  $i > 0$  has recently been of interest [19, 77]. Motivated by a theorem of Steele [80] and the computational work of Adams et al. [1], Schweinhart [77] introduced the following definition of the persistent homology dimension.

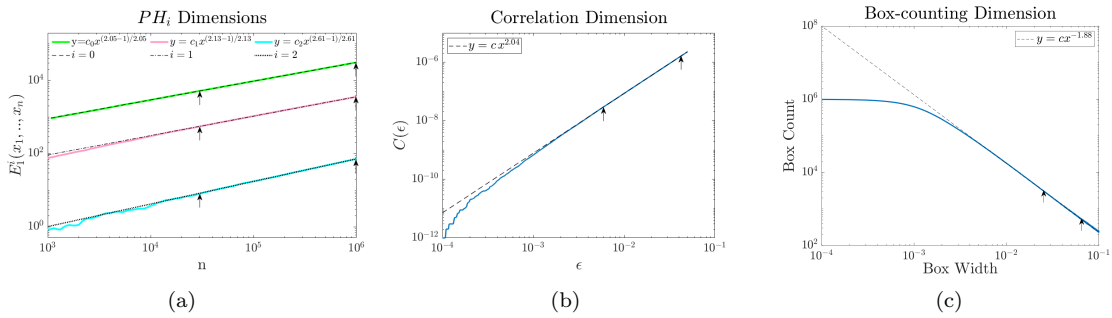


FIGURE 4. Dimension estimation for a sample of  $10^6$  points from the Lorenz attractor, using the methods described in Section 2. Power laws were fitted in the ranges denoted by the arrows.

**Definition 1.** Let  $X$  be a bounded subset of a metric space and  $\mu$  a measure defined on  $X$ . For each  $i \in \mathbb{N}$  and  $\alpha > 0$  we define the persistent homology dimension:

$$\dim_{PH_i^\alpha}(\mu) = \frac{\alpha}{1 - \beta},$$

where

$$\beta = \limsup_{n \rightarrow \infty} \frac{\log \left( \mathbb{E} \left( E_\alpha^i(x_1, \dots, x_n) \right) \right)}{\log(n)}.$$

We write this as the  $PH_i^\alpha$  dimension, and sometimes omit the  $i$  or  $\alpha$  when making general statements.

That is,  $\dim_{PH_i^\alpha}(\mu) = d$  if  $E_\alpha^i(x_1, \dots, x_n)$  scales as  $n^{\frac{d-\alpha}{d}}$ . Larger values of  $\alpha$  give relatively more weight to large intervals than to small ones. The case  $\alpha = 1$  is closely related to the dimension studied by Adams et al. [1], and agrees with it if defined. Weygaert et al. [87] defined a family of minimum spanning tree dimensions that are equivalent to the  $PH_0$  dimensions, and used heuristic arguments to claim that they coincide with the generalized Hausdorff dimensions for chaotic attractors. Martinez et al. [61] asserted that the  $\alpha \rightarrow 0$  of the  $PH_0$  dimension gives the Hausdorff dimension for point samples from chaotic attractors.

It is a corollary of Steele [80] that if  $\mu$  is a non-singular measure on  $\mathbb{R}^m$ , and  $0 < \alpha < m$  then  $\dim_{PH_0^\alpha}(\mu) = m$ . Schweinhart [77] proved that if  $\mu$  satisfies a fractal regularity hypothesis called Ahlfors regularity, then  $\dim_{PH_0^\alpha}(\mu)$  equals the Hausdorff dimension of the support of  $\mu$  (which coincides with the box-counting dimension under the regularity hypothesis). He also proved that if  $d$  equals the upper box-counting dimension of the support of  $\mu$  and  $\alpha < d$  then  $\dim_{PH_0^\alpha}(\mu) \leq d$ , as well as weaker results about the cases where  $i > 0$ .

Note that if  $\mu$  is supported on a  $k$ -dimensional subspace of  $\mathbb{R}^m$ , then  $H_i(X)$  is trivial for  $i \geq k$ . It follows that  $\dim_{PH_i^\alpha}(\mu) = 0$ . As such, even if  $\mu$  is regular, its  $PH_i$  dimension may not equal its Hausdorff dimension unless the Hausdorff dimension is sufficiently large. In particular, if  $\mu$  is a  $d$ -Ahlfors regular measure supported on a 2-dimensional subspace of  $\mathbb{R}^m$ ,  $d > 1.5$ , and  $\alpha < d$  then  $PH_1^\alpha$ -dimension of  $\mu$  equals  $d$ . [77]

Schweinhart's results show that the  $PH_0^\alpha$  and  $PH_1^\alpha$  dimensions of the natural measures on the Sierpinski triangle, Cantor set cross an interval, and Cantor dust equal the Hausdorff dimensions of those sets when  $\alpha$  is less than the true dimension.

**2.1.2. Computation of the PH dimension.** We use different methods to compute persistent homology for the cases  $i = 0$  and  $i > 0$ . In the 0-dimensional case, we use minimum spanning tree-based algorithms that work much more quickly, and are fast even for point samples with high ambient dimension. We use

the implementation of the dual-tree Boruvka Euclidean minimum spanning tree algorithm [59] included in the `mlpack` library [16] to compute the edges of the minimum spanning tree on a point sample. This algorithm was sufficiently fast for point clouds with  $10^6$  points in  $\mathbb{R}^8$ . See [59] for a brief survey of other algorithms to compute minimum spanning trees for point sets in Euclidean space and abstract metric spaces.

For  $i > 0$ , we use GUDHI to compute the persistent homology of the Alpha complex of the point sample. [60, 71] Both this computation and the previous one could be optimized by re-writing the data structures to support insertion of vertices.

Given a sample of  $n$  points  $x_1, \dots, x_n$ , we compute the  $\alpha$ -weighted sum  $E_\alpha^i(x_1, \dots, x_{c_i})$  for 100 logarithmically spaced values of  $c_i$  between 1,000 and  $n$ . Then, we use linear regression to fit a power law to the data  $(c_i, E_\alpha^0(x_1, \dots, x_{c_i}))$ . After some trial-and-error, we found that fitting the power law between  $x = c_n$  and  $x = c_{\lfloor n \rfloor / 2}$  provided a reasonable estimate for all examples tested. Using a smaller range sometimes produced better convergence when  $i > 0$ , but also introduced more oscillations in the estimate. Alternate non-linear regression methods did not seem to result in better performance.

Figure 4(a) shows the power law fits this method produces for a sample from the Lorenz attractor.

**2.1.3. *PH complexity.*** In some instances, the dimension of a metric space can be computed in terms of the persistent homology of the metric space itself. For example, as shown in Figure 3(c), the 1-dimensional persistent homology of the Sierpinski triangle contains intervals that scale as its dimension. This is captured by an alternate notion of *PH* dimension defined by MacPherson and Schweinhart [56], which measures the complexity of the connectivity of the shape rather than the classical dimension. Here, we refer to it as the “*PH* complexity” of a shape to differentiate it from notions of dimension. As we will see below, this quantity may be an indicator of when the dimension is hard to estimate using any of the methods presented here.

If  $X$  is a subset of a metric space, define the cumulative *PH* <sub>$i$</sub>  curve  $F_i$  by

$$F_i(X, \epsilon) = \# \{I \in PH_i(X) : |I| > \epsilon\}$$

Then the *PH* <sub>$i$</sub>  complexity of  $X$  is

$$\text{comp}_{PH_i}(X) = \lim_{\epsilon \rightarrow 0} \frac{-\log(F_i(X, \epsilon))}{\log(\epsilon)}.$$

Note that  $\text{comp}_{PH_i}(\mathbb{R}^n) = 0$  for any  $i$ . Also, if  $S$  is the Sierpinski triangle in Figure 2, our computation of the persistence diagram of  $S$  shows that  $\text{comp}_{PH_0}(S) = 0$  and  $\text{comp}_{PH_1}(S) = \frac{\log(3)}{\log(2)}$ .

We can estimate  $F_i(X, \epsilon)$  from samples; as the Hausdorff distance between  $\{x_1, \dots, x_n\}$  and  $X$  converges to zero, bottleneck stability implies that  $F_i(X, \epsilon)$  will converge for values of  $\epsilon$  large relative to the Hausdorff distance. In Figure 10 below we compute  $\text{comp}_{PH_1}$  of the Ikeda attractor.

**2.2. The Correlation Dimension.** The correlation dimension [34] is commonly used in applications because it is easy to implement and provides reasonable answers even for relatively small sample sizes. A probability measure  $\mu$  on a metric space  $X$  induces a probability measure  $\nu$  on the distance set of  $X$ . Define the correlation integral of  $X$  as the cumulative density function of  $\nu$ :

$$C(\epsilon) = \mathbb{P}(d(x, y) < \epsilon).$$

The correlation dimension equals the limit

$$\lim_{\epsilon \rightarrow 0} \frac{\log(C(\epsilon))}{\log(\epsilon)}$$

if it exists. There is an extensive literature on the estimation and properties of the correlation dimension; see for example [9, 24, 65, 79, 84].



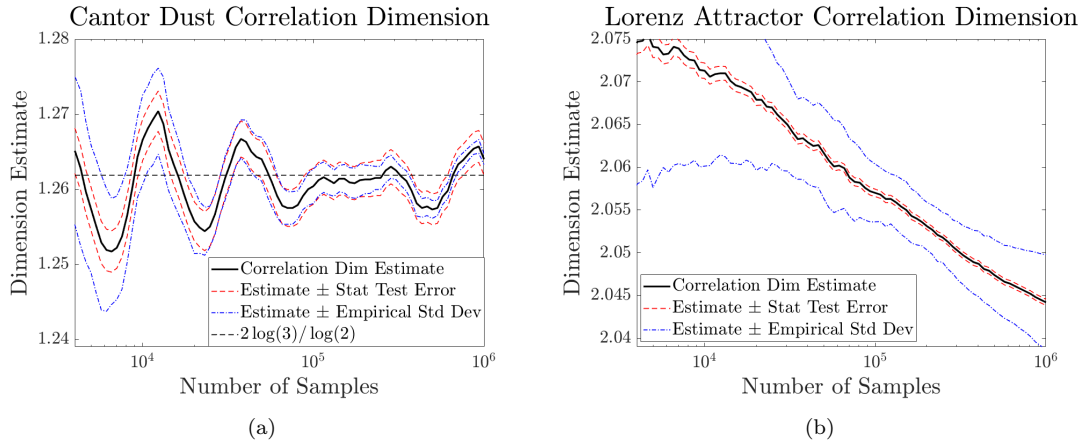


FIGURE 5. Correlation dimension estimates for (a) the Cantor dust and (b) the Lorenz attractor. The evenly dashed red lines indicate the statistical test error, and the unevenly dashed blue lines show the empirical standard deviation across 50 trials.

**2.3. Computation of the Correlation Dimension.** For a finite point sample  $x_1, \dots, x_n$  one can estimate the correlation dimension as

$$\lim_{\epsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{\log C(n, \epsilon)}{\log \epsilon}$$

where

$$C(n, \epsilon) = \frac{\#\{(x_i, x_j) : d(x_i, x_j) < \epsilon, i < j\}}{n(n-1)},$$

if  $n$  is taken to  $\infty$  appropriately as  $\epsilon \rightarrow 0$ . That is,  $C(n, \epsilon)$  measures the number of distances less than  $\epsilon$  in proportion to the number of all inter-point distances.

As with other dimension estimates, this limiting expression converges logarithmically slowly. To accelerate the convergence, we can essentially apply l'Hôpital's rule to find the limit by computing a slope. To do so, we fix a collection of logarithmically spaced values  $\epsilon_1 < \dots < \epsilon_m$ , and compute a linear regression through the data  $(\log \epsilon_i, \log C(n, \epsilon_i))$ . Our estimate for the correlation dimension is then given by the slope of the line of best fit. For a fixed point sample, the values of  $\epsilon_1$  and  $\epsilon_m$  are often chosen by hand to avoid outliers and edge effects.

While computing the  $\mathcal{O}(n^2)$  inter-point distances is prohibitive for very large  $n$ , such a calculation is not needed, as many of these distances are large and do not factor into the dimension calculation [83]. Using a  $kd$ -tree, one can quickly compute the  $\mathcal{O}(n)$  shortest distances with  $\mathcal{O}(n \log n)$  effort. After some trial-and-error, we settled upon the heuristic for choosing  $\epsilon_1$  and  $\epsilon_m$  as below:

$$C(N, \epsilon_1) \approx \frac{n^{.75}}{n(n-1)} \qquad C(N, \epsilon_m) \approx \frac{50n}{n(n-1)}.$$

See Figure 4(b). This heuristic provides a way to choose values of  $\epsilon_1$  and  $\epsilon_m$  consistently amongst different fractals and different sample sizes. It appears to give near-optimal convergence rates to the true dimension for self-similar fractals, and estimates that agree with the values in the literature for the Hénon, Ikeda, and Lorenz attractors [34, 79, 83].

As has been previously reported [83], the “statistical test error” from the linear regression calculation does not have much predictive value about the limiting dimension — see Figure 5. Also, it is much smaller than the empirical standard deviation of the dimension estimate between trials in some cases.

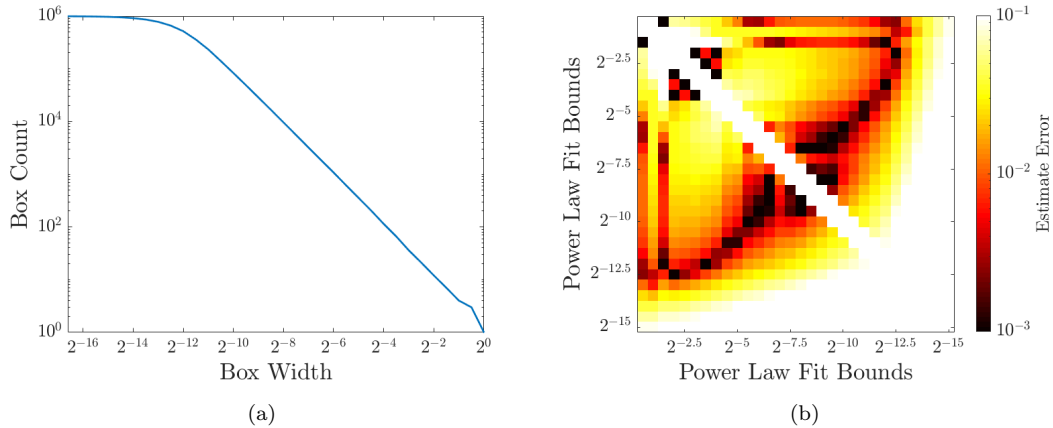


FIGURE 6. (a) Box-counts for a sample of  $10^6$  points from the Sierpinski triangle. (b) Error in the dimension estimate obtained by fitting a power law to the box counts between widths of the form  $2^{-i/2}$  and  $2^{-j/2}$ . Compare to the error of correlation dimension and  $PH_0^1$  dimension estimates in Figure 7(b), which are smaller than .002 for all samples of sizes between  $2 \cdot 10^5$  and  $10^6$  points. Only dark squares would yield comparably good estimates for the box-counting dimension.

**2.4. The Box-counting Dimension.** The box-counting dimension [10] of a bounded subset of  $X$  of  $\mathbb{R}^m$  is defined in terms of the number of cubes of width  $\delta$  needed to cover  $X$ . Let  $\{C_i^\delta\}_{i \in \mathbb{N}}$  be the cubes in the standard tiling of  $\mathbb{R}^m$  by cubes of width  $\delta$ , and let  $N_\delta(X)$  be the number of cubes in  $\{C_i^\delta\}_{i \in \mathbb{N}}$  that intersect  $X$ . Define the upper and lower box-counting dimensions by

$$\dim_{\overline{\text{box}}}(X) = \limsup_{\delta \rightarrow 0} -\frac{\log N_\delta(X)}{\log \delta} \quad \text{and} \quad \dim_{\underline{\text{box}}}(X) = \liminf_{\delta \rightarrow 0} -\frac{\log N_\delta(X)}{\log \delta},$$

respectively. If the upper and lower box-counting dimensions coincide, the shared value is called the box-counting dimension of  $X$  and is denoted  $\dim_{\text{box}}(X)$ . There are several equivalent definitions; see Falconer [25] for details. Many studies have investigated the properties and estimation of the box-counting dimension, including [50, 74, 82].

**2.4.1. Computation of the Box-counting Dimension.** We found it difficult to find a general method to compute the box-counting dimension that worked well for different examples and different numbers of samples. For example, computing box-counts of the form  $N_{2^{-i}}(\mathbf{x})$  and fitting a power law between  $2^{-i}$  and  $2^{-j}$  resulted in estimates that were sensitive to  $i$ ,  $j$ , and the number of samples. In some cases, it was easy to cherry-pick a specific choice just because it seemed to have the best convergence to the true dimension, though other choices resulted in power law fits that looked just as good. See Figure 6 for an example, which plots the error in the dimension estimate for many possible power law fits. Only very specific choices have errors as small as estimates of the correlation and  $PH_0$  dimensions for the same sample.

We settled on the following method, which produced reasonably good results for planar examples. It is based on the observation that if  $\{x_j\}_{j \in \mathbb{N}}$  are samples from  $X \subset \mathbb{R}^m$  and  $\delta > 0$  then the box count  $N_\delta(x_1, \dots, x_m)$  should converge to  $N_\delta(X)$  as  $n \rightarrow \infty$ . For a sample  $\{x_1, \dots, x_m\}$ , we estimate the box-counting dimension from the smallest boxes for which  $N_\delta(x_1, \dots, x_m)$  appears to have stabilized.

For a family of point samples  $\mathbf{x}_m$  in  $\mathbb{R}^n$  (where the sizes of  $\mathbf{x}_m$  are logarithmically spaced as before) we rescale and translate the point samples so they are contained in a unit cube. Then, we

Example	True Dim.	$\text{comp}_{PH_0}(X)$	$\text{comp}_{PH_1}(X)$	$\text{comp}_{PH_2}(X)$
$S$	$\frac{\log(3)}{\log(2)}$	0	$\frac{\log(3)}{\log(2)}$	—
$C \times I$	$1 + \frac{\log(2)}{\log(3)}$	$\frac{\log(2)}{\log(3)}$	0	—
$C \times C$	$\frac{2 \log(2)}{\log(3)}$	$\frac{2 \log(2)}{\log(3)}$	$\frac{2 \log(2)}{\log(3)}$	—
$M$	$\frac{\log(20)}{\log(3)}$	0	$\frac{\log(20)}{\log(3)}$	$\frac{\log(20)}{\log(3)}$

TABLE 1. Data for the four self-similar fractals studied here. As discussed in Section 3.2, lower values of  $\alpha$  appear to yield better dimension estimates when  $\text{comp}_{PH_i}(X) \neq 0$ .

	True	Correl.	Box	$PH_0^5$	$PH_0^1$	$PH_1^5$	$PH_1^1$	$PH_2^1$	$PH_2^2$
$S$	$\approx 1.585$	1.585	1.586	1.585	1.585	1.587	1.620	—	—
$C \times I$	$\approx 1.631$	1.633	1.618	1.629	1.623	1.634	1.642	—	—
$C \times C$	$\approx 1.262$	1.263	1.267	1.263	1.289	1.268	1.303	—	—
$M$	$\approx 2.727$	2.716	2.703	2.705	2.706	2.878	2.773	2.945	2.881

TABLE 2. Dimension estimates for self-similar fractals, averaged over 10 trials of  $10^6$  samples.

compute box-counts of the form  $N_{i/10000}(\mathbf{x}_m)$  for  $1 \leq i \leq 1000$ . We fit a power law to the data  $(i/10000, N_{i/10000}(\mathbf{x}_m))$  in the range  $(\lceil j/2 \rceil, j)$ , where  $j$  is the smallest index so that

$$N_{j/10000}(\mathbf{x}_m) = \min_{1 \leq k \leq 4} (N_{j/10000}(\mathbf{x}_{m-k})) .$$

We used linear regression to fit the power law; non-linear regression did not produce substantially different results. See Figure 4(c).

We tried several variations. For example fitting the power law in the range  $(\lceil .9j \rceil, j)$  produced estimates that converged faster with  $n$  for some examples but exhibited large oscillations in others.

### 3. RESULTS FOR SELF-SIMILAR FRACTALS

We compare the performance of the fractal dimension estimation procedures for four different self-similar fractals: the Sierpinski triangle ( $S$ ), the Cantor dust ( $C \times C$ ), the Cantor set cross an interval ( $C \times I$ ), and the Menger sponge ( $M$ ). The first three are subsets of  $\mathbb{R}^2$ , and  $M$  is contained in  $\mathbb{R}^3$ . We chose these sets to illustrate the observed relationship between  $\text{comp}_{PH_0}(X)$  and the performance of dimension estimation techniques; see Table 1. Definitions of the sets and sampling methods are contained in Appendix A.

For each set, we sample 50 trials of  $10^6$  points from the corresponding natural measure. We compute 100 dimension estimates for each trial, at logarithmically spaced numbers of points between  $10^3$  and  $10^6$ . To show the convergence of the dimension estimate to the true value, we plot number of points against the dimension estimate averaged across the trials, with thinner dotted lines one standard deviation above and below the estimate.

Dimension estimates for the four examples are plotted in Figure 7. In all examples, the  $PH_0$  dimension and the correlation dimension perform better than the  $PH_1$  or box-counting dimensions. The relatively poor performance of the  $PH_1$  dimension can be ascribed to the fact that the weighted sums  $E_\alpha^1(x_1, \dots, x_n)$  are smaller and noisier than the corresponding sums  $E_\alpha^0(x_1, \dots, x_n)$  (see Figure 4(a)). As mentioned previously, it was difficult to find an effective general method to produce box-counting estimates. This is illustrated for the Sierpinski triangle in Figure 6. While the plot of box width vs. box count looks linear for a large range on the log-log plot, only very specific choices of bounds will produce dimension estimates with errors on the same order as the  $PH_0$  and correlation dimensions. These bounds vary unpredictably with the sample size and example.

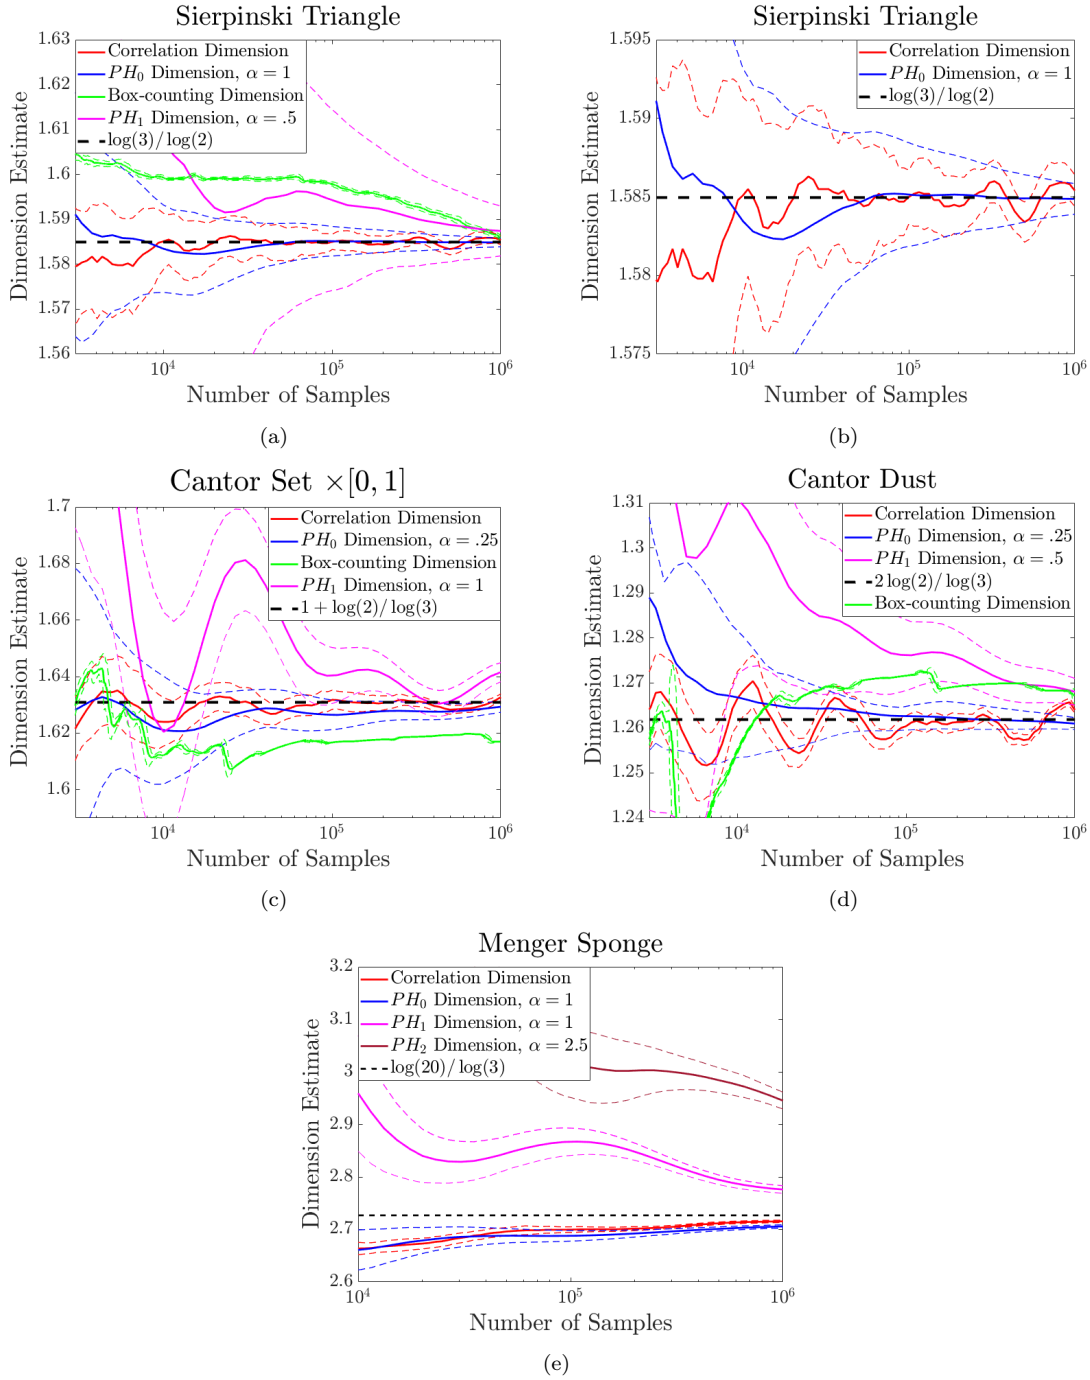


FIGURE 7. Dimension estimates for self-similar fractals. Note that (b) is a close-up of (a). We omitted the box-counting dimension in (e) because the automated procedure described in Section 2.4.1 did not result in stable estimates until  $n$  was larger than  $10^5$ .

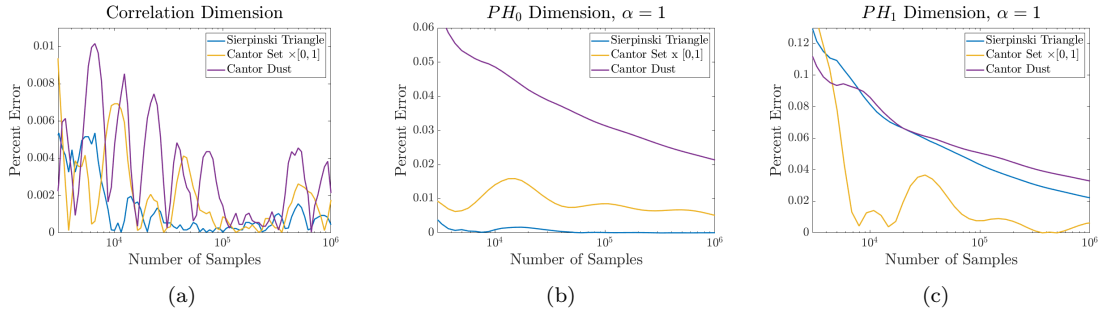


FIGURE 8. Percentage error of dimension estimates for three self-similar fractals.

The convergence of the  $PH_0$  and correlation dimensions to the true dimension is very similar in all four examples. The correlation dimension estimates exhibit oscillations as the sample size increases, a phenomenon often ascribed to the fractal’s lacunarity [83]. The  $PH_0$  dimension estimate appears less sensitive to oscillations, but this is likely due to how the power law fit is performed in the computations (as described in Section 2.1.2); choosing a narrower range over which to fit a power law results in a more oscillatory estimate. That said, for sizable samples, it is impractical to compute the correlation integral over a large range. The  $PH_0$  dimension provides a computationally practicable way to use information from multiple lengthscales. It also has the advantage of having the parameter  $\alpha$  which can be tuned to give a better estimate — see Section 3.2 for a discussion on the choice of  $\alpha$ . The variance of the correlation dimension estimate is slightly lower, but this doesn’t mean much when oscillations are present.

**3.1. Percentage Error.** An interesting pattern emerges when we compare the percentage error of the dimension estimates across the three planar examples, as in Figure 8. Both  $PH_0$  and correlation dimensions perform best for the Sierpinski triangle ( $\text{comp}_{PH_0}(S) = 0$ ), second best for the Cantor set cross an interval ( $\text{comp}_{PH_0}(C \times I) = \dim(C \times I) - 1 = \frac{\log(2)}{\log(3)}$ ), and worst for the Cantor dust ( $\text{comp}_{PH_0}(C \times C) = \dim(C \times C) = \frac{2\log(2)}{\log(3)}$ ). That is, dimension estimation appears to be more difficult when the connectivity of the underlying set is more complex. Note that  $\text{comp}_{PH_0}(X)$  can be computed (see [56]), and could be used as an indicator of whether additional caution is warranted when discussing dimension estimation results.

For the  $PH_1$  dimension, the situation is different and the difficulty of dimension estimation appears to (unsurprisingly) be related to  $\text{comp}_{PH_1}$  rather than  $\text{comp}_{PH_0}$ . The rate of convergence was fastest for the Cantor set cross an interval ( $\text{comp}_{PH_1}(C \times I) = 0$ ), but slower for the Cantor dust ( $\text{comp}_{PH_1}(C \times C) = \dim(C \times C) = \frac{2\log(2)}{\log(3)}$ ) and Sierpinski triangle ( $\text{comp}_{PH_1}(S) = \dim(S) = \frac{\log(3)}{\log(2)}$ ).

We exclude the Menger sponge from these figures, as the ambient dimension likely influences the difficulty of dimension estimation.

**3.2. Dependence on  $\alpha$ .** Estimates of the  $PH_0^\alpha$  and  $PH_1^\alpha$  dimensions for various choices of  $\alpha$  are shown in Figures 9. In the cases where  $\text{comp}_{PH_i}(X) = 0$ , there are only small differences between dimension estimates for different choices of  $\alpha$ , with perhaps a slight advantage for higher values of  $\alpha$  (Figures 9(a) and 9(d)). However, when  $\text{comp}_{PH_i}(X) > 0$ , dimension estimates for different values of  $\alpha$  are substantially different. Lower values of  $\alpha$  yield better estimates when  $\text{comp}_{PH_i}(X) > 0$  for planar examples, and middle values of  $\alpha$  (i.e. equal to about half the true dimension) seem to provide the best convergence for  $i = 1, 2$  (but convergence is slow, especially for  $i = 2$ ). Lower values of  $\alpha$  appear to give dimension estimates that have a higher variance between samples — see Figures 9(g) and 9(g) for the Sierpinski triangle, which is representative.

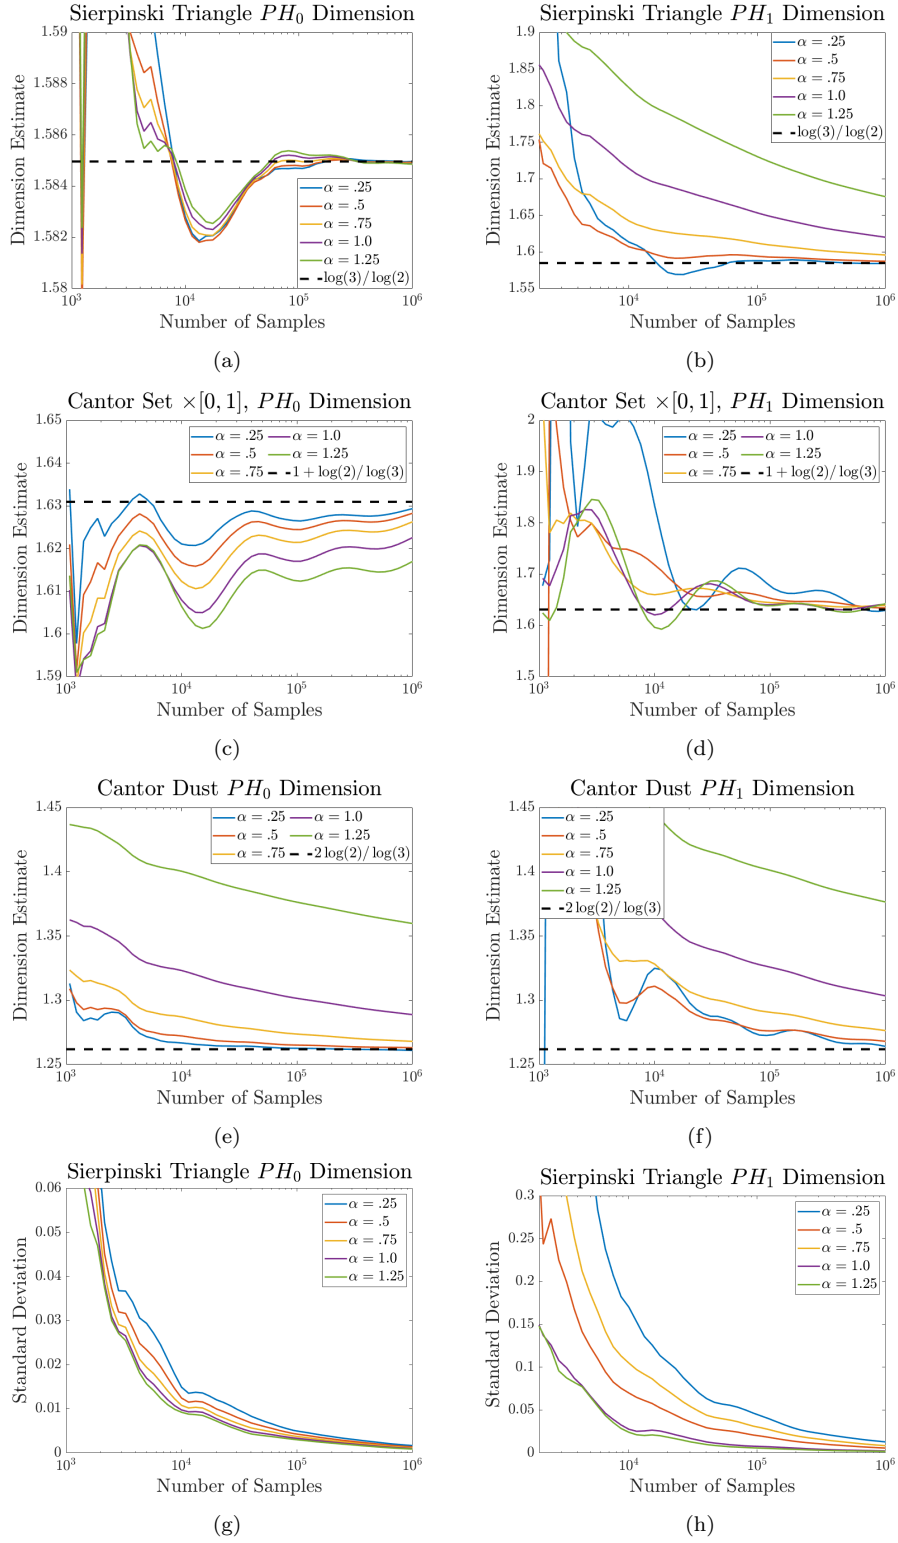


FIGURE 9.  $PH_i$  dimension estimates for various choices of  $\alpha$ . (g) and (h) show the standard deviation of the dimension estimate between trials.

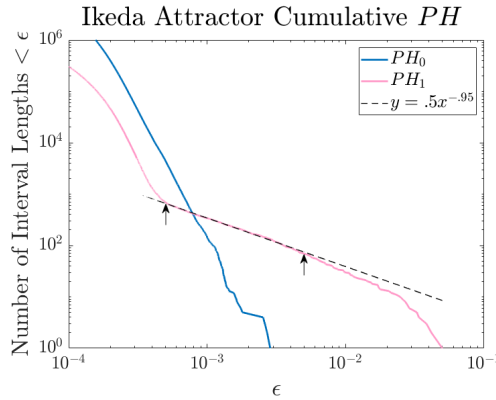


FIGURE 10. To estimate  $\text{comp}_{PH_1}$  of the Ikeda attractor, we fit a power law in the range marked by the arrows. Note that  $PH_1$  exhibits two regimes: noise from the point sample and persistent features of the attractor itself.  $PH_0$  is only noise, as the Ikeda attractor is connected.

	Correlation	Box-counting	$PH_0^5$	$PH_0^1$	$PH_1^5$	$PH_1^1$
Hénon	1.21	1.22	1.28	1.28	1.52	1.36
Ikeda	1.68	1.72	1.71	1.72	1.71	1.72
Rulkov	1.01	1.52	1.62	1.87	2.02	< 2.13
Lorenz	2.04	> 1.90	2.06	2.05	< 2.14	< 2.12
MG	3.04	> 2.45	3.59	3.70	—	—

TABLE 3. Dimension estimates for chaotic attractors, averaged over 10 trials of  $10^6$  samples.

#### 4. RESULTS FOR ATTRACTORS

For each chaotic attractor we sample 50 randomly selected initial conditions, and generate a time series of  $10^6$  points after discarding an initial transient trajectory (except the Mackey-Glass attractor, where we sample 25 initial conditions). For each time series, we compute 100 dimension estimates for each trial, at logarithmically spaced numbers of points between  $10^3$  and  $10^6$ . The various dynamical systems we studied are described in Appendix B and the dimension estimates we obtained for  $N = 10^6$  are summarized in Table 3 and Figure 11.

In all cases, we are skeptical that the box-counting dimension has approached a limiting value. The upper box dimension is known to be an upper bound for the  $PH_0^\alpha$  [77] and correlation dimensions, but most of our estimates of these dimensions tend to be higher than our estimates of the box-counting dimension. This comports with previous observations that box-counting estimates do not seem to converge well for strange attractors [36].

Of the two-dimensional maps we studied, the Ikeda map appears to be most well-behaved. The box-counting,  $PH_0$ , and even the  $PH_1$  dimension all seem to converge to a common value of 1.71, with a correlation dimension of 1.68 not that far off. The Ikeda attractor  $I$  was the only chaotic attractor we studied with non-trivial persistent homology at multiple lengthscales (rather than the persistent homology “noise” of point samples). Here, we compute its  $PH_1$  complexity,  $\text{comp}_{PH_1}(I)$ . Figure 10 shows the cumulative interval count function defined in 2.1.3, for a sample of  $10^6$  points from the Ikeda attractor. The  $PH_0$  curve only shows very small intervals, which come from the noise of point samples. This is as expected, because the Ikeda attractor is connected. However, the  $PH_1$  curve shows two different regimes, one that corresponds to noise and one that appears to be picking up on the elongated holes visible in Figure 1(e). The cumulative length plot appears to follow a power law in the range  $\epsilon = .0005$  to  $\epsilon = .005$ . Fitting a power law to this range gives an estimate of  $\text{comp}_{PH_1}(I) = .95$ . This is

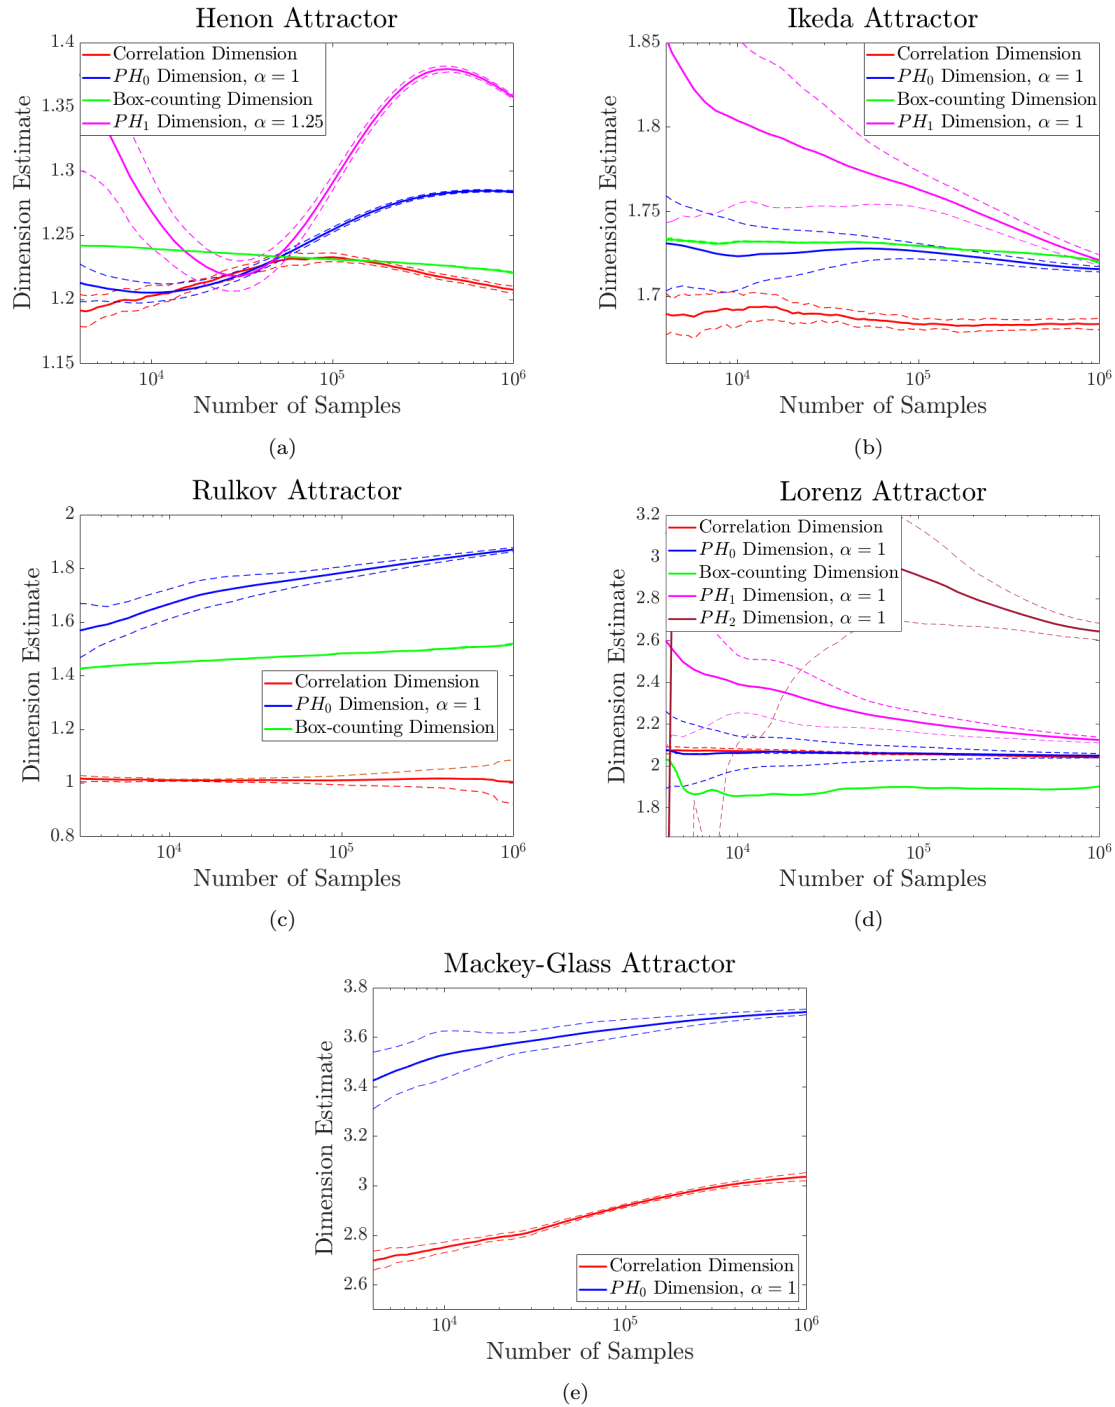
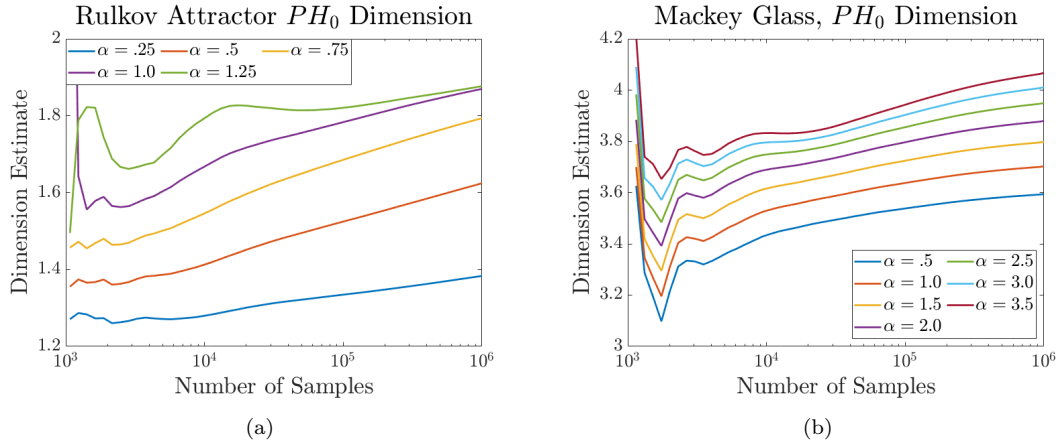


FIGURE 11. Average dimension estimates for various attractors, with dashed lines denoting  $\pm 1$  sample standard deviation. Note the different scales.




 FIGURE 12. Attractor  $PH_i$  dimension estimates for various choices of  $\alpha$ .

lower than the dimension estimates for the Ikeda attractor, indicating that the elongated holes in that set scale at a different rate than its other properties.

For the Hénon attractor, both  $PH_0$  and  $PH_1$  dimension estimates had large oscillations which make it hard to say whether they will converge to values different from either the box-counting or correlation dimensions. The correlation dimension exhibited smaller oscillations, which comports with previous observations by Theiler [85], who also warns of the danger of estimating the dimension of attractors with long-period oscillations. Previous studies [4] have claimed that the Hénon attractor is multifractal for the parameters chosen here, but we are not confident enough in our dimension estimates to make such an assertion.

The Rulkov attractor is non-homogeneous, as easily evidenced by Figure 1(f), and not surprisingly our fractal dimension estimates differed from each other, ranging from 1.00 (correlation dimension) to 2.13 ( $PH_1^1$  dimension). The  $PH_1^\alpha$  estimates began above 2 for small samples sizes and decreased toward 2 as the number of samples increased. For various values of  $\alpha$ , the  $PH_0^\alpha$  dimension estimates ranged from 1.38 for  $\alpha = 1$  to 1.88 for  $\alpha = 1.25$  (see Figure 12). As one may see in Figure 11(c), the variance of the correlation dimension estimate increases with sample size, suggesting the heuristic we chose for fitting the correlation integral requires further fine-tuning for this example. The Rulkov map is noninvertible, so it does not belong to the class of 2D maps considered in [90] where many of the fractal dimensions are known to coincide. While it is difficult to determine if numerical dimension estimates have converged, our results certainly suggest that the various fractal dimension definitions may not agree for the Rulkov attractor.

For the Lorenz attractor, the  $PH_0$  and the correlation dimensions performed well and both appear to converge toward 2.05. However, neither the  $PH_1$  nor the  $PH_2$  dimensions perform well. In fact, nearly  $10^5$  are needed before the  $PH_2$  dimension estimate less than 3, the ambient dimension in which the points reside! As in other cases, this is likely due to the fact that higher dimensional persistent homology contains less information — see Figure 4. The sum of the lengths of the  $PH_1$  and  $PH_2$  intervals of a point sample of  $10^6$  points are 10 and 400 times smaller than the sum of the lengths of the  $PH_0$  intervals, respectively. Also, the plot of  $n$  versus  $E_1^i(x_1, \dots, x_n)$  is much noisier for  $PH_2$  than  $PH_0$ .

Lastly we consider the Mackey-Glass equation, a delay differential equation for which it is still an open conjecture whether there exist parameters for which the system exhibits mathematically provable chaos. The phase space for this system is infinite dimensional, and we use a projection into  $\mathbb{R}^8$  for our fractal dimension calculations. In [26] Farmer reports the attractor's Lyapunov dimension to be 3.58, and

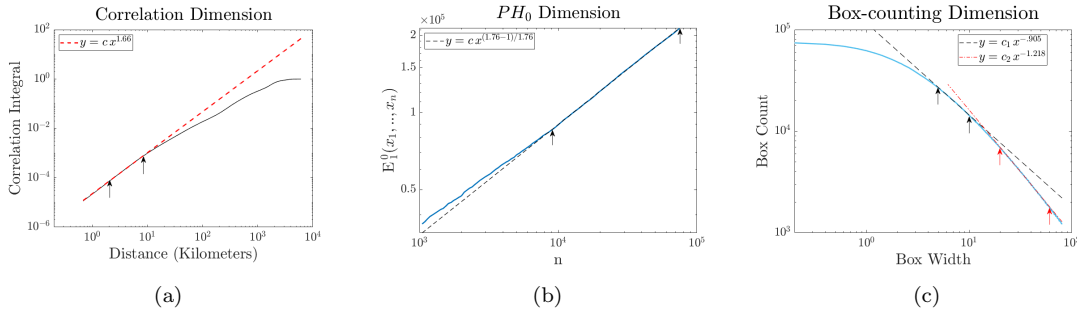


FIGURE 13. Dimension estimation for the earthquake hypocenter data set. Power laws were fitted in the ranges marked by the arrows.

we obtain a  $PH_0^{-5}$  dimension estimate close to this value. However, the spectrum of  $PH_0^\alpha$  dimension for various  $\alpha$  had the largest spread of any of the fractals we studied, except the Rulkov attractor, ranging between 3.59 for  $\alpha = .5$  and 4.06 for  $\alpha = 3.5$  (see Figure 12). While none of our dimension estimates appear to oscillate, all appear to converge slowly. The box-counting dimension estimation method described in Section 2.4.1 performed particularly poorly for this example, perhaps due to the high dimension and co-dimension. The estimate in Table 3 was computed by fitting a power law by hand. We did not attempt to compute either the  $PH_i$  dimensions for  $1 \leq i \leq 7$ , because the algorithm we use is impractical if the ambient dimension is greater than 3.

## 5. EARTHQUAKE DATA

We estimate the dimension of the Hauksson–Shearer Waveform Relocated Southern California earthquake catalog [40,51]. The hypocenter of an earthquake is the location beneath the earth’s surface where the earthquake originates. One can use the dimension estimate to study the geometry of earthquakes, i.e. by comparison to dimension estimates for fractures in rock surfaces or dislocations in crystals (see Section 6 of [47]). Previously, Harte [39] estimated the correlation dimension of earthquake hypocenters in New Zealand and Japan, and Kagan [47] studied the correlation dimension of earthquakes in southern California and developed extensive heuristics to correct for errors in that computation.

We downloaded the coordinates of the hypocenters of the earthquakes in the catalogue from the Southern California Earthquake Data Center Website [76], and selected the waveform-relocated earthquakes of magnitude greater than 2.0. This gave a data set of 74,929 earthquakes. The deepest earthquake in the data set is 33 kilometers below sea level and the depth distribution has a single peak, so we did not split the data set into multiple samples as Harte did. The data set had small relocation errors, with the error in distance between nearby hypocenters estimated to be less than .1 kilometer for 90% of such distances [40].

The presence of location errors means that a different methodology is required to estimate the correlation dimension than from the previously considered examples, where estimates at shorter lengthscales produce better results. To estimate the correlation dimension we plotted the correlation integral (Figure 13(a)) and determined that the best range to fit a power law was between 2 and 8.5 kilometers. This yields a correlation estimate of about 1.66. This can be compared to Kagan’s estimate of 1.5 for earthquakes from an earlier version of the same earthquake catalog that was processed with different methodology (Figure 8 of [47]), and to Harte’s estimates of 1.8 and 1.5 for shallow earthquakes in Kanto, Japan and Wellington, New Zealand, respectively.

Figure 13(b) shows the plot of  $n$  versus  $E_0^1(x_1, \dots, x_n)$  used to estimate the  $PH_0$  dimension. We estimated that the  $PH_0^1$  dimension is approximately 1.76 when  $\alpha = 1$ . Similar computations for  $\alpha = .5$  and  $\alpha = 1.5$  yielded dimension estimates of 1.75 and 1.83, respectively. From this, and the comparison with the correlation dimension, we have evidence that the earthquake probability distribution is not

regular. Also, note that the  $PH_0^1$  dimension plot (Figure 13(b)) appears to follow a power law in a large range, and the dimension can be estimated without fiddling with example-specific parameters. This precludes difficulties that arose in Harte’s analysis where the correlation dimension estimate was very sensitive to the scale at which it was measured (Figure 20 in [39]). Of course, the correlation dimension has an advantage of interpretability at different lengthscales — the  $x$ -axis in Figure 13(a) is kilometers, while it is number of samples in Figure 13(b).

The box-counting,  $PH_1$ , and  $PH_2$  dimensions did not produce good results. In the former case, it is unclear whether the plot of box width versus box count (Figure 13(c)) follows a power law at any range, but certainly no range shorter than 5 kilometers. Fitting between 5 and 10 kilometers yielded an estimate of .91, which is very different than the estimates computed with other methods. Fitting at larger length scales, between 20 and 50 kilometers produced an estimate of 1.22. It is not surprising that the  $PH_1$  and  $PH_2$  dimensions produced poor results, given the small sample size.

We also estimated that  $\text{comp}_{PH_i}(X)$  likely equals zero for the earthquake data, for  $i = 0, 1, 2$ .

## 6. CONCLUSION

Overall the performance of our  $PH_0$  dimension calculations is comparable to the correlation dimension, and gives dimension estimates that are often greater for non-regular sets. Both the  $PH_0$  and correlation dimensions were more reliable than either the box-counting,  $PH_1$ , or  $PH_2$  dimensions. The correlation dimension and the  $PH_0$  dimension are easy to implement and quick to compute. The  $PH_0$  dimension may be more user-friendly in the sense that a single choice of power law range worked well for all examples. Of these two, we do not view one dimension estimation technique as “better” than the other. Rather, they are complementary methods that — if the underlying set is not regular — can provide complementary information.

In many applications of persistent homology, small intervals are often ignored in favor of larger features and discarded as “noise”. However, as we have demonstrated here, there is sometimes signal in that noise. Adding the  $PH$  dimension to the current suite of  $PH$ -based techniques used in applications such as machine learning will offer an orthogonal data descriptor.

## REFERENCES

- [1] Henry Adams, Manuchehr Aminian, Elin Farnell, Michael Kirby, Joshua Mirth, Rachel Neville, Chris Peterson, Patrick Shipman, and Clayton Shonkwiler. A fractal dimension for measures via persistent homology. *Abel Symposia*, 2019.
- [2] Henry Adams, Sofya Chepushtanova, Tegan Emerson, Eric Hanson, Michael Kirby, Francis Motta, Rachel Neville, Chris Peterson, Patrick Shipman, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 2017.
- [3] David Aldous and J. Michael Steele. Asymptotics for Euclidean minimal spanning trees on random points. *Probability Theory and Related Fields*, 1992.
- [4] A. Arneodo, G. Grasseau, and E. Kostelich. Fractal dimensions and  $f(\alpha)$  spectrum of the Hénon attractor. *Physics Letters A*, 1987.
- [5] R. Badii and A. Politi. Renyi dimensions from local expansion rates. *Physical Review A*, 1987.
- [6] James W. Baish and Rakesh K. Jain. Fractals and cancer. *Perspectives in Cancer Research*, 2000.
- [7] Daniel Barbará and Ping Chen. Using the fractal dimension to cluster datasets. In *KDD '00 Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2000.
- [8] Vincent Beffara. The dimension of the SLE curves. *Annals of Probability*, 2008.
- [9] S. Borovkova, R. Burton, and H. Dehling. Consistency of the Takens estimator for the correlation dimension. *Annals of Applied Probability*, 1996.
- [10] M. G. Bouligand. Ensembles impropres et nombre dimensionnel. *Bulletin des Sciences Mathématiques*, 1928.
- [11] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research*, 2015.
- [12] Francesca Cagliaria and Claudia Landibc. Finiteness of rank invariants of multidimensional persistent homology groups. *Applied Mathematics Letters*, 2011.
- [13] Francesco Camastra and Antonino Staiano. Intrinsic dimension estimation: advances and open problems. *Information Sciences*, 2015.
- [14] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. *The Structure and Stability of Persistence Modules*. Springer, 2016.

- [15] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1), 2007.
- [16] R.R. Curtin, M. Edel, M. Lozhnikov, Y. Mentekidis, S. Ghaisas, and S. Zhang. mlpack 3: a fast, flexible machine learning library. *Journal of Open Source Software*, 2018.
- [17] S. Davies and P. Hall. Fractal analysis of surface roughness by using spatial data. *Journal of the Royal Statistical Society Series B*, 1999.
- [18] Vin De Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic and Geometric Topology*, 2007.
- [19] Vincent Divol and Wolfgang Polonik. On the choice of weight functions for linear representations of persistence diagrams. arXiv:1807.03678, July 2018.
- [20] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 2002.
- [21] Herbert Edelsbrunner and John Harer. Persistent homology — a survey. *Contemporary Mathematics*, 2008.
- [22] Herbert Edelsbrunner and Dmitriy Morozov. Persistent homology: Theory and practice. Technical Report LBNL-6037E, Lawrence Berkeley National Laboratory, 2013.
- [23] Gerald A. Edgar. *Classics on Fractals*. Studies in Nonlinearity. Westview Press, 2003.
- [24] Takens F. *Dynamical Systems and Bifurcations*, chapter On the numerical determination of the dimension of an attractor. Springer, 1985.
- [25] Kenneth Falconer. *Fractal Geometry: Mathematical Foundations and Applications*. Wiley, 2014.
- [26] J Doyne Farmer. Chaotic attractors of an infinite-dimensional dynamical system. *Physica D: Nonlinear Phenomena*, 4(3):366–393, 1982.
- [27] Paul Frederickson, James L Kaplan, Ellen D Yorke, and James A Yorke. The Liapunov dimension of strange attractors. *Journal of differential equations*, 49(2):185–207, 1983.
- [28] Marco Gabella, Sabastiano Pavone, , and Giovanni Perona. Errors in the estimate of the fractal correlation dimension of raindrop spatial distribution. *Journal of Applied Meteorology*, 2000.
- [29] Joshua Garland, Elizabeth Bradley, and James D Meiss. Exploring the topology of dynamical reconstructions. *Physica D: Nonlinear Phenomena*, 334:49–59, 2016.
- [30] Robert Ghrist. Barcodes: the persistent homology of data. *Bulletin of the American Mathematical Society*, 2008.
- [31] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences*, 2015.
- [32] Tilmann Gneiting, Hana Sevcikova, and Donald Percival. Estimators of fractal dimension: Assessing the roughness of time series and spatial data. *Statistical Science*, 2012.
- [33] Peter Grassberger. Generalizations of the Hausdorff dimension of fractal measures. *Physics Letters A*, 1984.
- [34] Peter Grassberger and Itamar Procaccia. Measuring the strangeness of strange attractors. *Physica D: Nonlinear Phenomena*, 1983.
- [35] Celso Grebogi, Edward Ott, Steven Pelikan, and James A Yorke. Strange attractors that are not chaotic. *Physica D: Nonlinear Phenomena*, 13(1-2):261–268, 1984.
- [36] H. S. Greenside, A. Wolf, J. Swift, and T. Pignataro. Impracticality of a box-counting algorithm for calculating the dimensionality of strange attractors. *Physical Review A*, 1982.
- [37] Luigi Guzzo, Angela Iovino, Guido Chincarini, Riccardo Giovanelli, and Martha P. Haynes. Scale-invariant clustering in the large-scale distribution of galaxies. *Astrophysical Journal, Part 2 - Letters*, 1991.
- [38] SM Hammel, CKRT Jones, and Jerome V Moloney. Global dynamical behavior of the optical field in a ring cavity. *JOSA B*, 2(4):552–564, 1985.
- [39] D. Harte. Dimension estimates of earthquake epicentres and hypocentres. *Journal of Nonlinear Science*, 1998.
- [40] Egill Hauksson, Wenzheng Yang, and Peter M. Sheare. Waveform relocated earthquake catalog for southern California (1981 to june 2011)r. *Bulletin of the Seismological Society of America*, 2012.
- [41] Felix Hausdorff. Dimension und äußeres maß. *Mathematische Annalen*, 1918.
- [42] Michel Hénon. A two-dimensional mapping with a strange attractor. In *The Theory of Chaotic Attractors*, pages 94–102. Springer, 1976.
- [43] Yasuaki Hiraoka, Takenobu Nakamura, Akihiko Hirata, Emerson G Escobar, Kaname Matsue, and Yasumasa Nishiura. Hierarchical structures of amorphous solids characterized by persistent homology. *Proceedings of the National Academy of Sciences*, 2016.
- [44] Song Hu, Xuexing Sun, Jun Xiang, Min Li, Peisheng Li, and Liqi Zhang. Correlation characteristics and simulations of the fractal structure of coal char. *Communications in Nonlinear Science and Numerical Simulation*, 2003.
- [45] Borja Ibarz, José Manuel Casado, and Miguel AF Sanjuán. Map-based models in neuronal dynamics. *Physics reports*, 501(1-2):1–74, 2011.
- [46] Àngel Jorba and Joan Carles Tatjer. A mechanism for the fractalization of invariant curves in quasi-periodically forced 1-d maps. *Discrete & Continuous Dynamical Systems-B*, 10:537–567, 2008.
- [47] Yan Y. Kagan. Earthquake spatial distribution: the correlation dimension. *Geophysical Journal International*, 2007.
- [48] James L Kaplan and James A Yorke. Chaotic behavior of multidimensional difference equations. In *Functional differential equations and approximation of fixed points*, pages 204–227. Springer, 1979.

- [49] Harry Kesten and Sungchul Lee. The central limit theorem for weighted minimal spanning trees on random points. *Annals of Applied Probability*, 1996.
- [50] Larry S. Liebovitch and Tibor Toth. A fast algorithm to determine fractal dimensions by box counting. *Physics Letters A*, 1989.
- [51] G. Lin, P. M. Shearer, and E. Hauksson. Applying a three-dimensional velocity model, waveform cross correlation, and cluster analysis to locate southern California seismicity from 1981 to 2005. *Journal of Geophysical Research: Solid Earth*, 2007.
- [52] R. Lopes and N. Betrouni. Fractal and multifractal analysis: A review. *Medical Image Analysis*, 2009.
- [53] Edward N Lorenz. Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2):130–141, 1963.
- [54] Shaun Lovejoy and Daniel Schertzer. Fractals, raindrops and resolution dependence of rain measurements. *Journal of Applied Meteorology*, 1990.
- [55] Michael C Mackey and Leon Glass. Oscillation and chaos in physiological control systems. *Science*, 197(4300):287–289, 1977.
- [56] R. D. MacPherson and B. Schweinhart. Measuring shape with topology. *Journal of Mathematical Physics*, 53(7), 2012.
- [57] Benoît Mandelbrot. *Fractals: Form, Chance and Dimension*. W.H.Freeman & Company, 1977.
- [58] Benoît Mandelbrot. *The Fractal Geometry of Nature*. W. H. Freeman and Co., 1982.
- [59] William B. March, Parikshit Ram, and Alexander G. Gray. Fast Euclidean minimum spanning tree: Algorithm, analysis, and applications. *KDD '10 Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2010.
- [60] Clément Maria. Persistent cohomology. In *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015.
- [61] Vicent J. Martinez, R. Dominguez-Tenreiro, and L. J. Roy. Hausdorff dimension from the minimal spanning tree. *Physical Review E*, 1992.
- [62] Konstantin Mischaikow, Marian Mrozek, J Reiss, and Andrzej Szymczak. Construction of symbolic dynamics from experimental time series. *Physical Review Letters*, 82(6):1144, 1999.
- [63] Dengyao Mo and Samuel H. Huang. Fractal-based intrinsic dimension estimation and its application in dimensionality reduction. *IEEE Transactions on Knowledge and Data Engineering*, 2012.
- [64] Audun Myers, Elizabeth Munch, and Firas A Khasawneh. Persistent homology of complex networks for dynamic state detection. *arXiv preprint arXiv:1904.07403*, 2019.
- [65] M. A. H. Nerenberg and Christopher Essex. Correlation dimension and systematic geometric effects. *Physical Review A*, 1990.
- [66] Xin Peng, Wei Qi, Mengfan Wang, Rongxin Su, and Zhimin He. Backbone fractal dimension and fractal hybrid orbital of protein structure. *Communications in Nonlinear Science and Numerical Simulation*, 2013.
- [67] Alfred Renyi. *Probability Theory*. North Holland Publishing Company, 1970.
- [68] Michel Rieu and Garrison Sposito. Fractal fragmentation, soil porosity, and soil water properties: I. theory. *Soil Science Society of America Journal*, 1991.
- [69] Vanessa Robins. *Computational topology at multiple resolutions: foundations and applications to fractals and dynamics*. PhD thesis, University of Colorado at Boulder, 2000.
- [70] Eric Rosenberg. *A Survey of Fractal Dimensions of Networks*. Springer, 2018.
- [71] Vincent Rouvreau. Alpha complex. In *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015.
- [72] Nikolai F Rulkov. Regularization of synchronized chaotic bursts. *Physical Review Letters*, 86(1):183, 2001.
- [73] Mohammad Saadatfar, Hiroshi Takeuchi, Vanessa Robins, Nicolas Francois, and Yasuaki Hiraoka. Pore configuration landscape of granular crystallization. *Nature Communications*, 8, 2017.
- [74] N. Sarkar and B.B. Chaudhuri. An efficient differential box-counting approach to compute fractal dimension of image. *IEEE Transactions on Systems, Man, and Cybernetics*, 1994.
- [75] Tim Sauer, James A Yorke, and Martin Casdagli. Embedology. *Journal of statistical Physics*, 65(3-4):579–616, 1991.
- [76] SCEDC. Southern california earthquake center., 2013. Caltech.Dataset.
- [77] B. Schweinhart. The persistent homology of random geometric complexes on fractals. arXiv:1808.02196.
- [78] Benjamin Schweinhart. Persistent homology and the upper box dimension. arXiv:1802.00533, 2018.
- [79] Julien Clinton Sprott and George Rowlands. Improved correlation dimension calculation. *International Journal of Bifurcation and Chaos*, 11(07):1865–1880, 2001.
- [80] J. Michael Steele. Growth rates of Euclidean minimal spanning trees with power weighted edges. *Annals of Probability*, 1988.
- [81] Floris Takens. *Dynamical Systems and Turbulence*, chapter Detecting Strange Attractors in Turbulence, pages 366–381. Springer, 1980.
- [82] Charles C. Taylor and James Taylor. Estimating the dimension of a fractal. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1991.
- [83] James Theiler. Estimating fractal dimension. *Journal of the Optical Society of America A*, 1990.
- [84] James Theiler. Statistical precision of dimension estimators. *Physical Review A*, 1990.
- [85] James Patrick Theiler. *Quantifying Chaos: Practical Estimation of the Correlation Dimension*. PhD thesis, California Institute of Technology, 1988.

[86] Caetano Traina, Agma Traina, Leejay Wu, and Christos Faloutsos. Fast feature selection using fractal dimension. *Journal of Information and Data Management*, 2010.

[87] Rien van de Weygaert, Bernard J.T. Jones, and Vincent J. Martinez. The minimal spanning tree as an estimator for generalized dimensions. *Physics Letters A*, 1992.

[88] Leopold Vietoris. Über den höheren zusammenhang kompakter räume und eine klasse von zusammenhangstreuen abbildungen. *Mathematische Annalen*, 1927.

[89] Kelin Xia and Guo-Wei Wei. Persistent homology analysis of protein structure, flexibility, and folding. *International journal for numerical methods in biomedical engineering*, 2014.

[90] Lai-Sang Young. Dimension, entropy and Lyapunov exponents. *Ergodic theory and dynamical systems*, 2(1):109–124, 1982.

[91] Lai-Sang Young. Mathematical theory of Lyapunov exponents. *Journal of Physics A: Mathematical and Theoretical*, 46(25):254001, 2013.

[92] Boming Yu. Analysis of flow in fractal porous media. *Applied Mechanics Reviews*, 2008.

[93] J.E. Yukich. Asymptotics for weighted minimal spanning trees on random points. *Stochastic Processes and their Applications*, 2000.

APPENDIX A. SELF-SIMILAR FRACTALS

**A.1. The Sierpinski Triangle.** The Sierpinski triangle (Figure 1(a)) is defined by iteratively removing equilateral triangles from a larger equilateral triangle. Begin by sub-dividing the equilateral triangle formed by  $(0, 0)$ ,  $(1, 0)$ , and  $(\sqrt{3}/2, 1/2)$  into four congruent triangles, and removing the center triangle. Repeat this process on the remaining three triangles, and continue ad infinitum. The Sierpinski triangle equals three copies of itself rescaled by the factor  $1/2$  so the dimension of the resulting set is  $\frac{\log(3)}{\log(2)}$  (for a self-similar set equal to  $m$  copies of itself rescaled by a factor  $r$ , the self-similarity dimension equals  $\frac{\log(m)}{\log(1/r)}$ ).

We sampled points from the natural measure on the Sierpinski Triangle by sampling random integers  $\{a_1, a_2, \dots\} \in \{0, 1, 2\}$  and computing

$$(x, y) = \left( \sum_{i:a_i=1} 2^{-i} + \sum_{i:a_i=2} 2^{-i-1}, \sum_{i:a_i=2} \sqrt{3} 2^{-i-1} \right).$$

In practice, we end this procedure at  $i = 64$ .

**A.2. The Cantor Dust.** The standard middle-thirds Cantor set  $C$  is the set formed by removing the interval  $(1/3, 2/3)$  from the closed interval  $[0, 1]$ , and iteratively removing the middle third of the remaining intervals. The dimension of  $C$  is  $\frac{\log(2)}{\log(3)}$ . The Cantor dust (Figure 1(b)) is the product  $C \times C$ ; its dimension is  $\frac{2\log(2)}{\log(3)}$ .

We sampled points from the natural measure on  $C$  by sampling random integers  $\{a_1, a_2, \dots\} \in \{0, 1\}$  and computing

$$x = 2 \sum_{i=1}^{\infty} a_i 3^{-i}.$$

In practice, we truncated the summation at  $i = 64$ . A point  $(x_1, x_2)$  from the natural measure on the Cantor dust can then be sampled by independently sampling  $x_1$  and  $x_2$  as above.

**A.3. The Cantor Set Cross an Interval.** Consider the set  $C \times [0, 1]$ , where  $C$  is the Cantor set defined in the previous section. The dimension of this set is  $1 + \frac{\log(2)}{\log(3)}$ , one greater than the dimension of the Cantor set.

We sampled points  $(x, y)$  from the natural measure on the Cantor set cross an interval by sampling a random point  $x$  on the Cantor set by the procedure described in the previous section, and a random real number  $y$  from the uniform distribution on  $[0, 1]$ .

**A.4. The Menger Sponge.** The Menger sponge is defined by iteratively removing cubes from a larger cube. Begin by sub-dividing the unit cube in  $\mathbb{R}^3$  into nine smaller sub-cubes, and remove nine of them: one from the center of the original cube one from the center of each of the eight faces. Repeat this process on each of the 20 remaining sub-cubes and continue ad infinitum. The resulting structure is equal to 20 copies of itself rescaled by a factor  $1/3$  so the dimension of the resulting set is  $\frac{\log(3)}{\log(20)}$ .

For  $j \in \infty$  we sample a uniform three-tuple of integers  $(x_i, y_i, z_i) \in \{0, 1, 2\}$  throwing out and re-sampling any three-tuple for which two or more of the coordinates equals one. To sample a point from the natural measure on the Menger sponge, we form the sum

$$(x, y, z) = \sum_{i=0}^{\infty} 3^{-i} (x_i, y_i, z_i) .$$

In practice, we end this procedure at  $i = 64$ .

## APPENDIX B. CHAOTIC ATTRACTORS

**B.1. Hénon Map.** The Hénon map [42] is given by  $(x_n, y_n) \mapsto (x_{n+1}, y_{n+1})$  where:

$$\begin{aligned} x_{n+1} &= 1 - ax_n^2 + y_n \\ y_{n+1} &= bx_n. \end{aligned}$$

We used the parameters of  $a = 1.4$  and  $b = 0.3$ . We generated 50 time series using randomly chosen initial conditions, computing a trajectory of length  $1.1 \cdot 10^7$  and discarding the initial  $10^6$  points in the series.

**B.2. Ikeda Map.** The complex Ikeda map [38] is given by  $z_n \mapsto z_{n+1}$  where:

$$z_{n+1} = a + R \exp \left[ i \left( \phi - \frac{p}{1 + |z_n|^2} \right) \right] z_n.$$

We used the parameters:

$$a = 1 \qquad R = 0.9 \qquad \phi = 0.4 \qquad p = 6.$$

We generated 50 time series using randomly chosen initial conditions, computing a trajectory of length  $1.1 \cdot 10^6$  and discarding the initial  $10^5$  points in the series.

**B.3. Rulkov Map.** The chaotic Rulkov map [45,72] is given by  $(x_n, y_n) \mapsto (x_{n+1}, y_{n+1})$  where:

$$\begin{aligned} x_{n+1} &= \frac{\alpha}{1 + x_n^2} + y_n \\ y_{n+1} &= y_n - \mu(x_n - \sigma) \end{aligned}$$

We use the parameters

$$\mu = 0.0001 \qquad \alpha = 3.75 \qquad \sigma = -1 \qquad I = 0$$

We generated 50 time series using randomly chosen initial conditions, computing a trajectory of length  $1.1 \cdot 10^6$  and discarding the initial  $10^5$  points in the series.

**B.4. Lorenz System.** The Lorenz system [53] given by the system of ordinary differential equations below:

$$\begin{aligned} \dot{x} &= \sigma(y - x) \\ \dot{y} &= x(\rho - z) - y \\ \dot{z} &= xy - \beta z \end{aligned}$$

We use the parameters:

$$\rho = 28 \qquad \sigma = 10 \qquad \beta = 8/3.$$

We generated 50 time series using randomly chosen initial conditions, integrating forward for  $1.01 \times 10^5$  units of time, and discarding the initial  $10^3$  units of time. We used MATLAB's `ode45` integrator with a relative error tolerance of  $10^{-9}$  and an absolute error tolerance of  $10^{-9}$ . We sampled our trajectories at a rate of 10Hz, producing time series with  $10^6$  points each.

**B.5. Mackey-Glass.** The Mackey-Glass equation [55] is given by the scalar delay differential equation below:

$$(1) \quad y'(t) = -ay(t) + b \frac{y(t-\tau)}{1 + y(t-\tau)^n}.$$

We take parameters:

$$a = 1 \qquad b = 2 \qquad \tau = 3 \qquad n = 10.$$

A natural phase space for this dynamical system is the Banach space of continuous functions  $C = C^0([-\tau, 0], \mathbb{R})$ . For our dimension calculations we discretize this space using a projection map  $\pi_m : C \rightarrow \mathbb{R}^m$  (for  $m = 8$ ) which evaluates a function  $y \in C$  at  $m$  maximally spaced points on the interval  $[-\tau, 0]$ .

We generated 25 time series using randomly chosen initial conditions and integrating forward for  $\frac{\tau}{m-1} \cdot 10^6 + 10^4$  units of time. We used MATLAB's `dde23` integrator with the default error tolerances for the first 5,000 units of time, and the remainders of each trajectory were computed using a relative error tolerance of  $10^{-5}$  and an absolute error tolerance of  $10^{-9}$ . The initial transient periods of  $10^4$  units of time were discarded, and the remainders were sampled at a rate of  $\frac{\tau}{m-1}$ Hz, producing time series with  $10^6$  points each.

(J. Jaquette) DEPARTMENT OF MATHEMATICS, BRANDEIS UNIVERSITY WALTHAM, MA 02453

*Email address:* `jjaquette@brandeis.edu`

(B. Schweinhart) DEPARTMENT OF MATHEMATICS, OHIO STATE UNIVERSITY COLUMBUS, OH 43210

*Email address:* `schweinhart.2@osu.edu`