Fractal Dimension Estimation with Persistent Homology

J. Jaquette [1]

Brandeis University

CRM Applied Mathematics Seminar

November 11, 2019

[1] This material is based upon work supported by grants 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.

From topology to dynamics and back again

2017 *On ε approximations of persistence diagrams*
J. & Miroslav Kramar

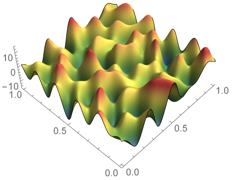

2018 *Counting and Discounting Slowly Oscillating Periodic Solutions to Wright's Equation*
J. et al.

2019 *Fractal Dimension Estimation with Persistent Homology: A Comparative Study*
J. & Benjamin Schweinhart

Presentation Outline

1. Background: Fractal Dimension and Persistent Homology
2. Previous Work and Definitions
3. Computational Results

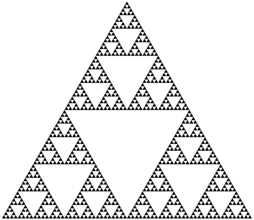

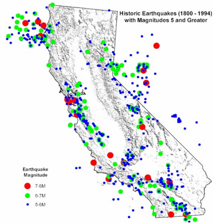

- Fractal dimension measures how the properties of a shape depend on scale.
- The first notion of a fractional dimension was proposed by Hausdorff in 1918. Since then, several other definitions have been proposed, including the box-counting, packing, and correlation dimensions.
- These dimensions agree on a wide class of “regular” sets.

Filtrations

$\epsilon = 0$ $\qquad \epsilon = \frac{1}{32\sqrt{3}}$ $\qquad \epsilon = \frac{1}{16\sqrt{3}}$ $\qquad \epsilon = \frac{1}{8\sqrt{3}}$ $\qquad \epsilon = \frac{1}{4\sqrt{3}}$

A **filtration** of topological spaces is family $\{X_\alpha\}_{\alpha \in I}$ with ordered index set I, together with inclusions $i_{\alpha,\beta} : X_\alpha \hookrightarrow X_\beta$ for $\alpha < \beta$.

Example: if $S \subset \mathbb{R}^2$ we have the filtration of ϵ-neighborhoods $\{S_\epsilon\}_{\epsilon \in \mathbb{R}^+}$.

Filtrations

$\epsilon = 0$ $\quad$ $\epsilon = \frac{1}{32\sqrt{3}}$ $\quad$ $\epsilon = \frac{1}{16\sqrt{3}}$ $\quad$ $\epsilon = \frac{1}{8\sqrt{3}}$ $\quad$ $\epsilon = \frac{1}{4\sqrt{3}}$

A **filtration** of topological spaces is family $\{X_\alpha\}_{\alpha \in I}$ with ordered index set I, together with inclusions $i_{\alpha,\beta} : X_\alpha \hookrightarrow X_\beta$ for $\alpha < \beta$.
Example: if $S \subset \mathbb{R}^2$ we have the filtration of ϵ-neighborhoods $\{S_\epsilon\}_{\epsilon \in \mathbb{R}^+}$.

Persistent Homology

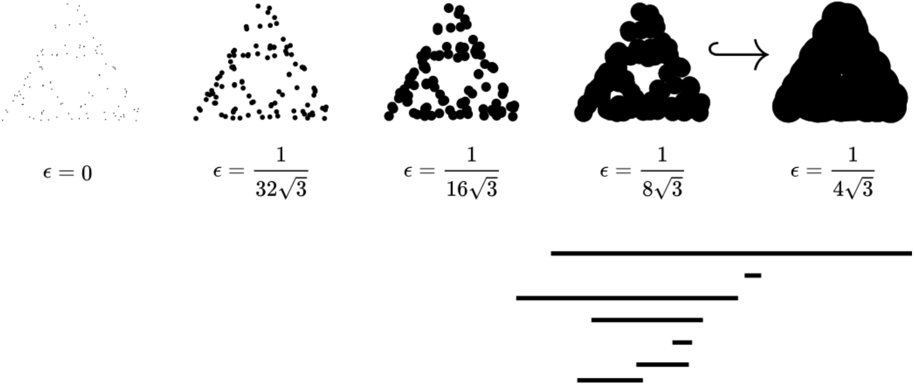

Persistent Homology (PH) tracks how the homology changes through a filtration. PH_i is a set of intervals corresponding to homology generators that are born and die in this process.

Definition of Persistent Homology

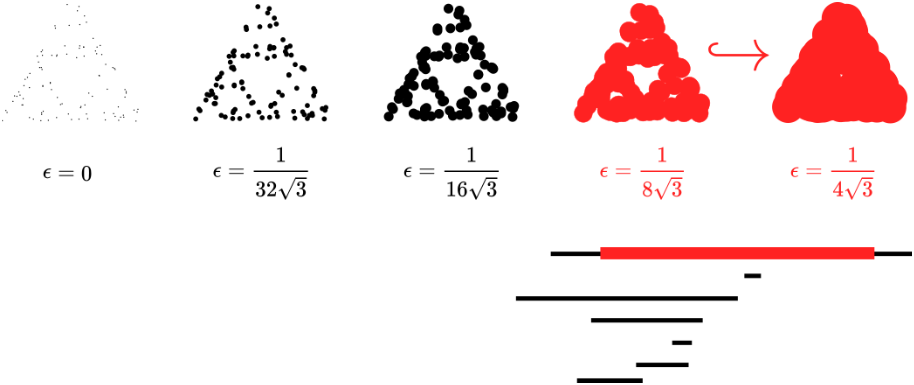

Given a filtration $\{X_\alpha\}$, $PH_i(X_\alpha)$ is the unique set of intervals so that

$$\mathrm{rank}(i_{\alpha,\beta} : H_i(X_\alpha) \to H_i(X_\beta)) = \#\left\{ I \in PH_i(X_\alpha) : [\alpha, \beta] \subseteq I \right\}.$$

Persistent Homology

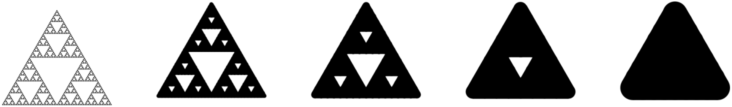

$\epsilon = 0$ $\quad \epsilon = \frac{1}{32\sqrt{3}}$ $\quad \epsilon = \frac{1}{16\sqrt{3}}$ $\quad \epsilon = \frac{1}{8\sqrt{3}}$ $\quad \epsilon = \frac{1}{4\sqrt{3}}$

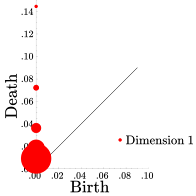

The information in *PH* is often summarized by a persistence diagram: a scatter plot of (birth, death) for each interval.

Persistent Homology

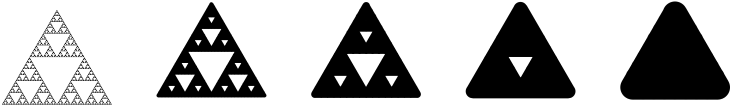

$\epsilon = 0$ $\quad \epsilon = \frac{1}{32\sqrt{3}}$ $\quad \epsilon = \frac{1}{16\sqrt{3}}$ $\quad \epsilon = \frac{1}{8\sqrt{3}}$ $\quad \epsilon = \frac{1}{4\sqrt{3}}$

Infinitely many holes at $\epsilon = 0$

Infinitely many intervals with Birth=0

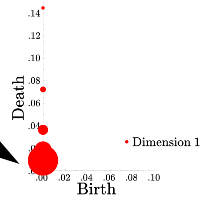

$PH_1(S_\epsilon)$ is a set of intervals, one for each component of the complement that disappears as ϵ increases (by Alexander duality).

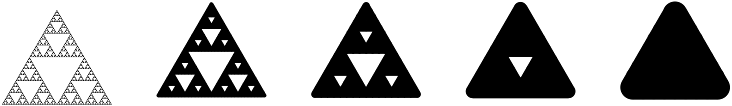

$\epsilon = 0$ $\quad \epsilon = \frac{1}{32\sqrt{3}}$ $\quad \epsilon = \frac{1}{16\sqrt{3}}$ $\quad \epsilon = \frac{1}{8\sqrt{3}}$ $\quad \epsilon = \frac{1}{4\sqrt{3}}$

1 hole disappears at $\epsilon = \frac{1}{4\sqrt{3}}$

1 interval $(0, \frac{1}{4\sqrt{3}})$

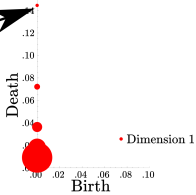

$PH_1(S_\epsilon)$ is a set of intervals, one for each component of the complement that disappears as ϵ increases (by Alexander duality).

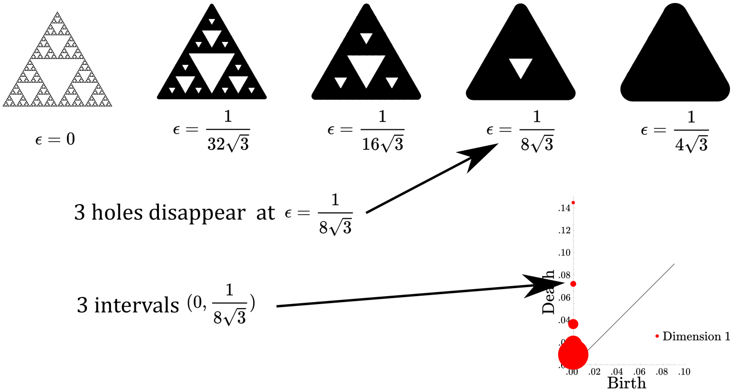

$PH_1(S_\epsilon)$ has one interval for each bounded component of the complement of S (by Alexander duality).

Persistent Homology of a Sample

$\epsilon = 0$ $\quad$ $\epsilon = \frac{1}{32\sqrt{3}}$ $\quad$ $\epsilon = \frac{1}{16\sqrt{3}}$ $\quad$ $\epsilon = \frac{1}{8\sqrt{3}}$ $\quad$ $\epsilon = \frac{1}{4\sqrt{3}}$

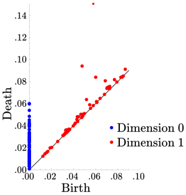

Persistent Homology of a Sample

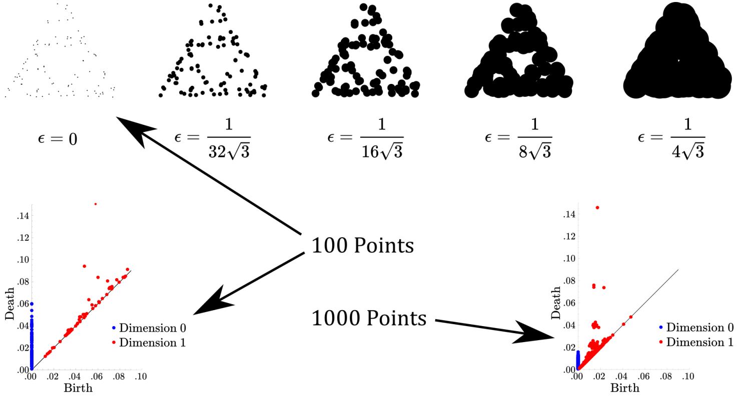

If we take the persistent homology of larger and larger samples, the diagram begins to approach that of the support.

Persistent Homology of a Sample

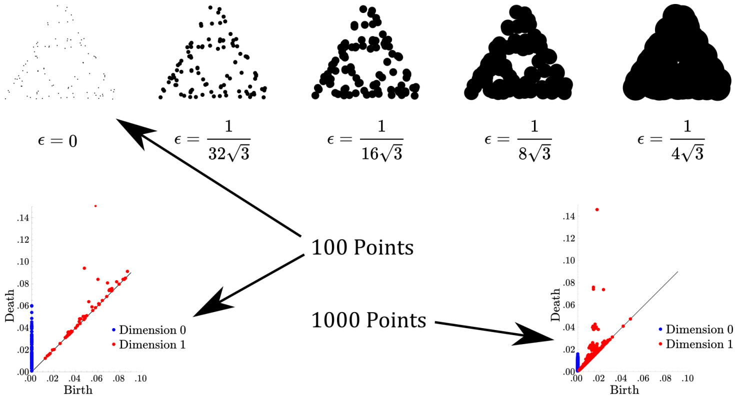

We also have a cluster of small intervals that are usually written off as "noise." We can use this noise to estimate fractal dimension!

Minimum Spanning Trees

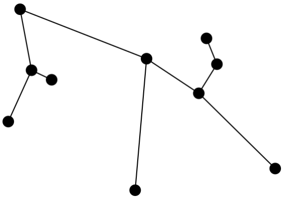

Definition (Minimum Spanning Tree)

Let $\mathbf{x}$ be a finite metric space (e.g. a weighted graph). The **minimum spanning tree** on $\mathbf{x}$, denoted $T(\mathbf{x})$ is the connected graph with vertex set $\mathbf{x}$ that minimizes the sum of the length of the edges.

In fact, for any $\alpha > 0$, $T(\mathbf{x})$ minimizes the weighted sum

$$\Sigma_{e \in T(\mathbf{x})} |e|^{\alpha} .$$

Minimum Spanning Trees

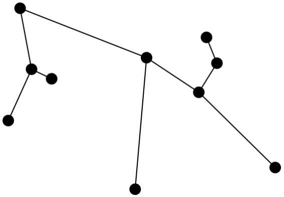

Definition (Minimum Spanning Tree)

Let $\mathbf{x}$ be a finite metric space (e.g. a weighted graph). The **minimum spanning tree** on $\mathbf{x}$, denoted $T(\mathbf{x})$ is the connected graph with vertex set $\mathbf{x}$ that minimizes the sum of the length of the edges.

In fact, for any $\alpha > 0$, $T(\mathbf{x})$ minimizes the weighted sum

$$\Sigma_{e \in T(\mathbf{x})} |e|^{\alpha} .$$

Minimum Spanning Trees and PH_0

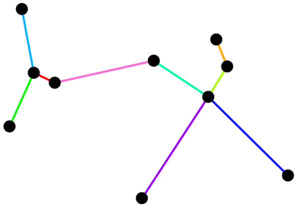

$T(\mathbf{x})$

$PH_0(\mathbf{x})$

Theorem

If $\mathbf{x}$ is a finite metric space, then there is a bijection between the edges of $T(\mathbf{x})$ and the intervals of $PH_0(\mathbf{x})$. An edge corresponds to an interval of half the length.

*This depends on whether persistent homology is taken with respect to the Rips/Čech complex.

Minimum Spanning Trees and PH_0

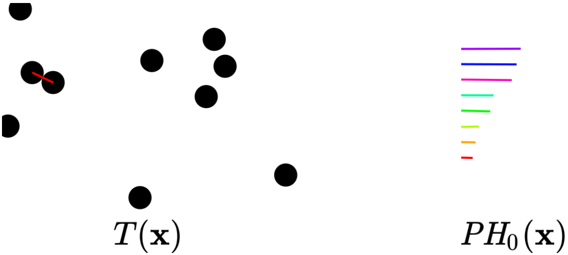

The proof follows Kruskal's algorithm to compute the MST
– Expand balls until two from different components overlap

Minimum Spanning Trees and PH_0

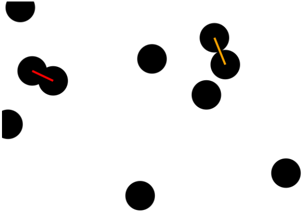

$T(\mathbf{x})$

$PH_0(\mathbf{x})$

The proof follows Kruskal's algorithm to compute the MST
– Expand balls until two from different components overlap

Minimum Spanning Trees and PH_0

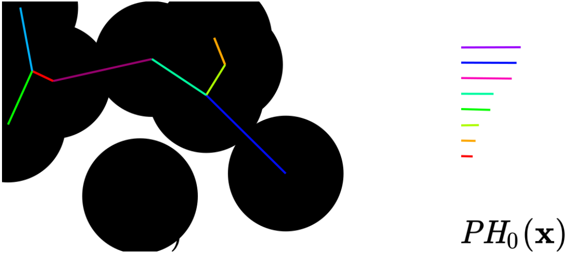

The proof follows Kruskal's algorithm to compute the MST
– Expand balls until two from different components overlap

Minimum Spanning Trees and PH_0

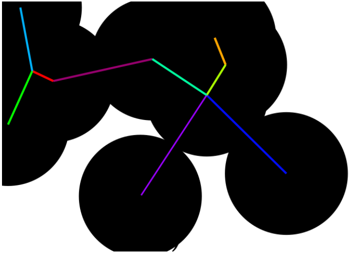

$PH_0(\mathbf{x})$

The proof follows Kruskal's algorithm to compute the MST
– Expand balls until two from different components overlap

Main Questions

Question

Can the fractal dimension of a metric measure space be estimated from the persistent homology of random point samples?

Question

How does the practical performance of the PH_i-dimension compare to classical methods such as box-counting or the correlation algorithm?

Presentation Outline

1. Background: Fractal Dimension and Persistent Homology

2. Previous Work and Definitions

3. Computational Results

Previous Work on Fractal Dimension and *PH*

Several authors have defined fractal dimensions based on *PH*, and compared computational estimates with known dimensions:

2000 Robins. *PhD thesis*: Persistent Betti numbers of fractals, proved results for H_0 of disconnected sets.

2012 MacPherson & Schweinhart. *Measuring Shape with Topology*: *PH* of shapes, studied probability distributions of polymers.

2018 Schweinhart. *PH and the Upper Box Dimension & The persistent homology of random geometric complexes on fractals*: First results relating *PH* to a classically defined fractal dimension.

2019 Adams et. al. *A Fractal Dimension for Measures via PH*: *PH* of random point samples; definition very similar to the one here. Computational experiments that motivated the current work.

Previous Work on Fractal Dimension and *PH*

Several authors have defined fractal dimensions based on *PH*, and compared computational estimates with known dimensions:

2000 Robins. *PhD thesis*: Persistent Betti numbers of fractals, proved results for H_0 of disconnected sets.

2012 MacPherson & Schweinhart. *Measuring Shape with Topology*: *PH* of shapes, studied probability distributions of polymers.

2018 Schweinhart. *PH and the Upper Box Dimension & The persistent homology of random geometric complexes on fractals*: First results relating *PH* to a classically defined fractal dimension.

2019 Adams et. al. *A Fractal Dimension for Measures via PH*: *PH* of random point samples; definition very similar to the one here. Computational experiments that motivated the current work.

Previous Work on Fractal Dimension and *PH*

Several authors have defined fractal dimensions based on *PH*, and compared computational estimates with known dimensions:

2000 Robins. *PhD thesis*: Persistent Betti numbers of fractals, proved results for H_0 of disconnected sets.

2012 MacPherson & Schweinhart. *Measuring Shape with Topology*: *PH* of shapes, studied probability distributions of polymers.

2018 Schweinhart. *PH and the Upper Box Dimension & The persistent homology of random geometric complexes on fractals*: First results relating *PH* to a classically defined fractal dimension.

2019 Adams et. al. *A Fractal Dimension for Measures via PH*: *PH* of random point samples; definition very similar to the one here. Computational experiments that motivated the current work.

Previous Work on Fractal Dimension and *PH*

Several authors have defined fractal dimensions based on *PH*, and compared computational estimates with known dimensions:

2000 Robins. *PhD thesis*: Persistent Betti numbers of fractals, proved results for H_0 of disconnected sets.

2012 MacPherson & Schweinhart. *Measuring Shape with Topology*: *PH* of shapes, studied probability distributions of polymers.

2018 Schweinhart. *PH and the Upper Box Dimension* & *The persistent homology of random geometric complexes on fractals*: First results relating *PH* to a classically defined fractal dimension.

2019 Adams et. al. *A Fractal Dimension for Measures via PH*: *PH* of random point samples; definition very similar to the one here. Computational experiments that motivated the current work.

Previous Work on Fractal Dimension and *PH*

Several authors have defined fractal dimensions based on *PH*, and compared computational estimates with known dimensions:

2000 Robins. *PhD thesis*: Persistent Betti numbers of fractals, proved results for H_0 of disconnected sets.

2012 MacPherson & Schweinhart. *Measuring Shape with Topology*: *PH* of shapes, studied probability distributions of polymers.

2018 Schweinhart. *PH and the Upper Box Dimension* & *The persistent homology of random geometric complexes on fractals*: First results relating *PH* to a classically defined fractal dimension.

2019 Adams et. al. *A Fractal Dimension for Measures via PH*: *PH* of random point samples; definition very similar to the one here. Computational experiments that motivated the current work.

Persistent Homology Dimension

Definition (α-Weighted Lifetime Sum)

If X is a bounded metric space, define

$$E_\alpha^i(X) = \sum_{(b,d)\in PH_i(X)} (d-b)^\alpha .$$

When $i = 0$ and X is finite the sum can be taken over the edges of the minimum spanning tree on X : $E_\alpha^0(X) = \sum_{e\in T(X)} \left(\frac{|e|}{2}\right)^\alpha$

Definition (Persistent Homology Dimension)

Let μ be a probability measure on a metric space, $\{x_i\}_{i\in\mathbb{N}}$ be i.i.d. samples from μ, and $\alpha > 0$

If the support of μ is d-dimensional, then $E_\alpha^i(x_1,\ldots,x_n)$ should scale as $n^{\frac{d-\alpha}{d}}$.

$$\dim_{PH_i^\alpha}(\mu) := \frac{\alpha}{1-\beta}, \qquad \beta := \limsup_{n\to\infty} \frac{\log(\mathbb{E}(E_\alpha^i(x_1,\ldots,x_n)))}{\log(n)} .$$

Persistent Homology Dimension

Definition (α-Weighted Lifetime Sum)

If X is a bounded metric space, define

$$E_\alpha^i(X) = \sum_{(b,d)\in PH_i(X)} (d-b)^\alpha .$$

When $i=0$ and X is finite the sum can be taken over the edges of the minimum spanning tree on X : $E_\alpha^0(X) = \sum_{e\in T(X)} \left(\frac{|e|}{2}\right)^\alpha$

Definition (Persistent Homology Dimension)

Let μ be a probability measure on a metric space, $\{x_i\}_{i\in\mathbb{N}}$ be i.i.d. samples from μ, and $\alpha > 0$

If the support of μ is d-dimensional, then $E_\alpha^i(x_1,\dots,x_n)$ should scale as $n^{\frac{d-\alpha}{d}}$.

$$\dim_{PH_i^\alpha}(\mu) := \frac{\alpha}{1-\beta}, \qquad \beta := \limsup_{n\to\infty} \frac{\log(\mathbb{E}(E_\alpha^i(x_1,\dots,x_n)))}{\log(n)} .$$

Persistent Homology Dimension

Definition (α-Weighted Lifetime Sum)

If X is a bounded metric space, define

$$E_\alpha^i(X) = \sum_{(b,d)\in PH_i(X)} (d-b)^\alpha \, .$$

When $i = 0$ and X is finite the sum can be taken over the edges of the minimum spanning tree on X : $\quad E_\alpha^0(X) = \sum_{e\in T(X)} \left(\frac{|e|}{2}\right)^\alpha$

Definition (Persistent Homology Dimension)

Let μ be a probability measure on a metric space, $\{x_i\}_{i\in\mathbb{N}}$ be i.i.d. samples from μ, and $\alpha > 0$

If the support of μ is d-dimensional, then $E_\alpha^i(x_1,\ldots,x_n)$ should scale as $n^{\frac{d-\alpha}{d}}$.

$$\dim_{PH_i^\alpha}(\mu) := \frac{\alpha}{1-\beta}, \qquad \beta := \limsup_{n\to\infty} \frac{\log(\mathbb{E}(E_\alpha^i(x_1,\ldots,x_n)))}{\log(n)} \, .$$

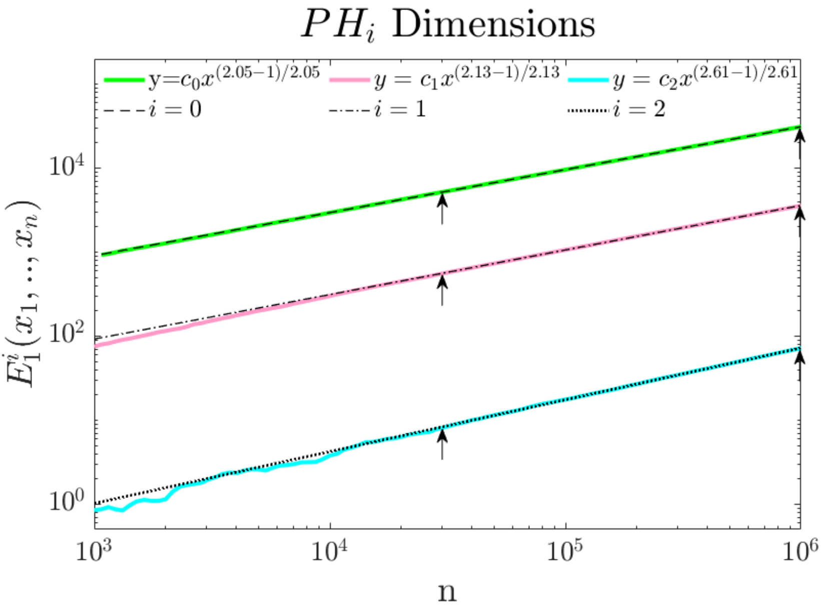
PH_i Dimensions
$y = c_0 x^{(2.05-1)/2.05}$
$y = c_1 x^{(2.13-1)/2.13}$
$y = c_2 x^{(2.61-1)/2.61}$
$i = 0$
$i = 1$
$i = 2$
10^4
10^2
10^0
$E_1^i(x_1, \ldots, x_n)$
10^3
10^4
10^5
10^6
n

Ahlfors Regularity

Ahlfors regularity is a standard hypothesis that implies that the fractal dimension of a measure is well defined. That is, the various classical notions of dimension coincide and equal d.

Definition (Ahlfors Regularity)

A probability measure μ supported on a metric space X is d-Ahlfors regular if there exist positive real numbers c and r_0 so that

$$\frac{1}{c}\, r^d \leq \mu(B_r(x)) \leq c\, r^d$$

r for all $x \in X$ and $r < r_0$.

Examples include:

- Bounded probability densities
- The natural measures on the Cantor set, Sierpinski triangle

Equivalence of Persistent Homology Dimension

Theorem (Schweinhart, 2018)

Let μ be a d-Ahlfors regular measure on a metric space.
If $0 < \alpha < d$ then,

$$dim_{PH_0^\alpha}(\mu) = d\,.$$

Higher dimensional results are more difficult.
Cleanest result is for $\mathbb{R}^2$ (for the Čech complex):

Theorem (Schweinhart, 2018)

Let μ be a d-Ahlfors regular measure on $\mathbb{R}^2$ with $d > 1.5$.
If $0 < \alpha < d$, then,

$$dim_{PH_1^\alpha}(\mu) = d\,.$$

Equivalence of Persistent Homology Dimension

Theorem (Schweinhart, 2018)

Let μ be a d-Ahlfors regular measure on a metric space.
If $0 < \alpha < d$ then,

$$dim_{PH_0^\alpha}(\mu) = d\,.$$

Higher dimensional results are more difficult.
Cleanest result is for $\mathbb{R}^2$ (for the Čech complex):

Theorem (Schweinhart, 2018)

Let μ be a d-Ahlfors regular measure on $\mathbb{R}^2$ with $d > 1.5$.
If $0 < \alpha < d$, then,

$$dim_{PH_1^\alpha}(\mu) = d\,.$$

Presentation Outline

Computational Results

Question

How does the practical performance of the PH_i-dimension compare to classical methods?

Computational Results

Question

How does the practical performance of the PH_i-dimension compare to classical methods?

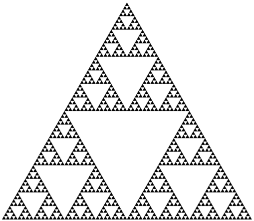

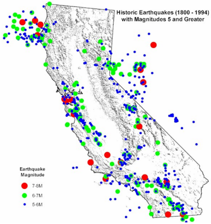

We compared the performance of algorithms to estimate the PH_i, box-counting, and correlation dimensions, for three classes of examples: self-similar fractals, chaotic attractors, and empirical earthquake data.

Classical Fractal Dimension Definitions

Definition (Correlation Dimension)

A probability measure μ on a space X induces a probability measure ν on the distance set of X. Define the correlation integral and correlation dimension of X

$$C(\epsilon) = \mathbb{P}(d(x,y) < \epsilon), \qquad \dim_{\text{corr}}(\mu) = \lim_{\epsilon \to 0} \frac{\log(C(\epsilon))}{\log(\epsilon)}$$

Definition (Box Counting Dimension)

Fix compact $X \subseteq \mathbb{R}^m$. Let $\left\{C_i^\delta\right\}_{i \in \mathbb{N}}$ be the cubes in the standard tiling of $\mathbb{R}^m$ by cubes of width δ. Let $N_\delta(X)$ be the number of cubes in $\left\{C_i^\delta\right\}_{i \in \mathbb{N}}$ that intersect X. Define the upper and lower box-counting dimensions by

$$\dim_{\overline{\text{box}}}(X) = \limsup_{\delta \to 0} \frac{\log(N_\delta(X))}{\log(1/\delta)} \quad \text{and} \quad \dim_{\underline{\text{box}}}(X) = \liminf_{\delta \to 0} \frac{\log(N_\delta(X))}{\log(1/\delta)},$$

Classical Fractal Dimension Definitions

Definition (Correlation Dimension)

A probability measure μ on a space X induces a probability measure ν on the distance set of X. Define the correlation integral and correlation dimension of X

$$C(\epsilon) = \mathbb{P}(d(x,y) < \epsilon), \qquad \dim_{\text{corr}}(\mu) = \lim_{\epsilon \to 0} \frac{\log(C(\epsilon))}{\log(\epsilon)}$$

Definition (Box Counting Dimension)

Fix compact $X \subseteq \mathbb{R}^m$. Let $\left\{C_i^\delta\right\}_{i \in \mathbb{N}}$ be the cubes in the standard tiling of $\mathbb{R}^m$ by cubes of width δ. Let $N_\delta(X)$ be the number of cubes in $\left\{C_i^\delta\right\}_{i \in \mathbb{N}}$ that intersect X. Define the upper and lower box-counting dimensions by

$$\dim_{\overline{\text{box}}}(X) = \limsup_{\delta \to 0} \frac{\log(N_\delta(X))}{\log(1/\delta)} \qquad \text{and} \qquad \dim_{\underline{\text{box}}}(X) = \liminf_{\delta \to 0} \frac{\log(N_\delta(X))}{\log(1/\delta)},$$

Opportunities for Error

Quote (Brandstater & Swinney, 1987)

It is not difficult to develop an algorithm that will yield numbers that can be called dimension, but it is far more difficult to be confident that those numbers truly represent the dynamics of the system.

r

- Finite sampling, order of limits
- Statistical Error
- Noise
- Discretization
- Edge effects
- Lacunarity and oscillations
- Non-uniform sampling
- Autocorrelation

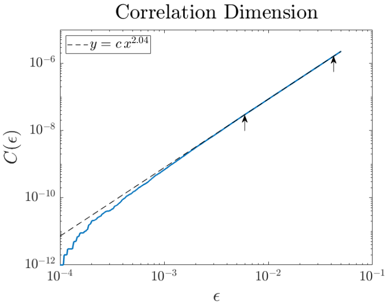

Lorenz Attractor

Opportunities for Error

Quote (Brandstater & Swinney, 1987)

It is not difficult to develop an algorithm that will yield numbers that can be called dimension, but it is far more difficult to be confident that those numbers truly represent the dynamics of the system.

r

- Finite sampling, order of limits
- Statistical Error
- Noise
- Discretization
- Edge effects
- Lacunarity and oscillations
- Non-uniform sampling
- Autocorrelation

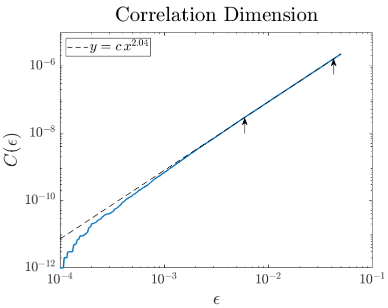

Lorenz Attractor

Lorenz Attractor

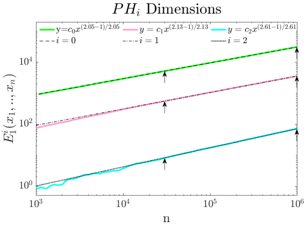

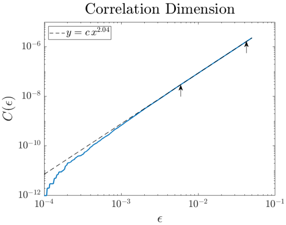

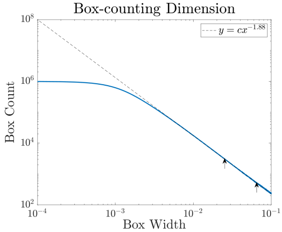

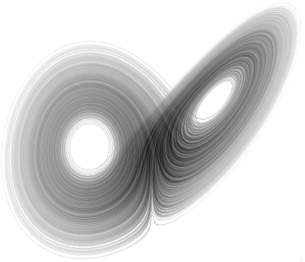

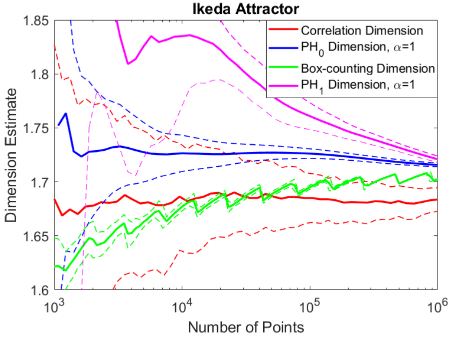

Fractal dimension definitions may disagree,
... or we haven't taken n large enough,
... or we are making a systematic error.

Ikeda attractor

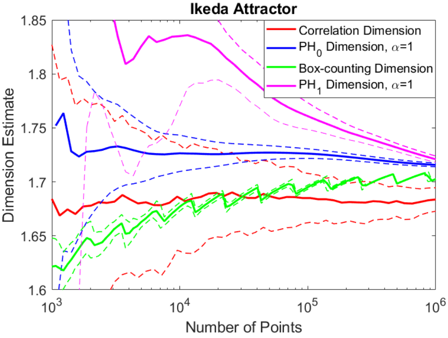

Fractal dimension definitions may disagree,

... or we haven't taken n large enough,

... or we are making a systematic error.

Ikeda attractor

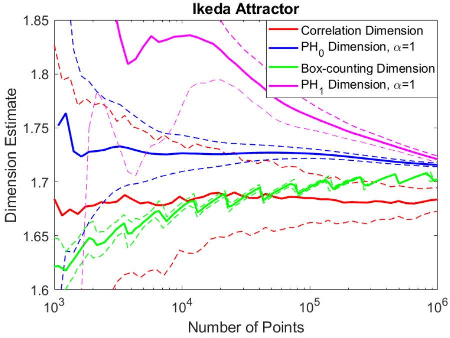

Fractal dimension definitions may disagree,

… or we haven't taken n large enough,

… or we are making a systematic error.

Ikeda attractor

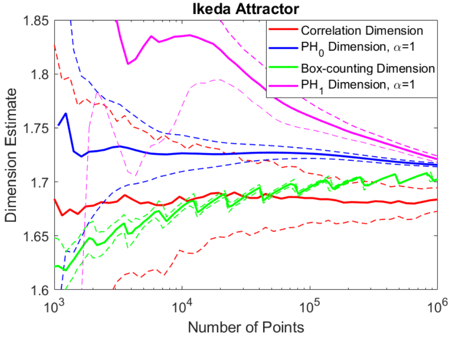

Fractal dimension definitions may disagree,

... or we haven't taken *n* large enough,

... or we are making a systematic error.

Sierpinski Triangle

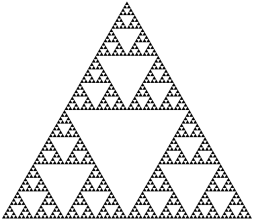

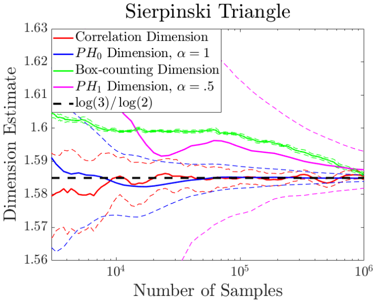

In general, the PH_0 and correlation dimensions perform comparably well. In cases where the true dimension is known, they approach it at about the same rate. In most cases, the box-counting and higher PH_i dimensions perform worse.

Sierpinski Triangle

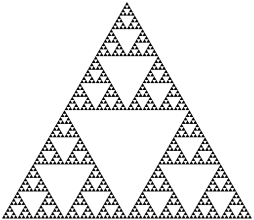

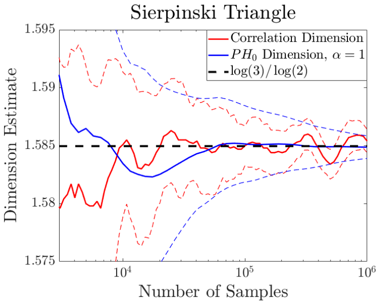

We found one simple rule for fitting a power law to estimate the PH_0 which worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension.

Earthquake Data

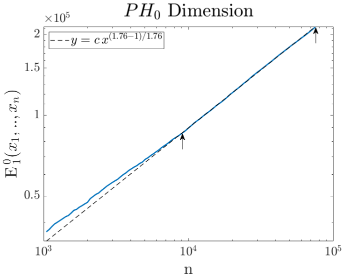

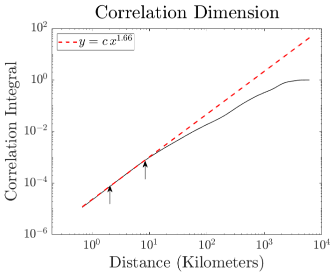

We applied the dimension estimation algorithms to the Hauksson–Shearer Southern California earthquake catalog, and found a PH_0 dimension estimate of 1.76 and a correlation dimension estimate of 1.66. This is in line with previous studies.

Dependence on α: Best Convergence

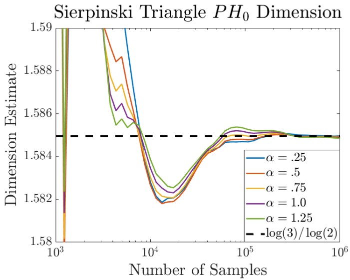

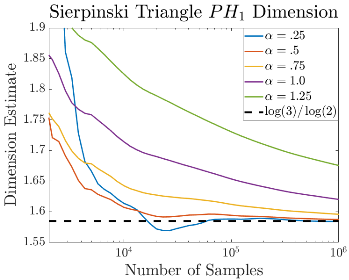

For d-Ahlfors regular measures μ, then $\dim_{PH_0^\alpha}(\mu) = d$ for all α

We choose an α which gives the best convergence

Dependence on α: Multi-Fractals

Different notions of dimension may disagree for non-regular sets.
For non-regular sets, different values of α may give different values for $\dim_{PH}^{i}$.

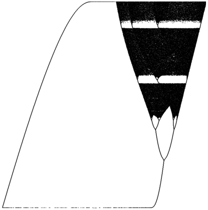

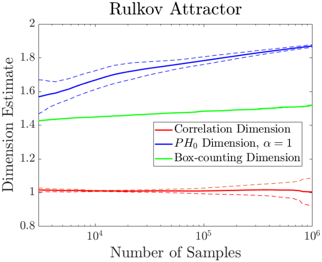

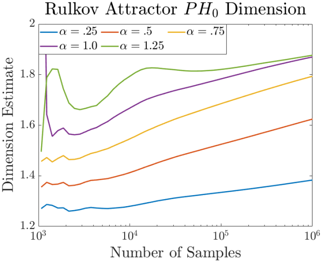

	Correlation	Box-counting	$PH_0^{.5}$	PH_0^{1}	$PH_1^{.5}$	PH_1^{1}	Lyapunov
Rulkov	1.01	1.52	1.62	1.87	2.02	< 2.13	1.19

Slower convergence for higher dimensional data

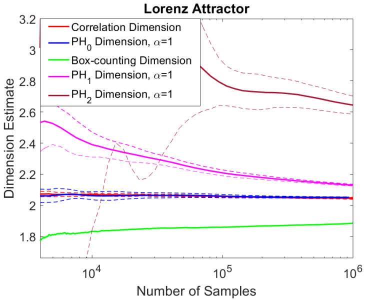

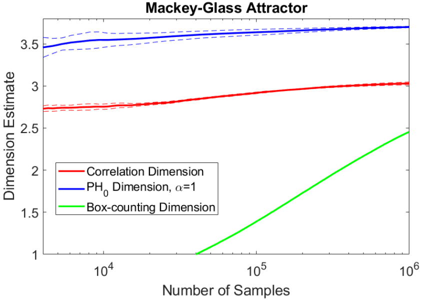

	Correlation	Box-counting	$PH_0^{.5}$	PH_0^1	PH_1^1	Lyapunov
Lorenz	2.04	> 1.90	2.06	2.05	< 2.12	2.06
MG	3.04	> 2.45	3.59	3.70	–	3.58

- **Effectiveness**.
  - PH_0 and correlation dimensions perform comparably well
  - Box-counting, PH_1, and PH_2 dimensions perform worse
- **Efficiency**.
  - Computing PH_0 dimension is fast and comparable with the correlation and box-counting dimensions
  - PH_1 dimension is reasonably fast for subsets of $\mathbb{R}^2$
  - PH_1 and PH_2 are slow/impractical for higher ambient dimensions
- **Equivalence**.
  - For a large class of regular fractals PH dimension coincides with classical definitions of fractal dimension
  - If the different fractal dimension estimates disagree, it may be due to (i) slow convergence (ii) systematic error (iii) definitions truly disagree
- **Error**. Error estimates do not meaningfully reflect the difference between the dimension estimate and the true dimension
- **Ease-of-use**. One simple rule for fitting a power law to estimate the PH_0 worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension

- **Effectiveness**.
  - PH_0 and correlation dimensions perform comparably well
  - Box-counting, PH_1, and PH_2 dimensions perform worse
- **Efficiency**.
  - Computing PH_0 dimension is fast and comparable with the correlation and box-counting dimensions
  - PH_1 dimension is reasonably fast for subsets of $\mathbb{R}^2$
  - PH_1 and PH_2 are slow/impractical for higher ambient dimensions
- **Equivalence**.
  - For a large class of regular fractals PH dimension coincides with classical definitions of fractal dimension
  - If the different fractal dimension estimates disagree, it may be due to (i) slow convergence (ii) systematic error (iii) definitions truly disagree
- **Error**. Error estimates do not meaningfully reflect the difference between the dimension estimate and the true dimension
- **Ease-of-use**. One simple rule for fitting a power law to estimate the PH_0 worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension

- **Effectiveness**.
  - PH_0 and correlation dimensions perform comparably well
  - Box-counting, PH_1, and PH_2 dimensions perform worse
- **Efficiency**.
  - Computing PH_0 dimension is fast and comparable with the correlation and box-counting dimensions
  - PH_1 dimension is reasonably fast for subsets of $\mathbb{R}^2$
  - PH_1 and PH_2 are slow/impractical for higher ambient dimensions
- **Equivalence**.
  - For a large class of regular fractals PH dimension coincides with classical definitions of fractal dimension
  - If the different fractal dimension estimates disagree, it may be due to (i) slow convergence (ii) systematic error (iii) definitions truly disagree
- **Error**. Error estimates do not meaningfully reflect the difference between the dimension estimate and the true dimension
- **Ease-of-use**. One simple rule for fitting a power law to estimate the PH_0 worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension

- **Effectiveness**.
  - PH_0 and correlation dimensions perform comparably well
  - Box-counting, PH_1, and PH_2 dimensions perform worse
- **Efficiency**.
  - Computing PH_0 dimension is fast and comparable with the correlation and box-counting dimensions
  - PH_1 dimension is reasonably fast for subsets of $\mathbb{R}^2$
  - PH_1 and PH_2 are slow/impractical for higher ambient dimensions
- **Equivalence**.
  - For a large class of regular fractals PH dimension coincides with classical definitions of fractal dimension
  - If the different fractal dimension estimates disagree, it may be due to (i) slow convergence (ii) systematic error (iii) definitions truly disagree
- **Error**. Error estimates do not meaningfully reflect the difference between the dimension estimate and the true dimension
- **Ease-of-use**. One simple rule for fitting a power law to estimate the PH_0 worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension

- **Effectiveness**.
  - PH_0 and correlation dimensions perform comparably well
  - Box-counting, PH_1, and PH_2 dimensions perform worse
- **Efficiency**.
  - Computing PH_0 dimension is fast and comparable with the correlation and box-counting dimensions
  - PH_1 dimension is reasonably fast for subsets of $\mathbb{R}^2$
  - PH_1 and PH_2 are slow/impractical for higher ambient dimensions
- **Equivalence**.
  - For a large class of regular fractals PH dimension coincides with classical definitions of fractal dimension
  - If the different fractal dimension estimates disagree, it may be due to (i) slow convergence (ii) systematic error (iii) definitions truly disagree
- **Error**. Error estimates do not meaningfully reflect the difference between the dimension estimate and the true dimension
- **Ease-of-use**. One simple rule for fitting a power law to estimate the PH_0 worked well for all examples, in contrast to the correlation dimension and (especially) the box-counting dimension