Linear-Size Approximations to the Vietoris-Rips Filtration

Don Sheehy
Geometrica Group
INRIA Saclay

This work appeared at SoCG 2012

The goal of topological data analysis is to extract meaningful topological information from data.

Use powerful ideas from computational geometry to speed up persistent homology computation when the data is intrinsically low-dimensional.

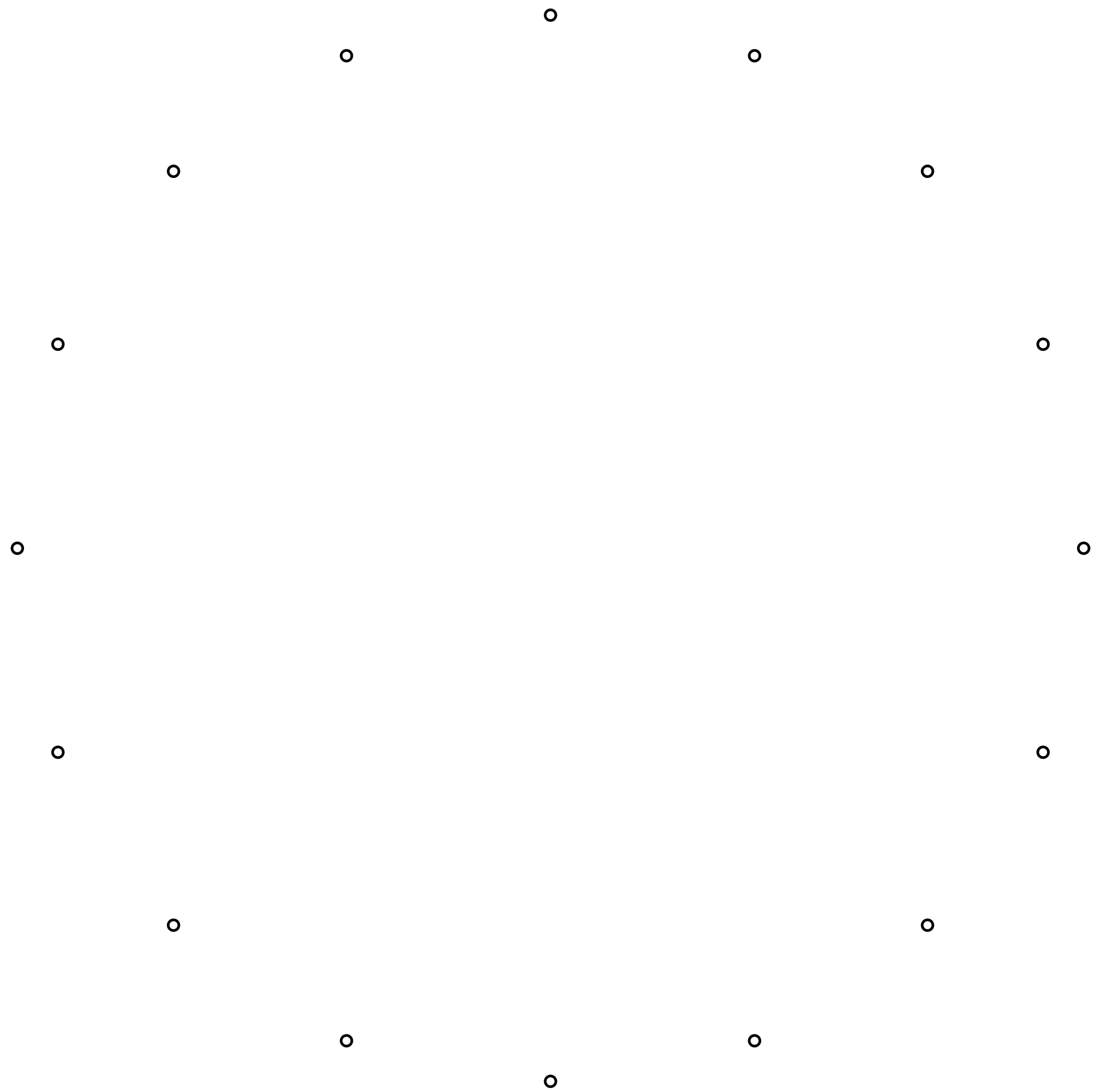

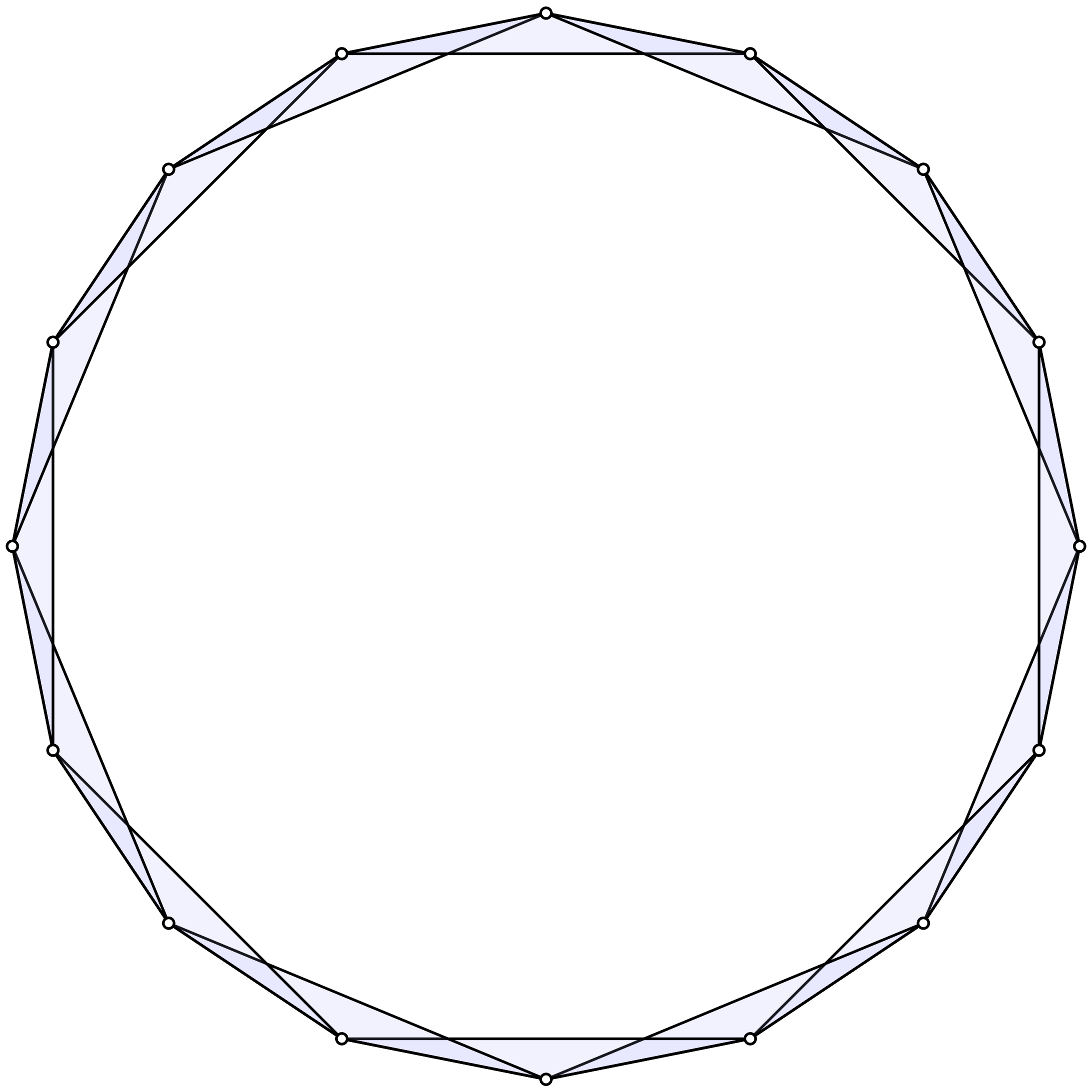

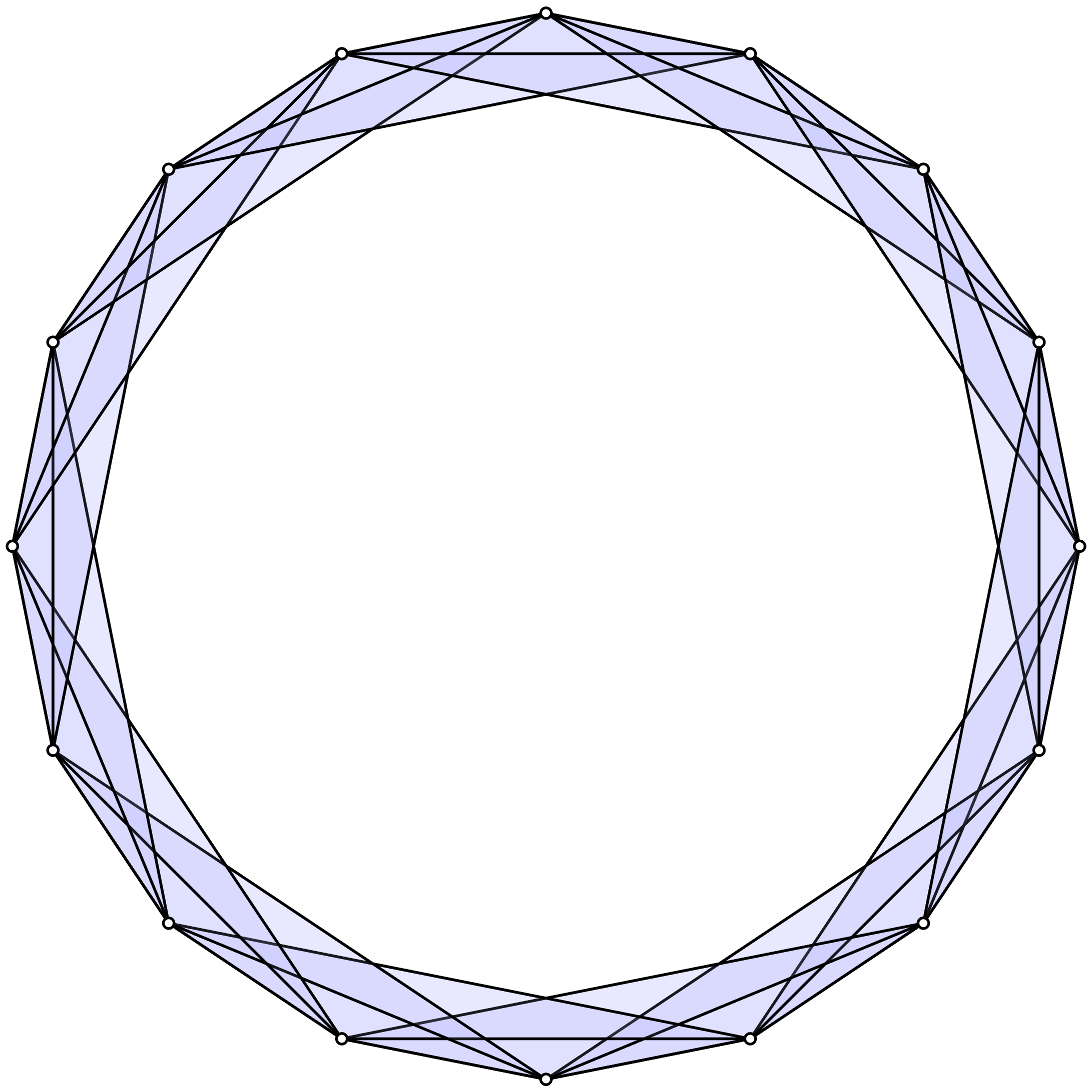

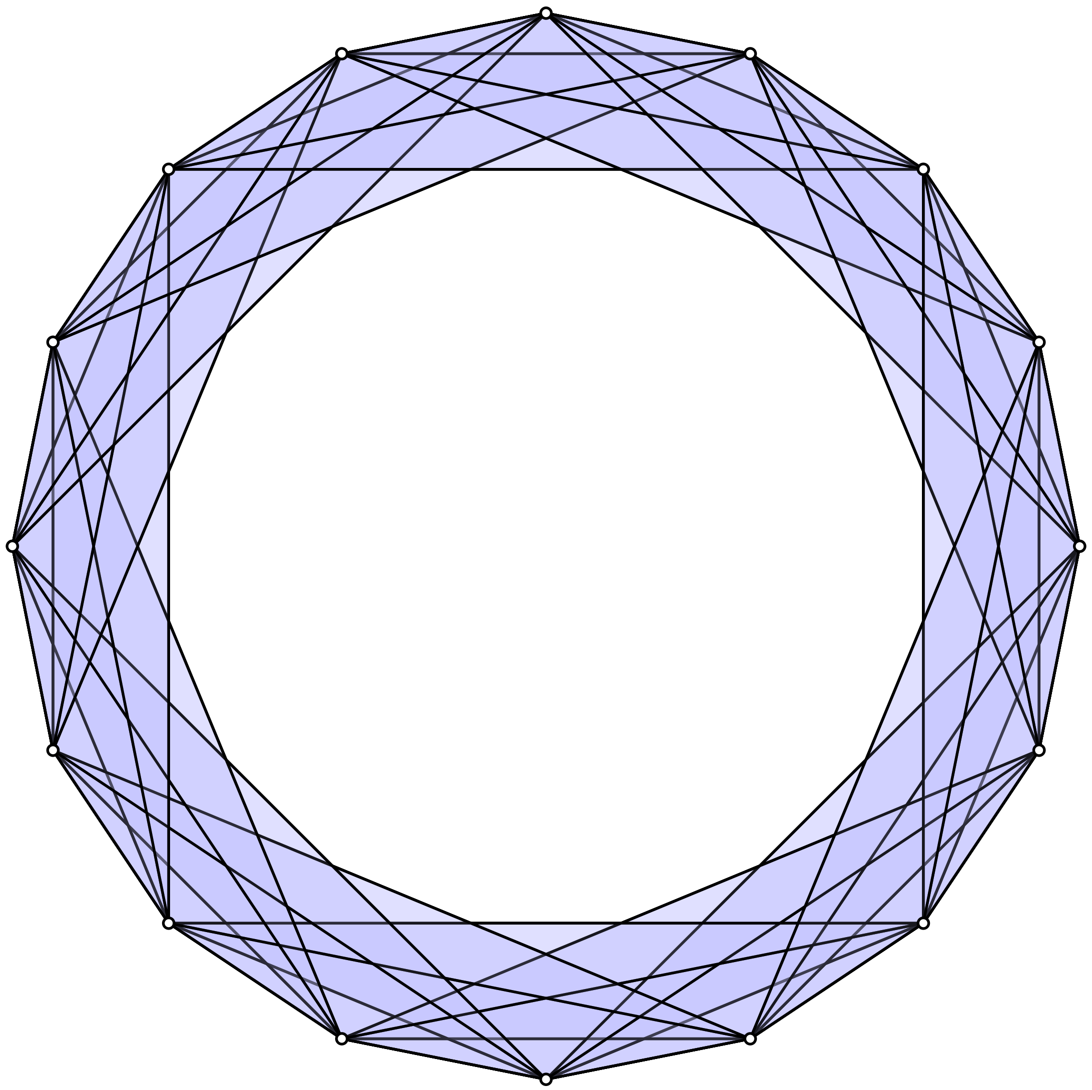

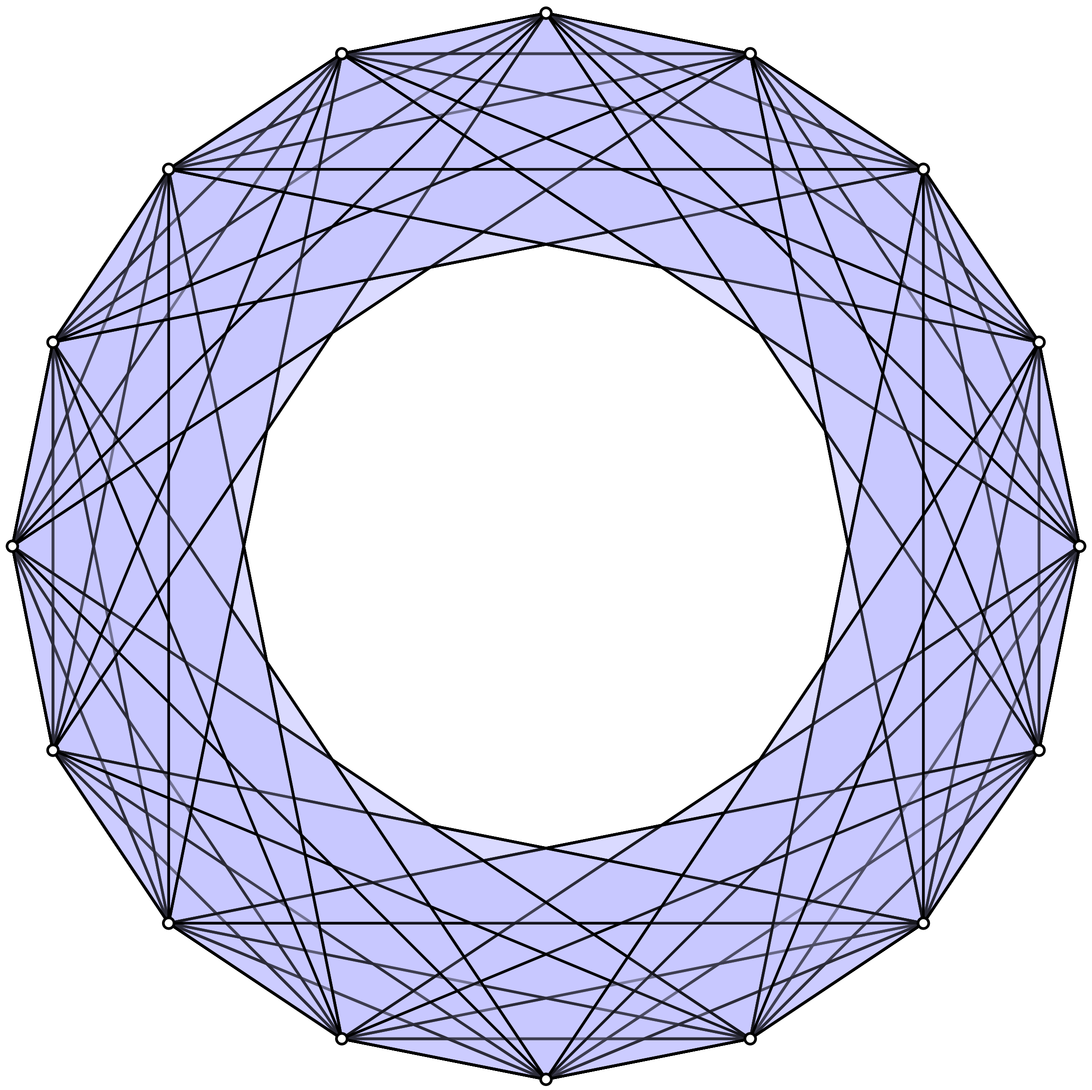

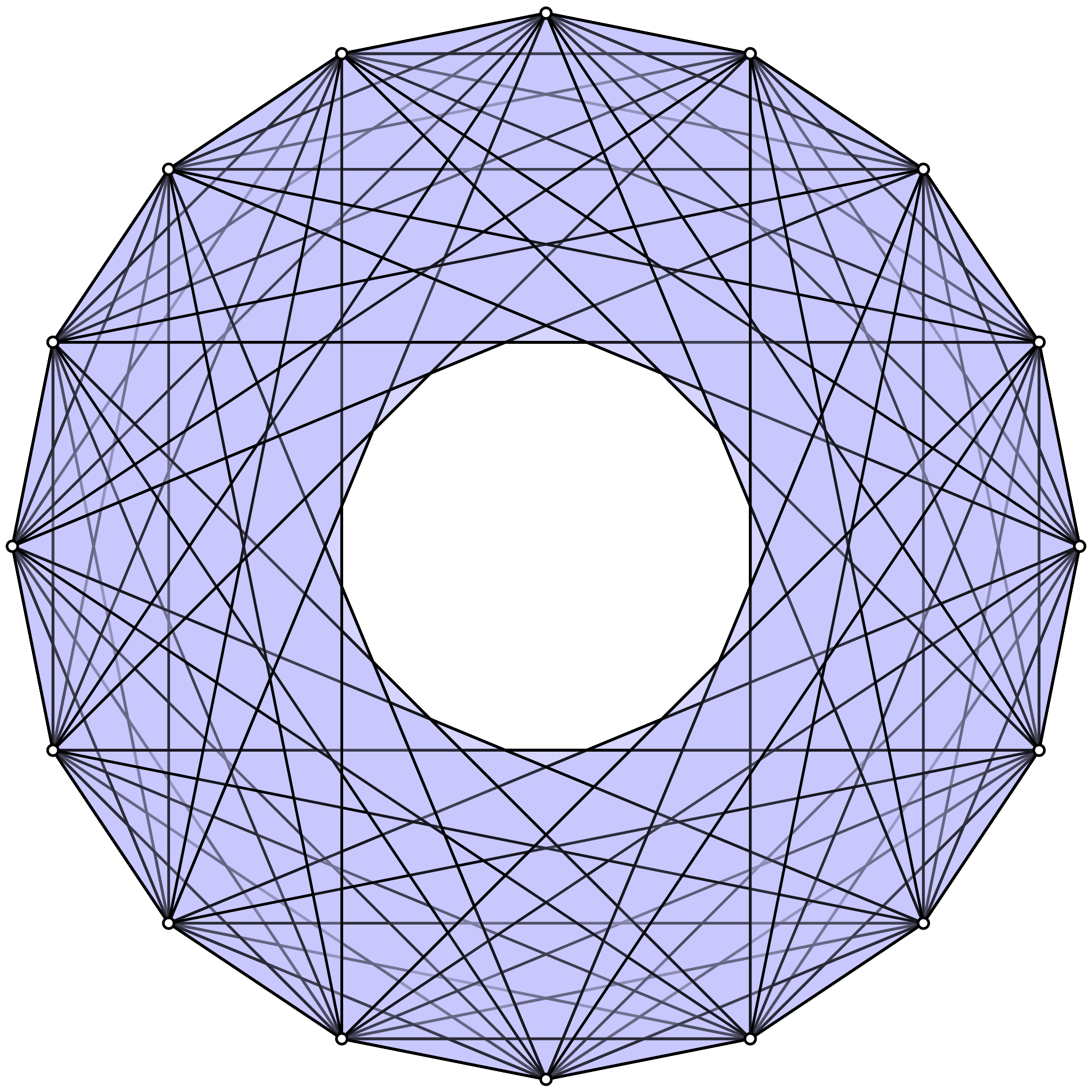

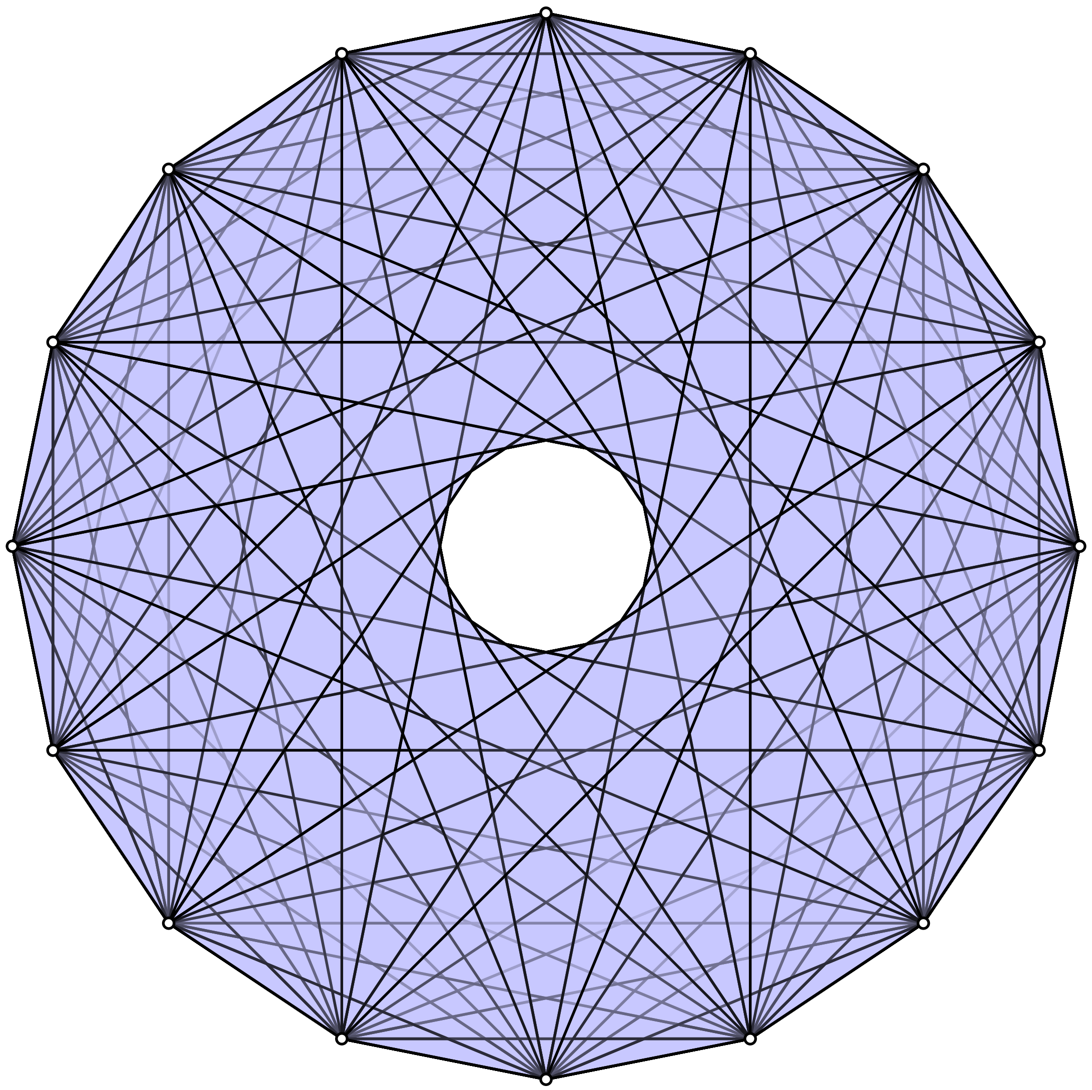

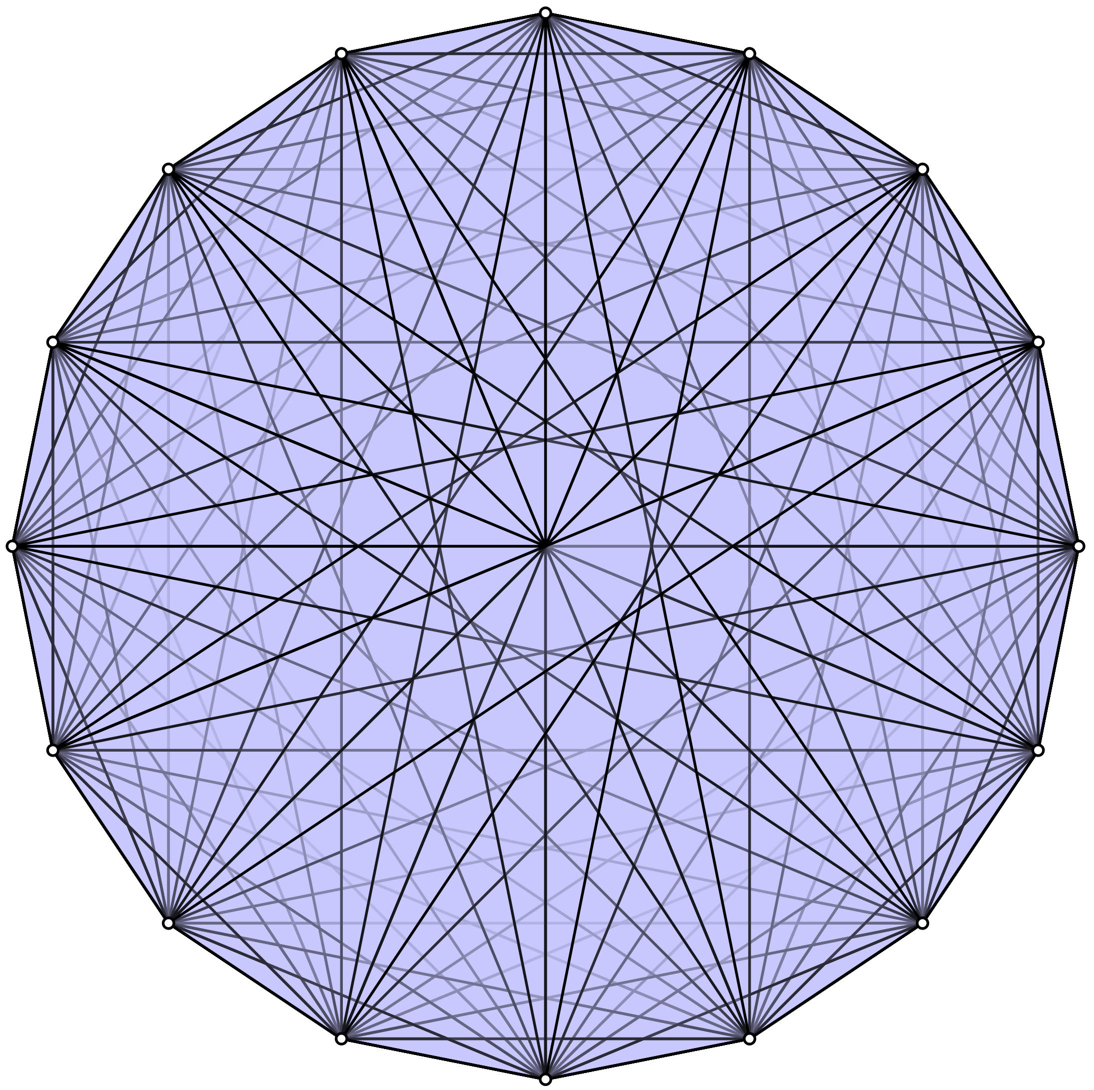

A **filtration** is a growing sequence of spaces.

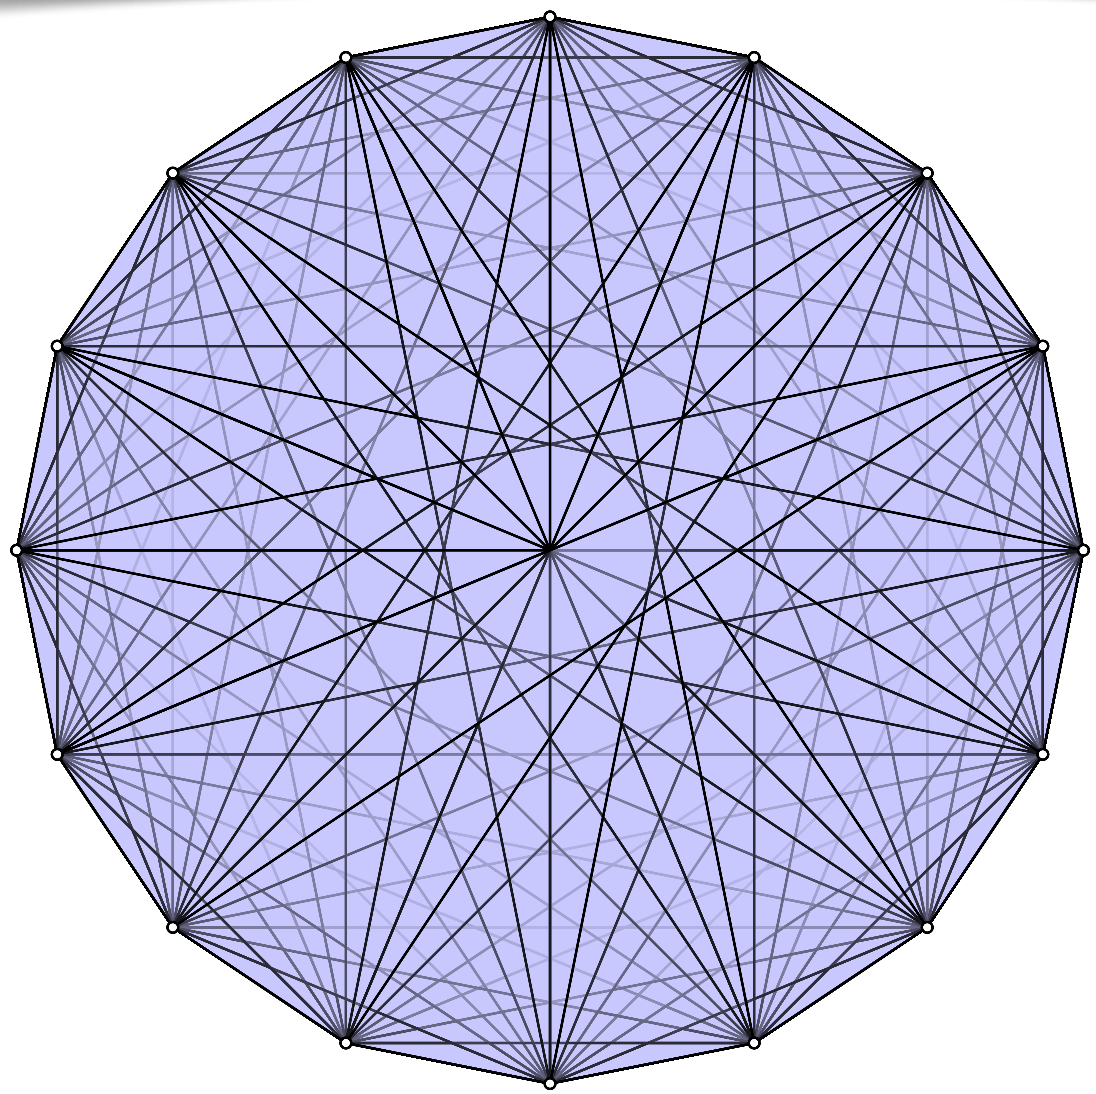

A **filtration** is a growing sequence of spaces.

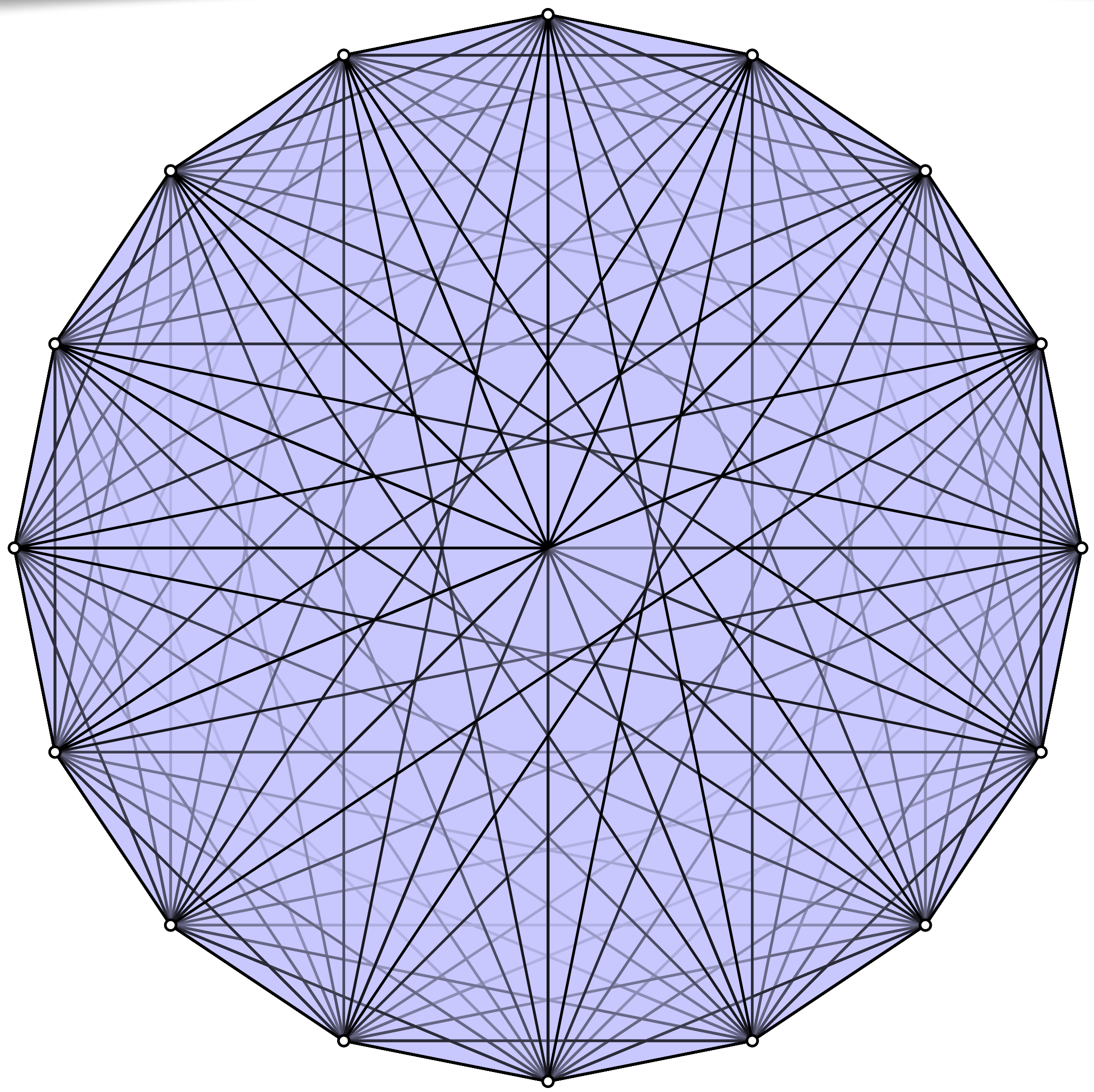

A **persistence diagram** describes the topological changes over time.

A **filtration** is a growing sequence of spaces.

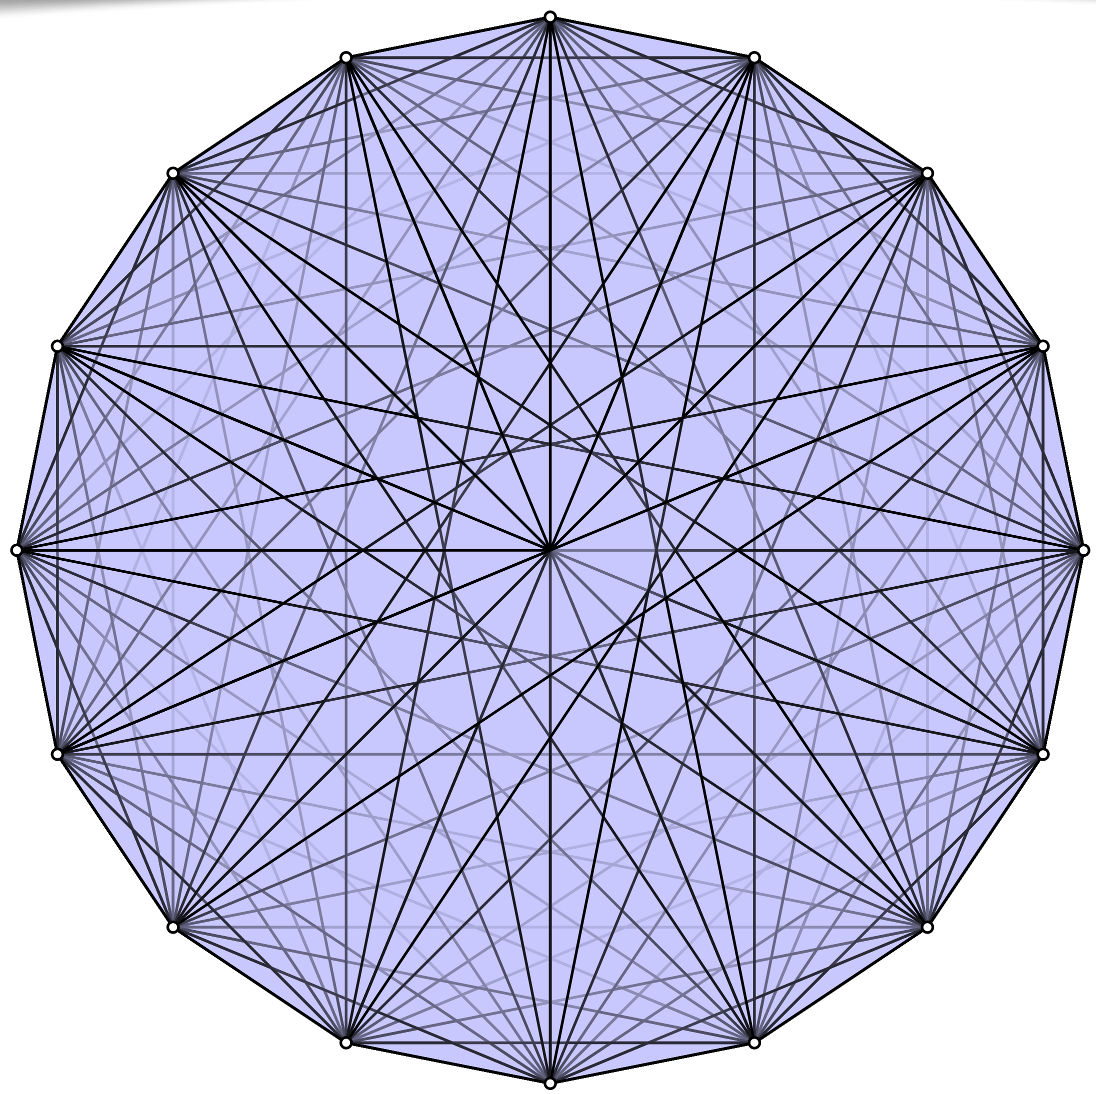

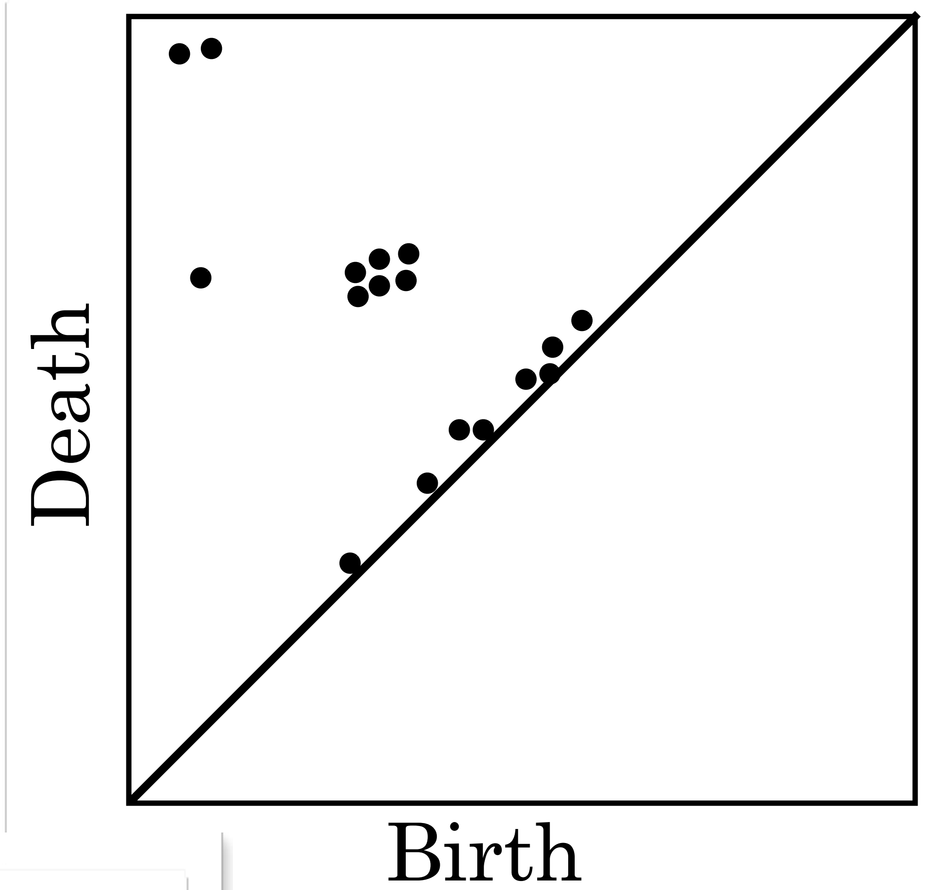

A **persistence diagram** describes the topological changes over time.

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p, q) \leq 2\alpha$ for all $p, q \in \sigma$.

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p, q) \leq 2\alpha$ for all $p, q \in \sigma$.

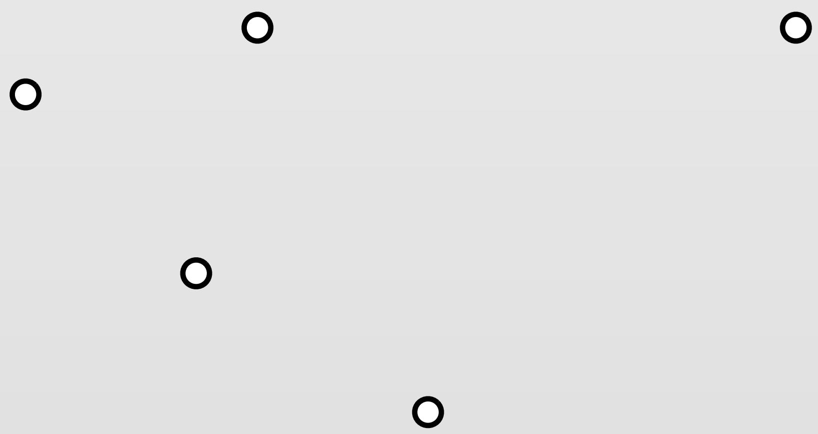

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$.

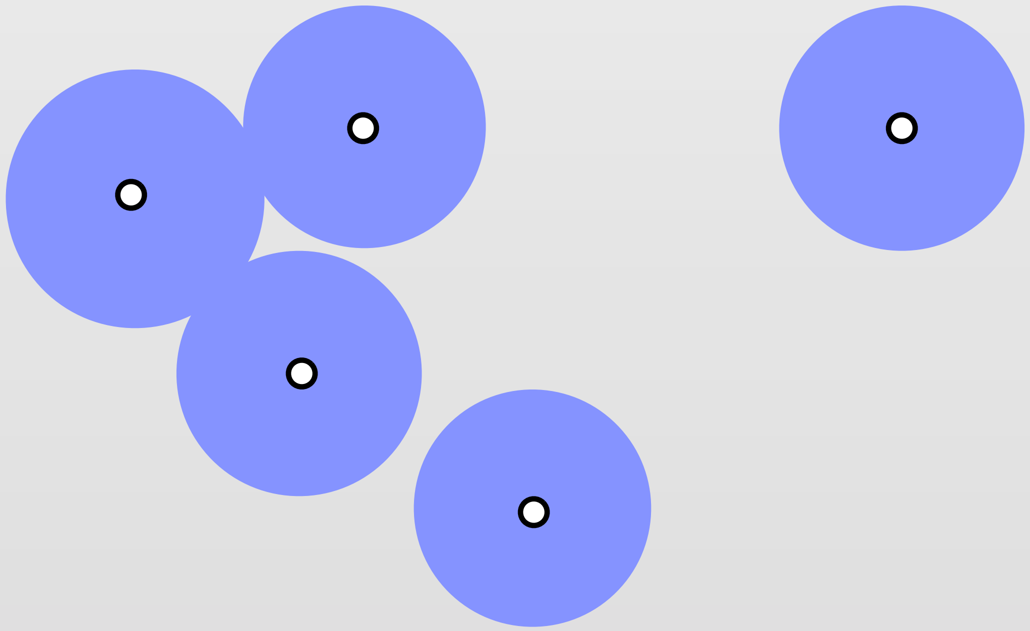

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$.

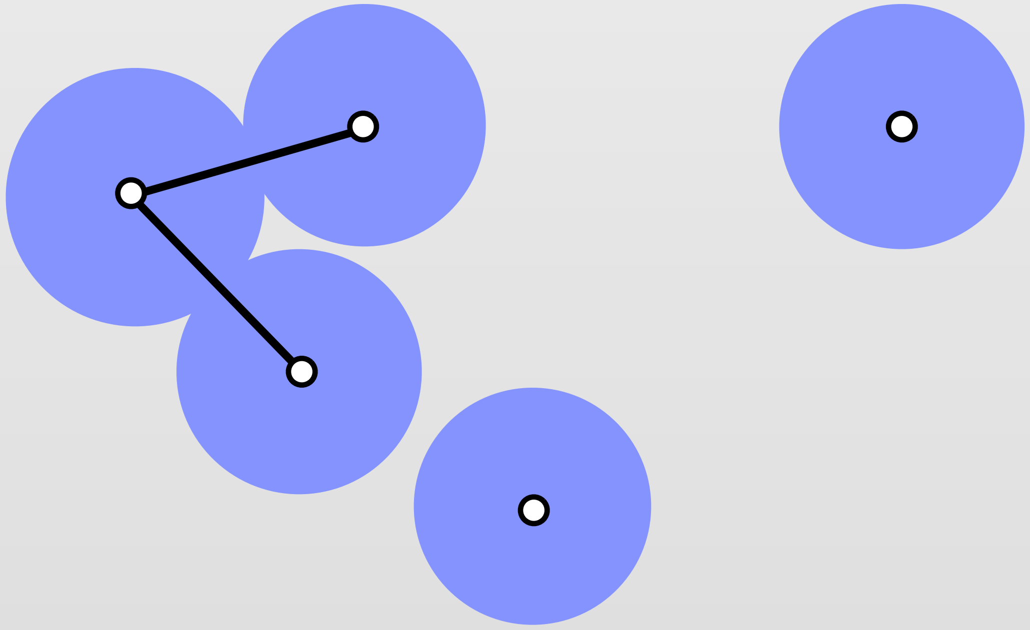

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$.

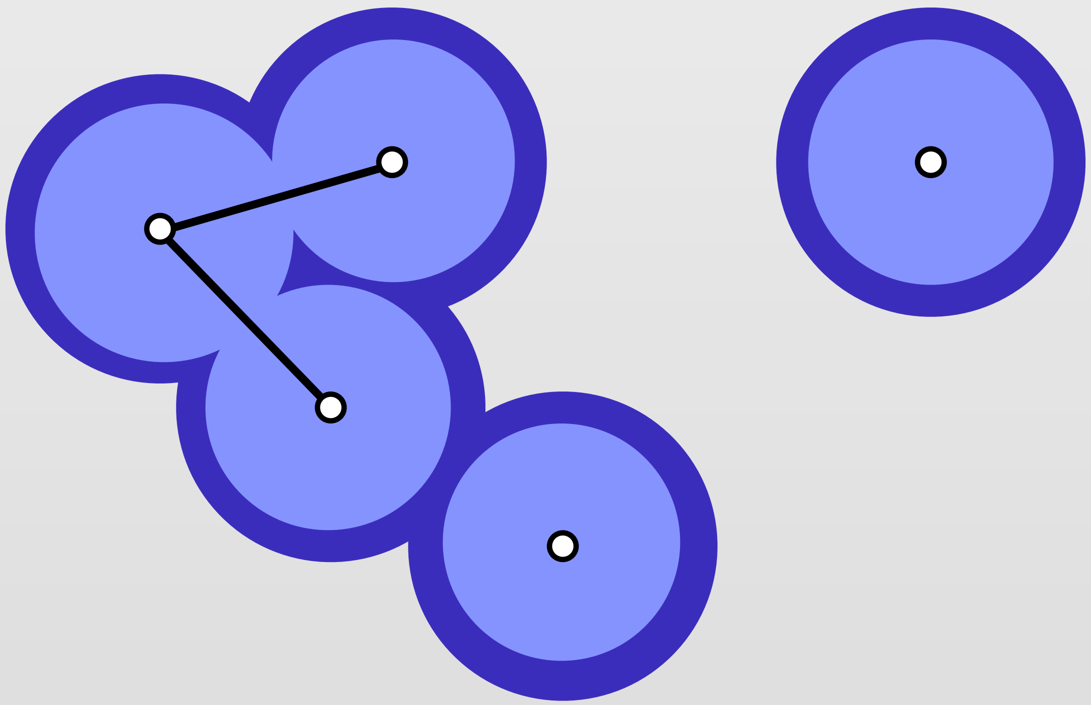

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$.

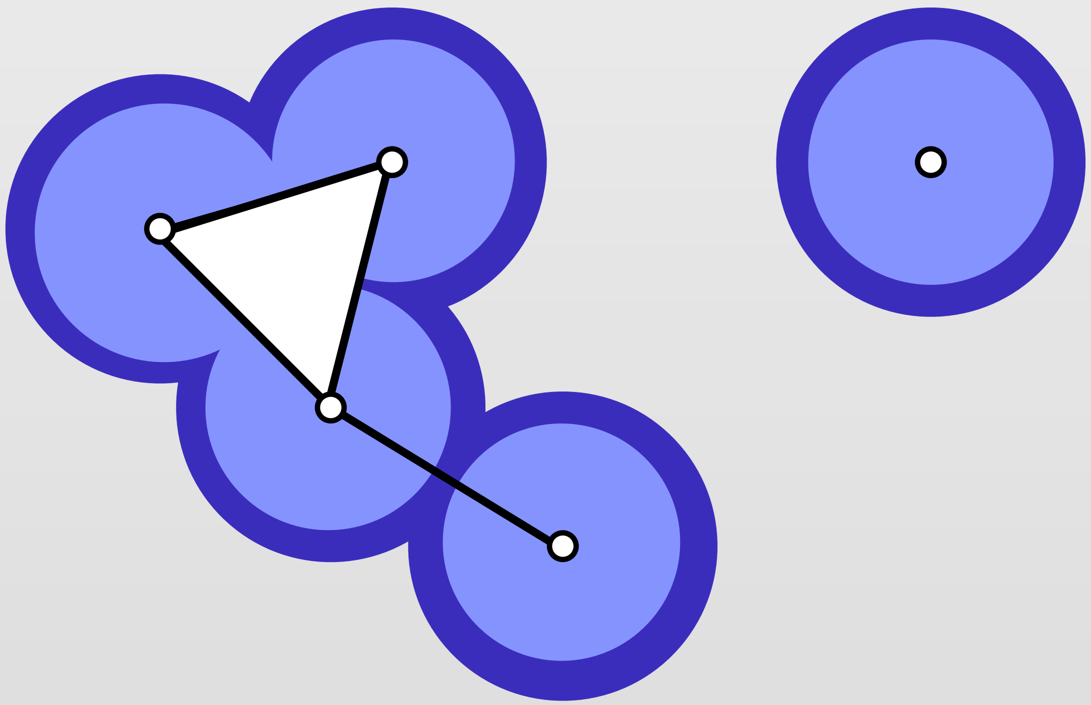

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$.

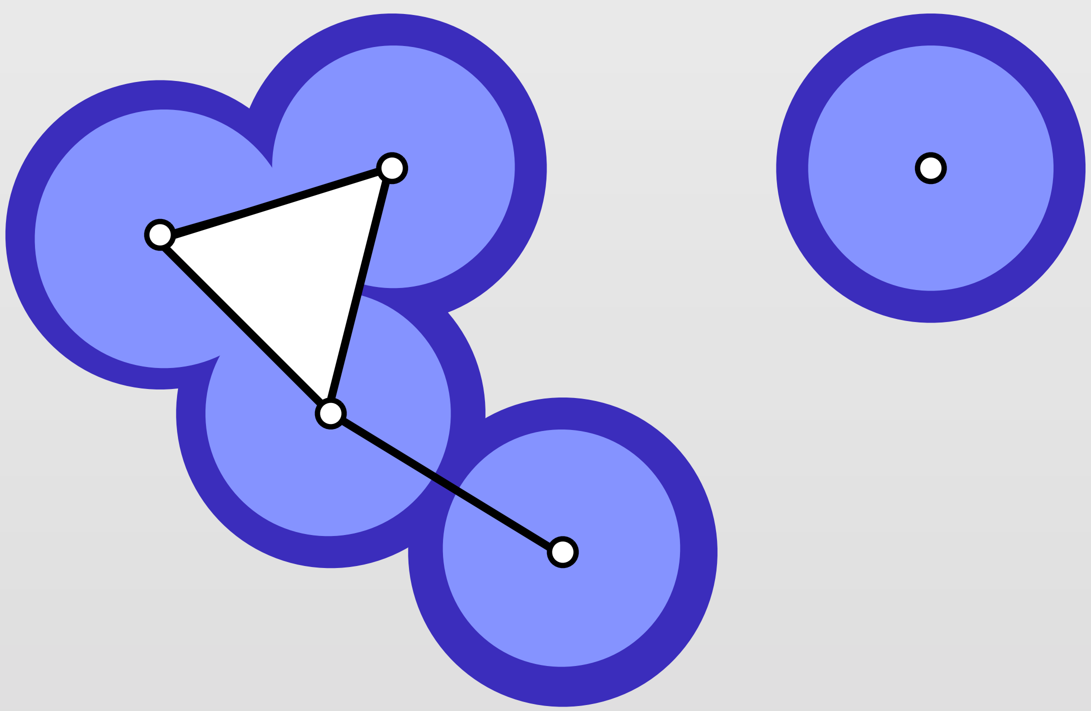

R_∞ is the powerset 2^P.

The **Vietoris-Rips Filtration** encodes the topology of a metric space when viewed at different scales.

Input: A finite metric space $(P, \mathbf{d})$.
Output: A sequence of simplicial complexes $\{R_\alpha\}$ such that $\sigma \in R_\alpha$ iff $\mathbf{d}(p, q) \leq 2\alpha$ for all $p, q \in \sigma$.

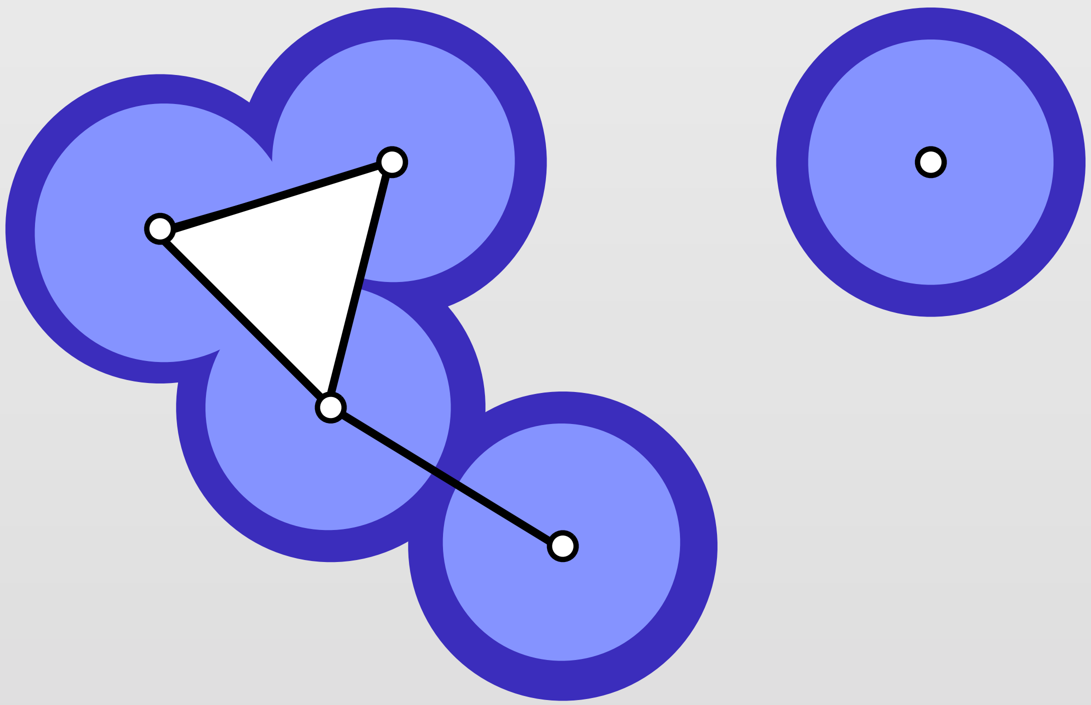

R_∞ is the powerset 2^P.

This is too big!

Persistence Diagrams describe the changes in topology corresponding to changes in scale.

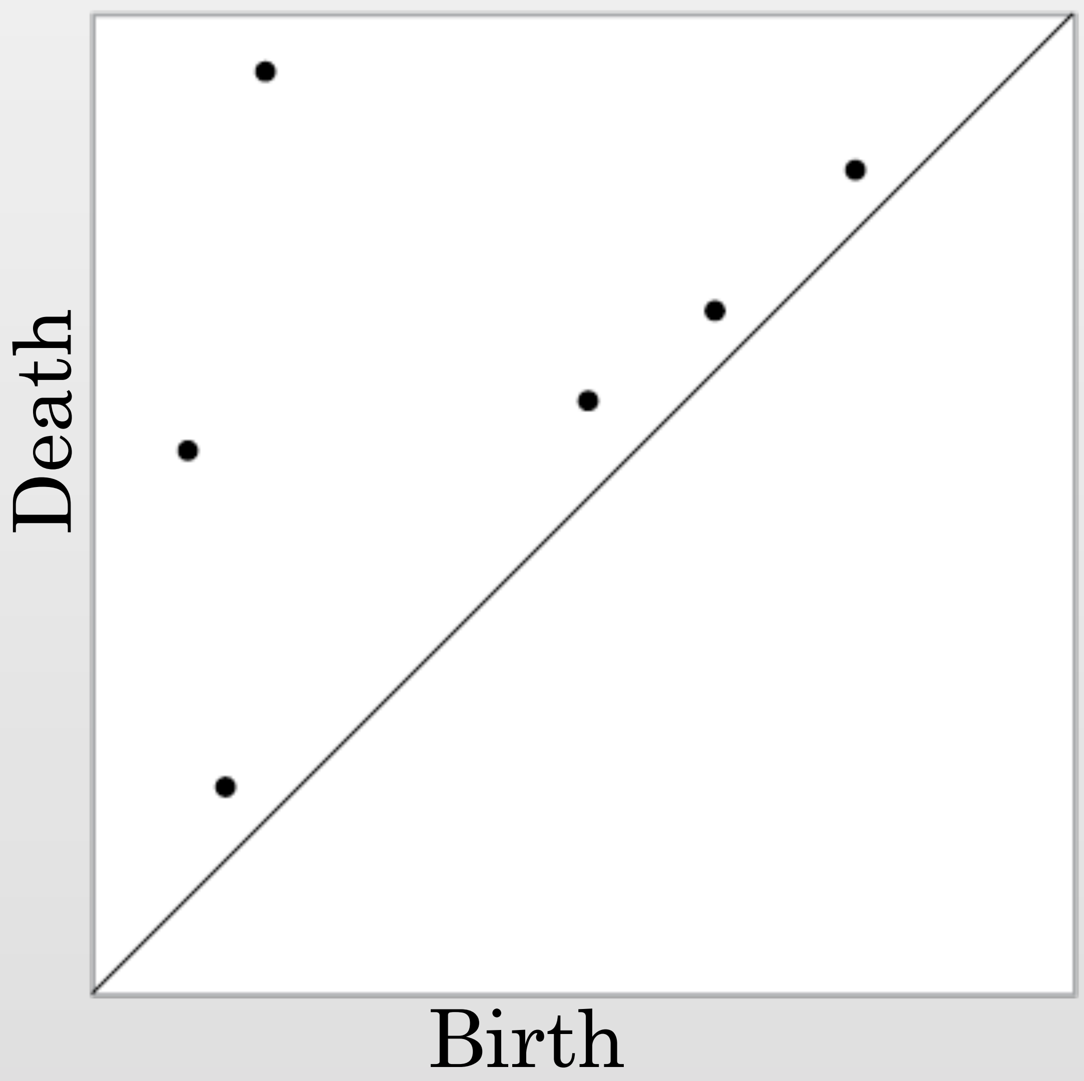

Persistence Diagrams describe the changes in topology corresponding to changes in scale.

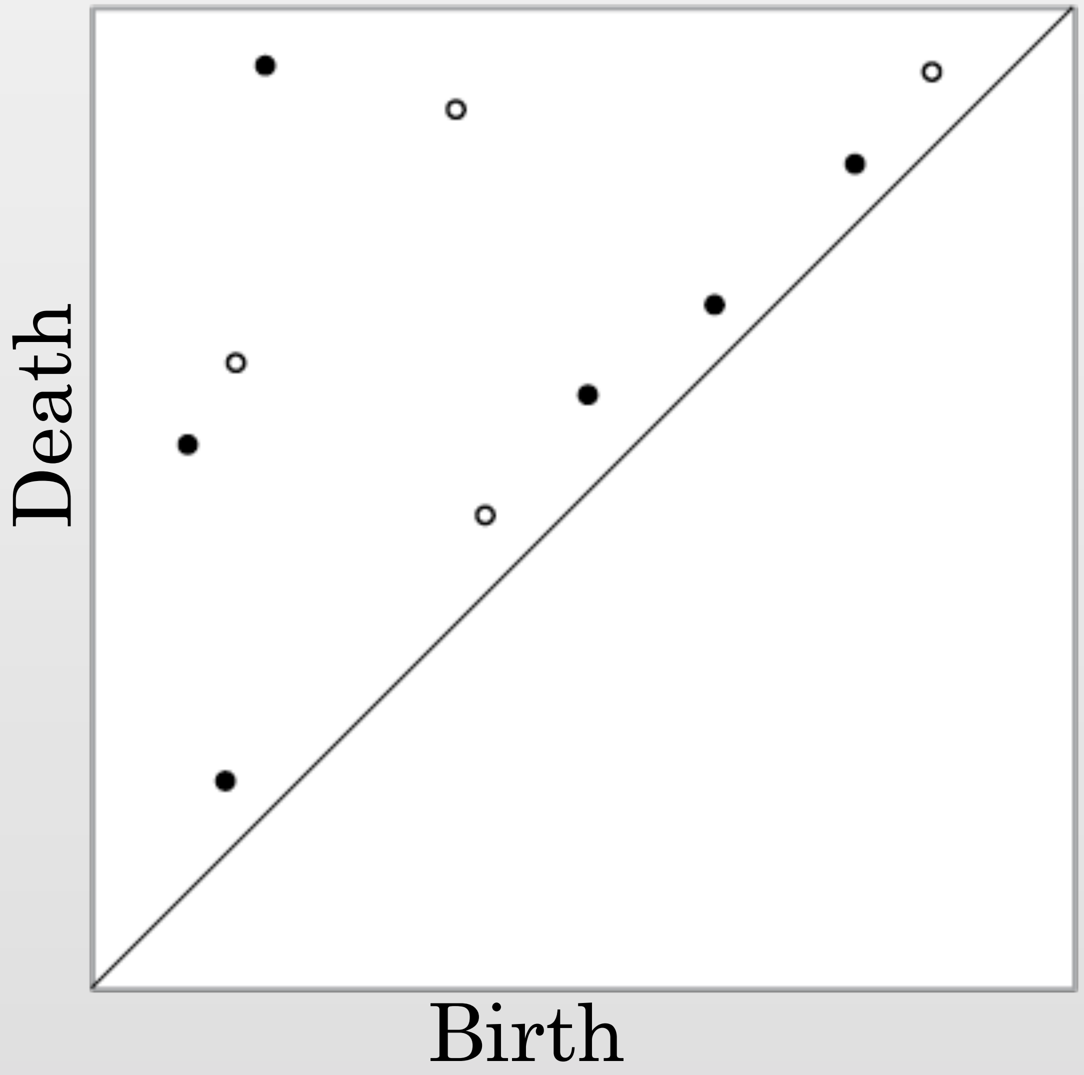

Persistence Diagrams describe the changes in topology corresponding to changes in scale.

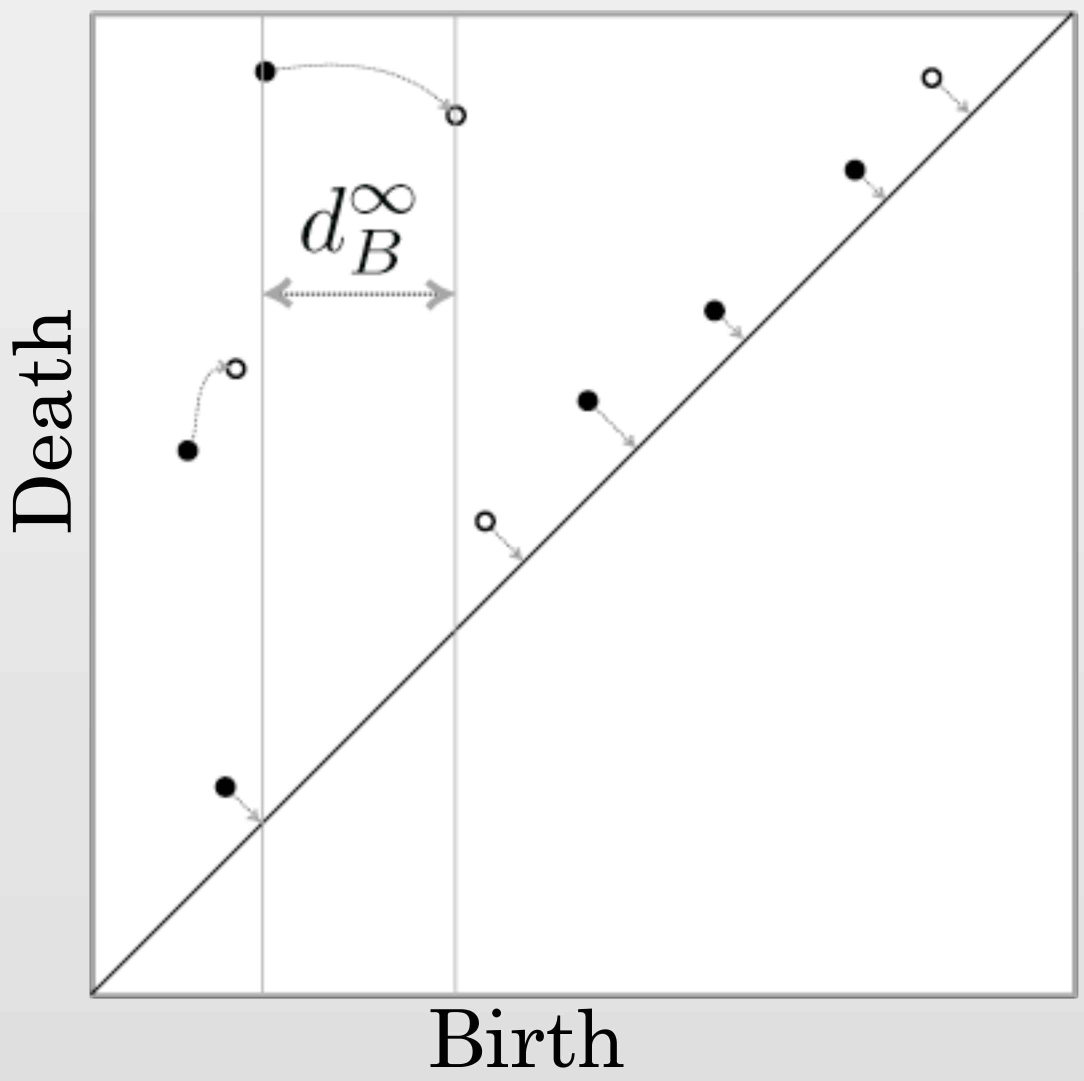

Bottleneck Distance

$$d_B^\infty = \max_i |p_i - q_i|_\infty$$

Persistence Diagrams describe the changes in topology corresponding to changes in scale.

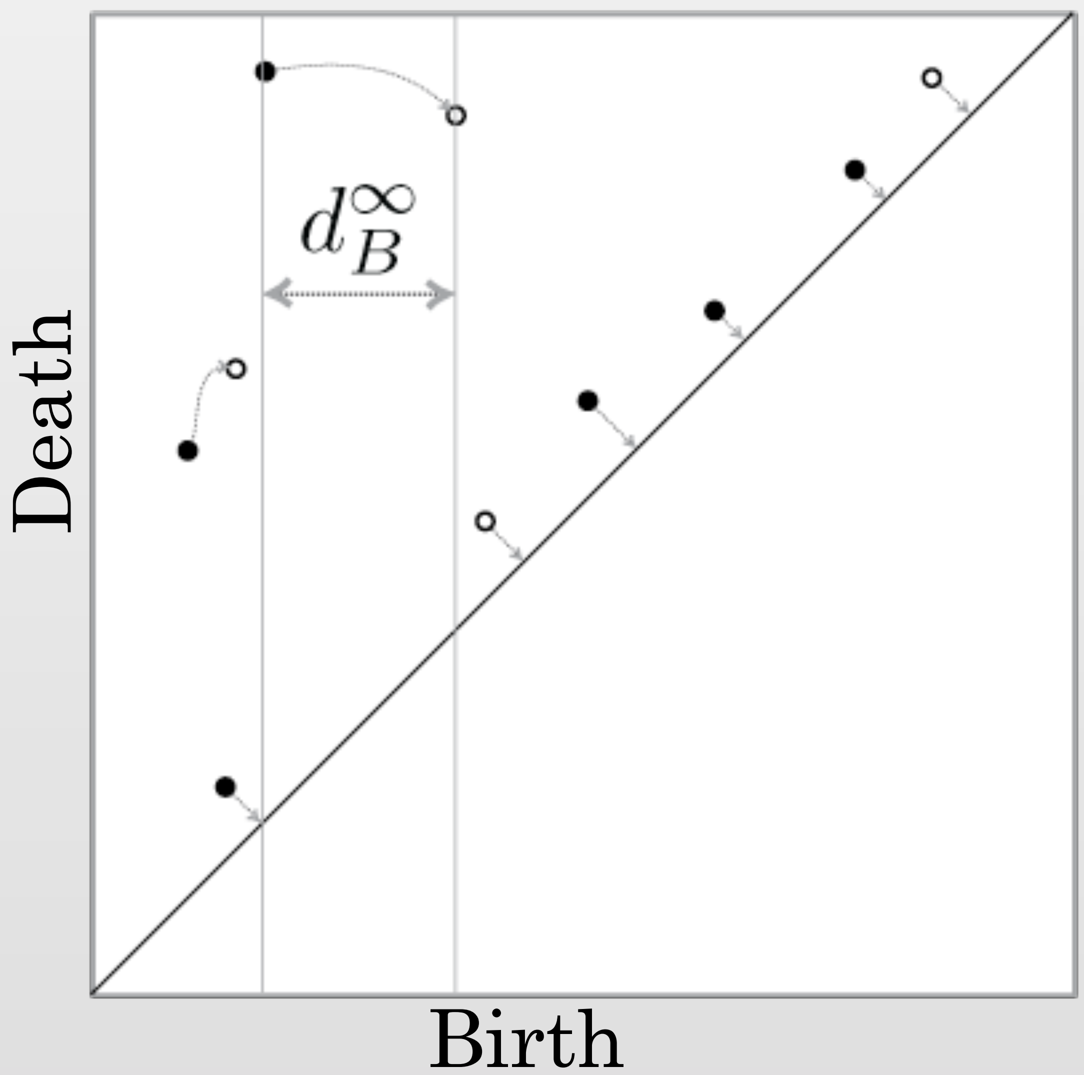

Bottleneck Distance

$$d_B^\infty = \max_i |p_i - q_i|_\infty$$

In *approximate persistence diagrams*, birth and death times differ by at most a constant factor.

Persistence Diagrams describe the changes in topology corresponding to changes in scale.

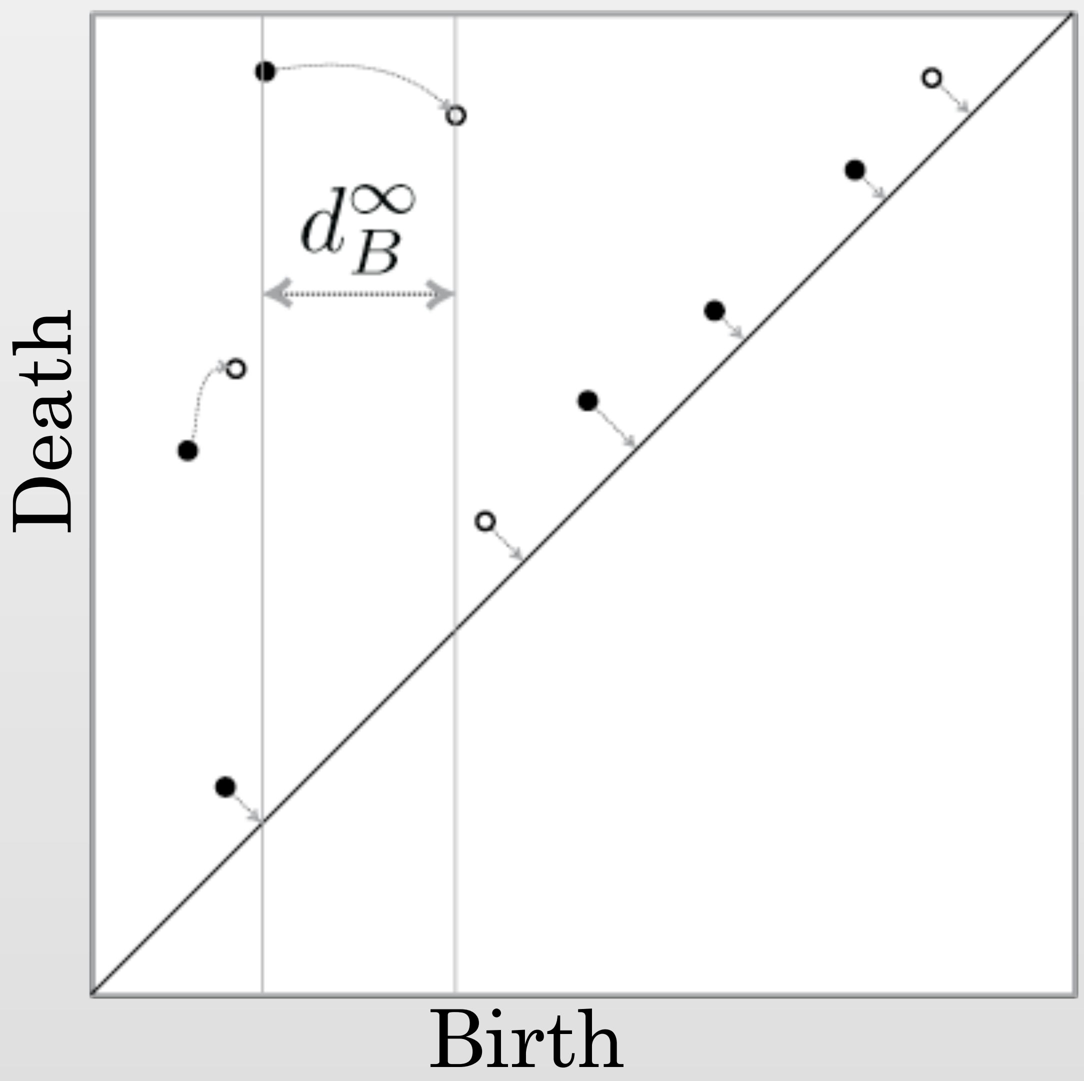

Bottleneck Distance

$$d_B^\infty = \max_i |p_i - q_i|_\infty$$

In *approximate persistence diagrams*, birth and death times differ by at most a constant factor.

This is just the bottleneck distance of the log-scale diagrams.

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

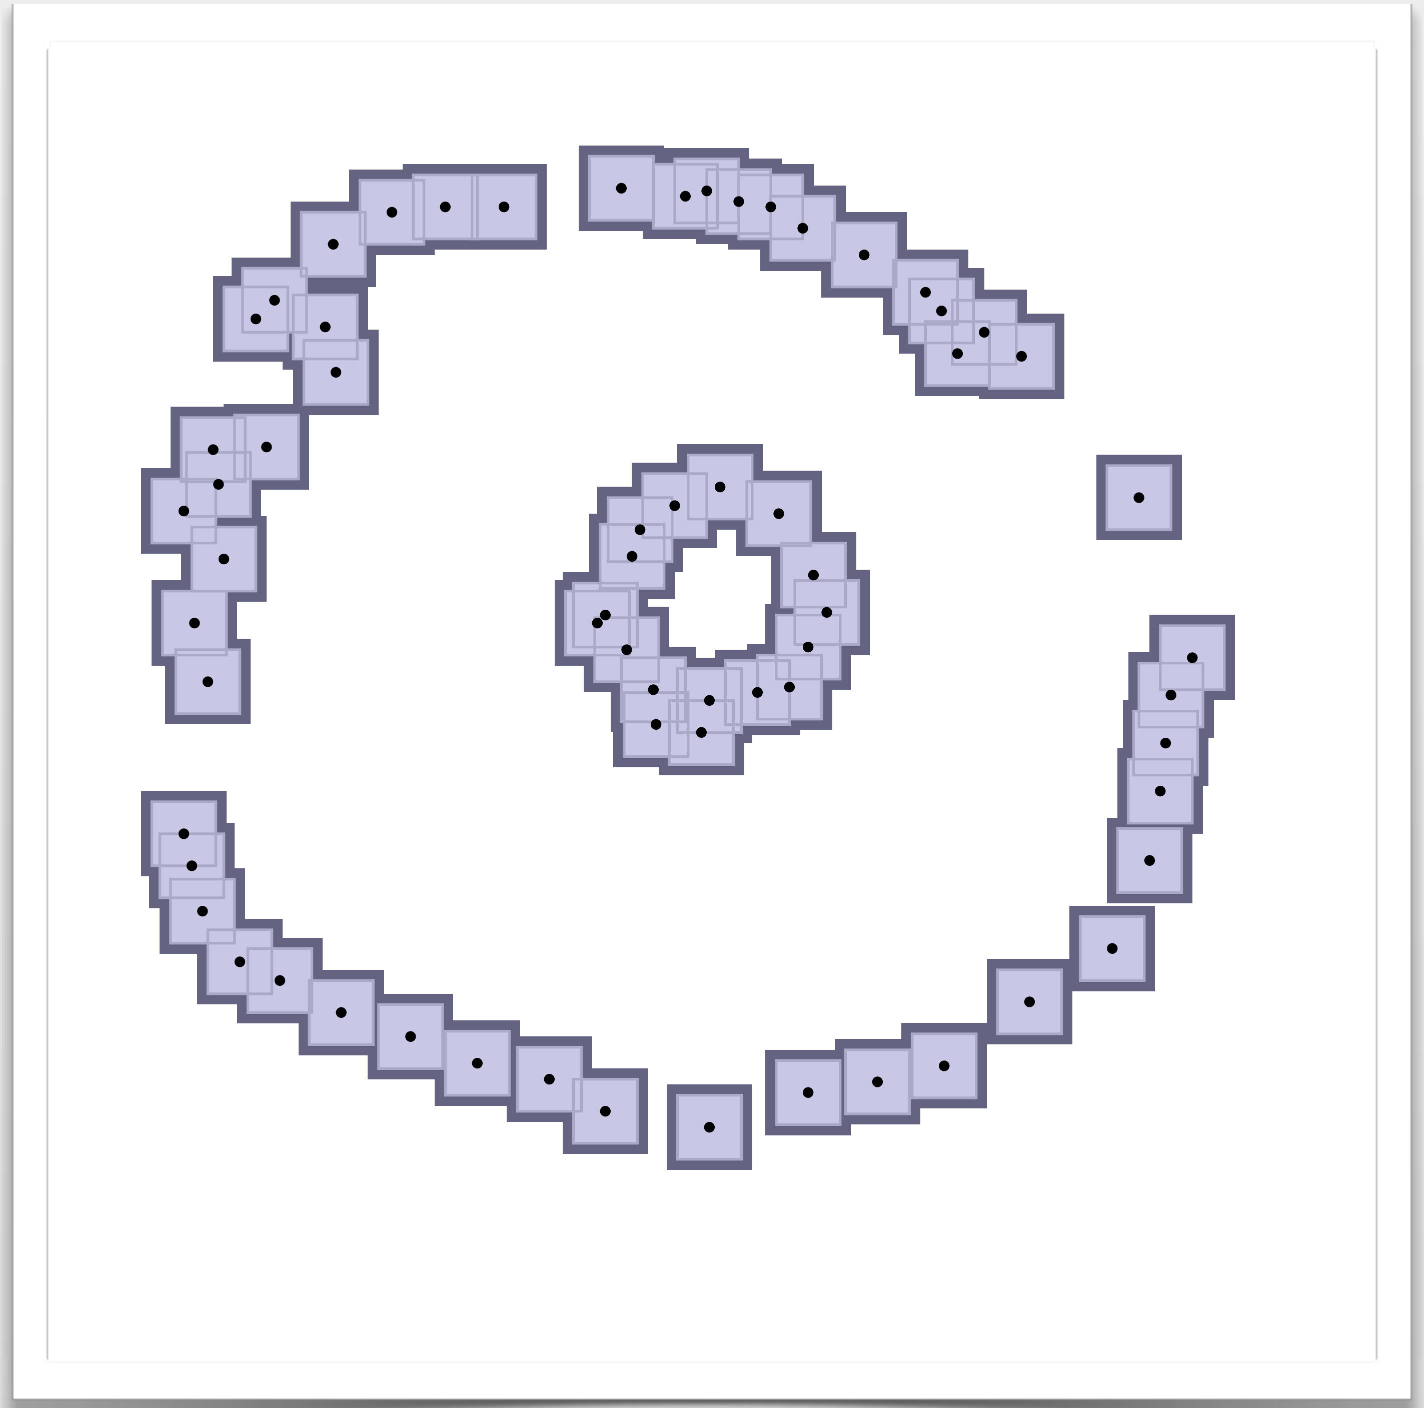

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

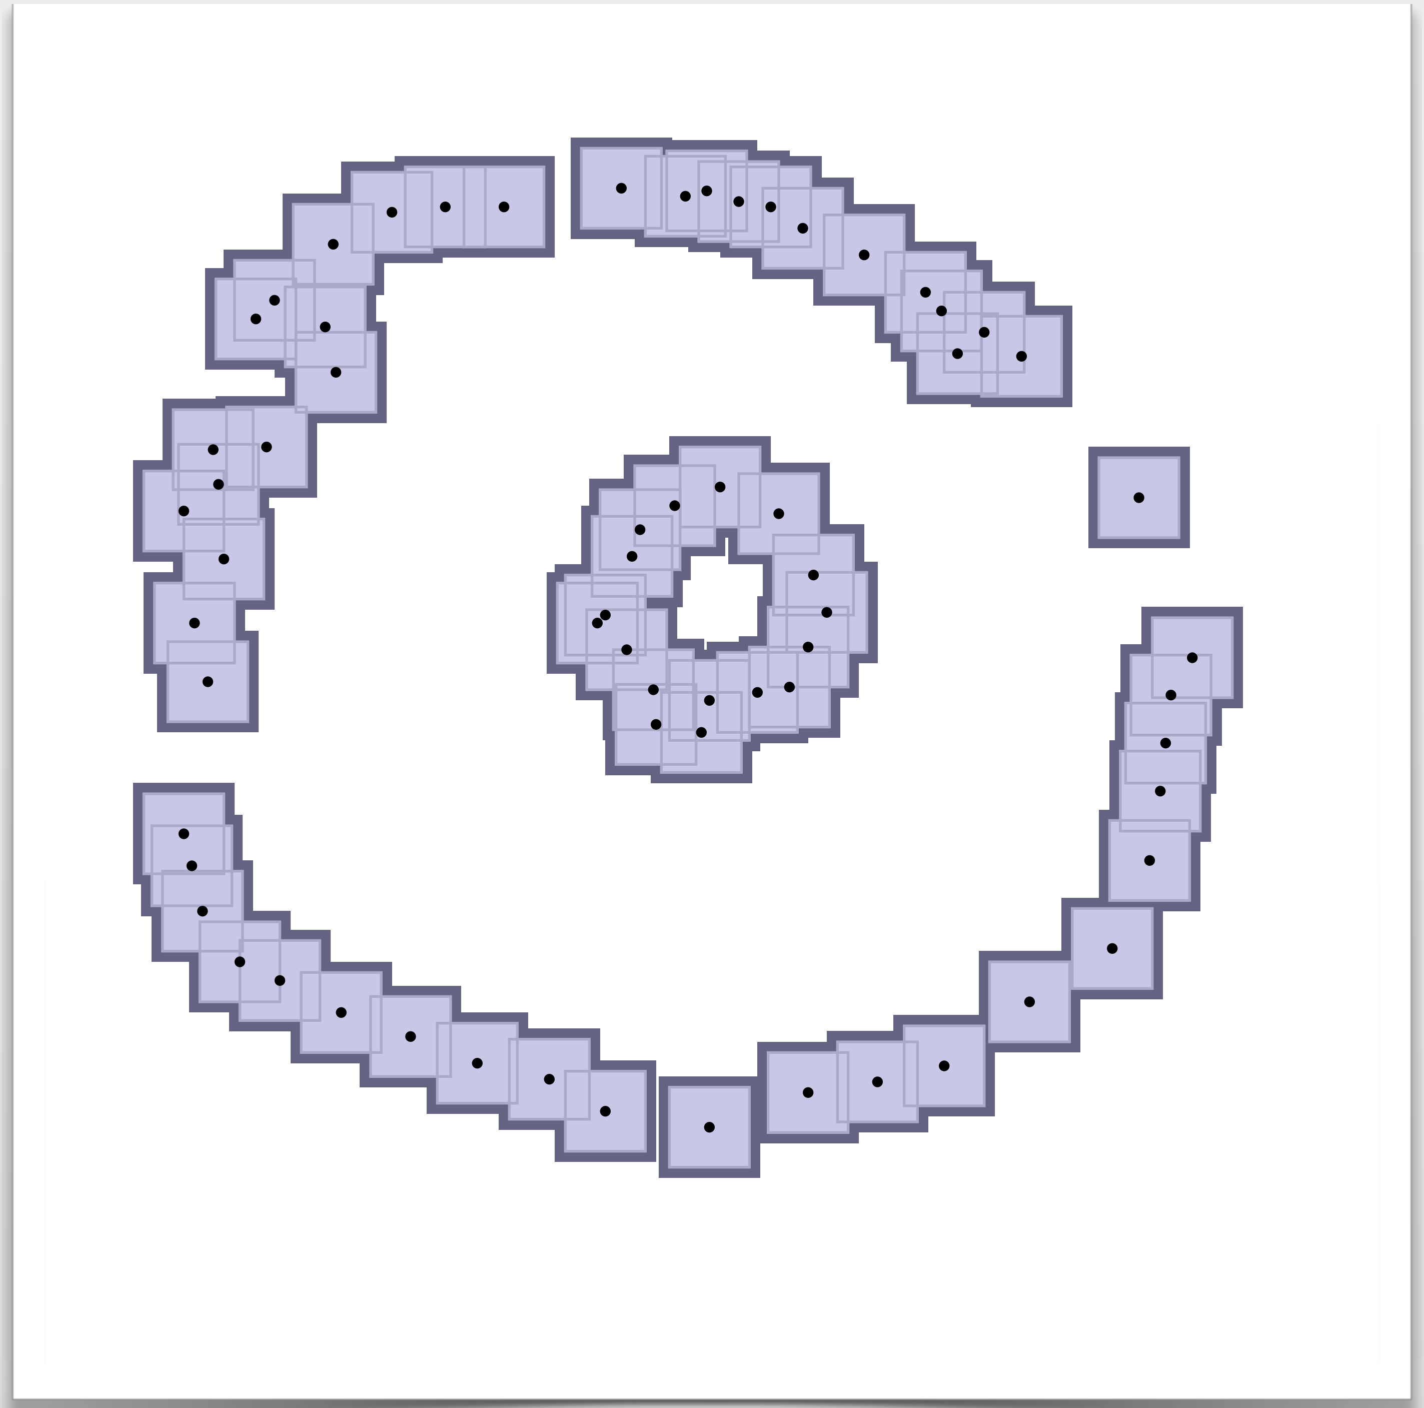

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

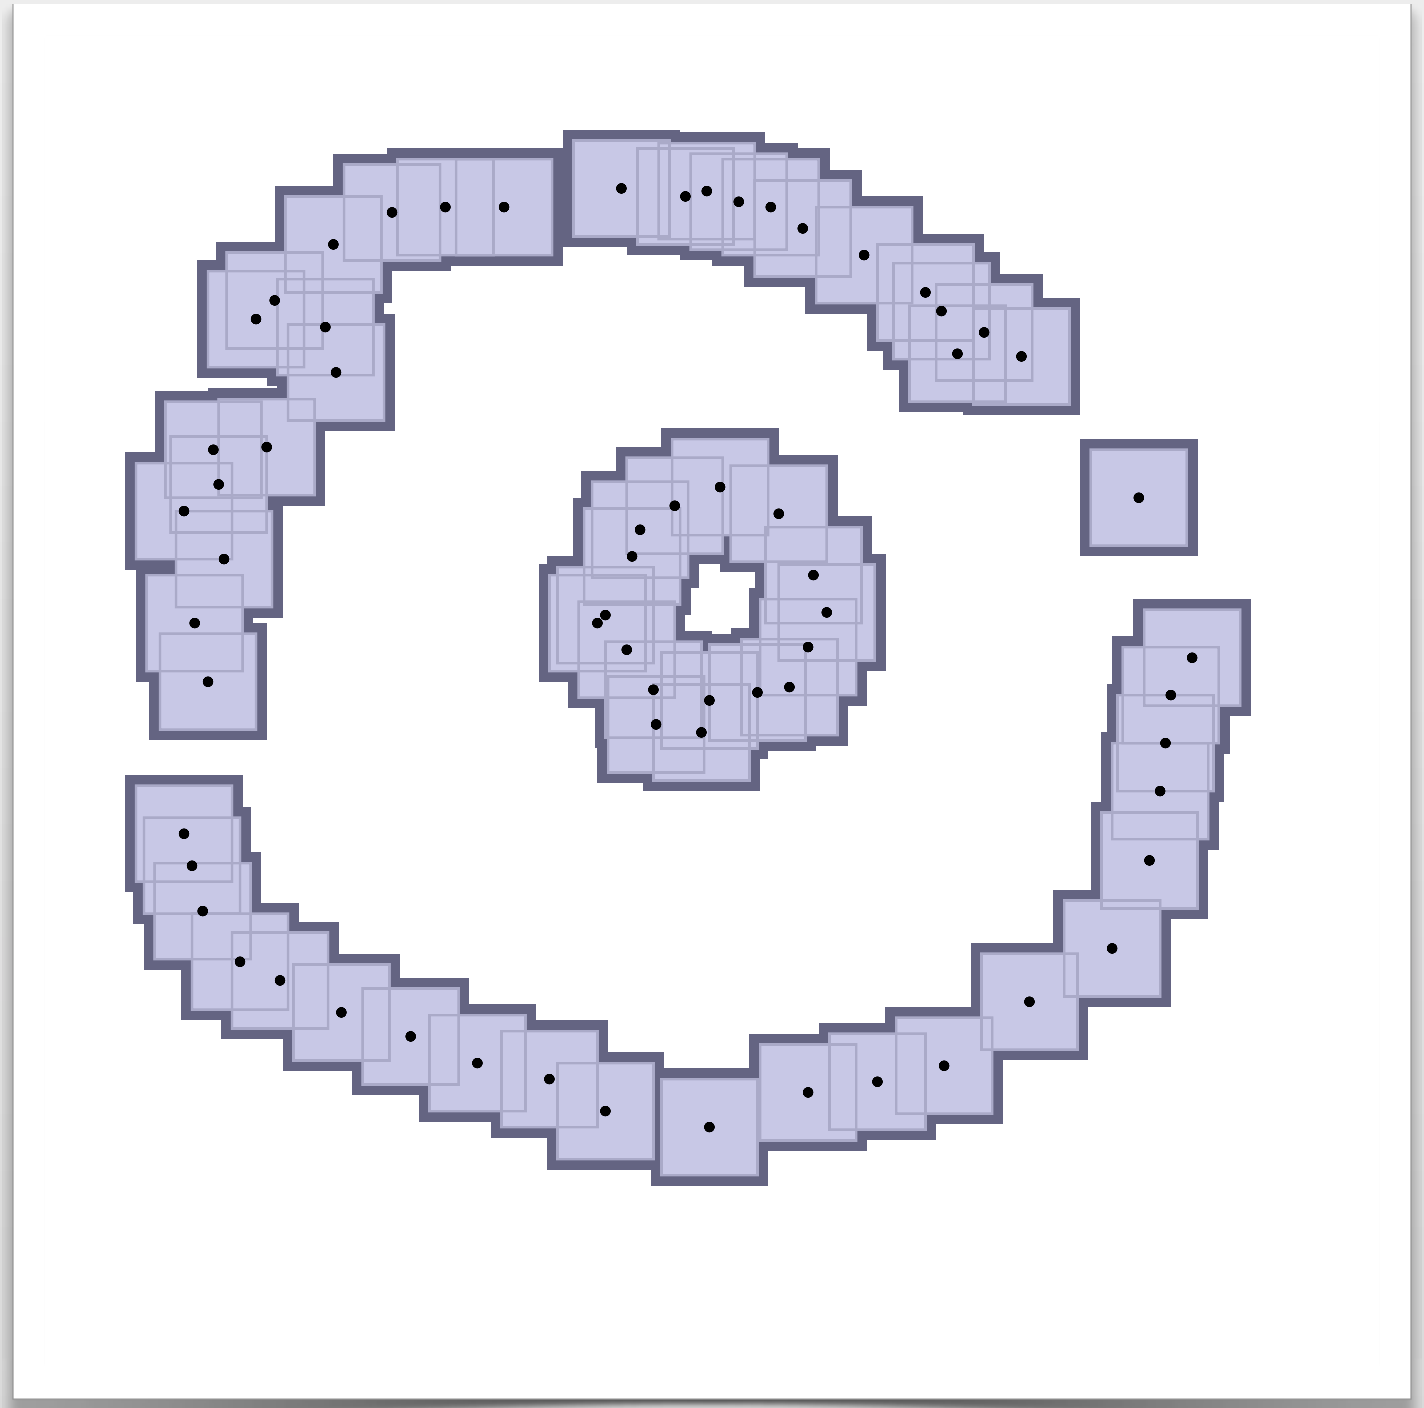

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

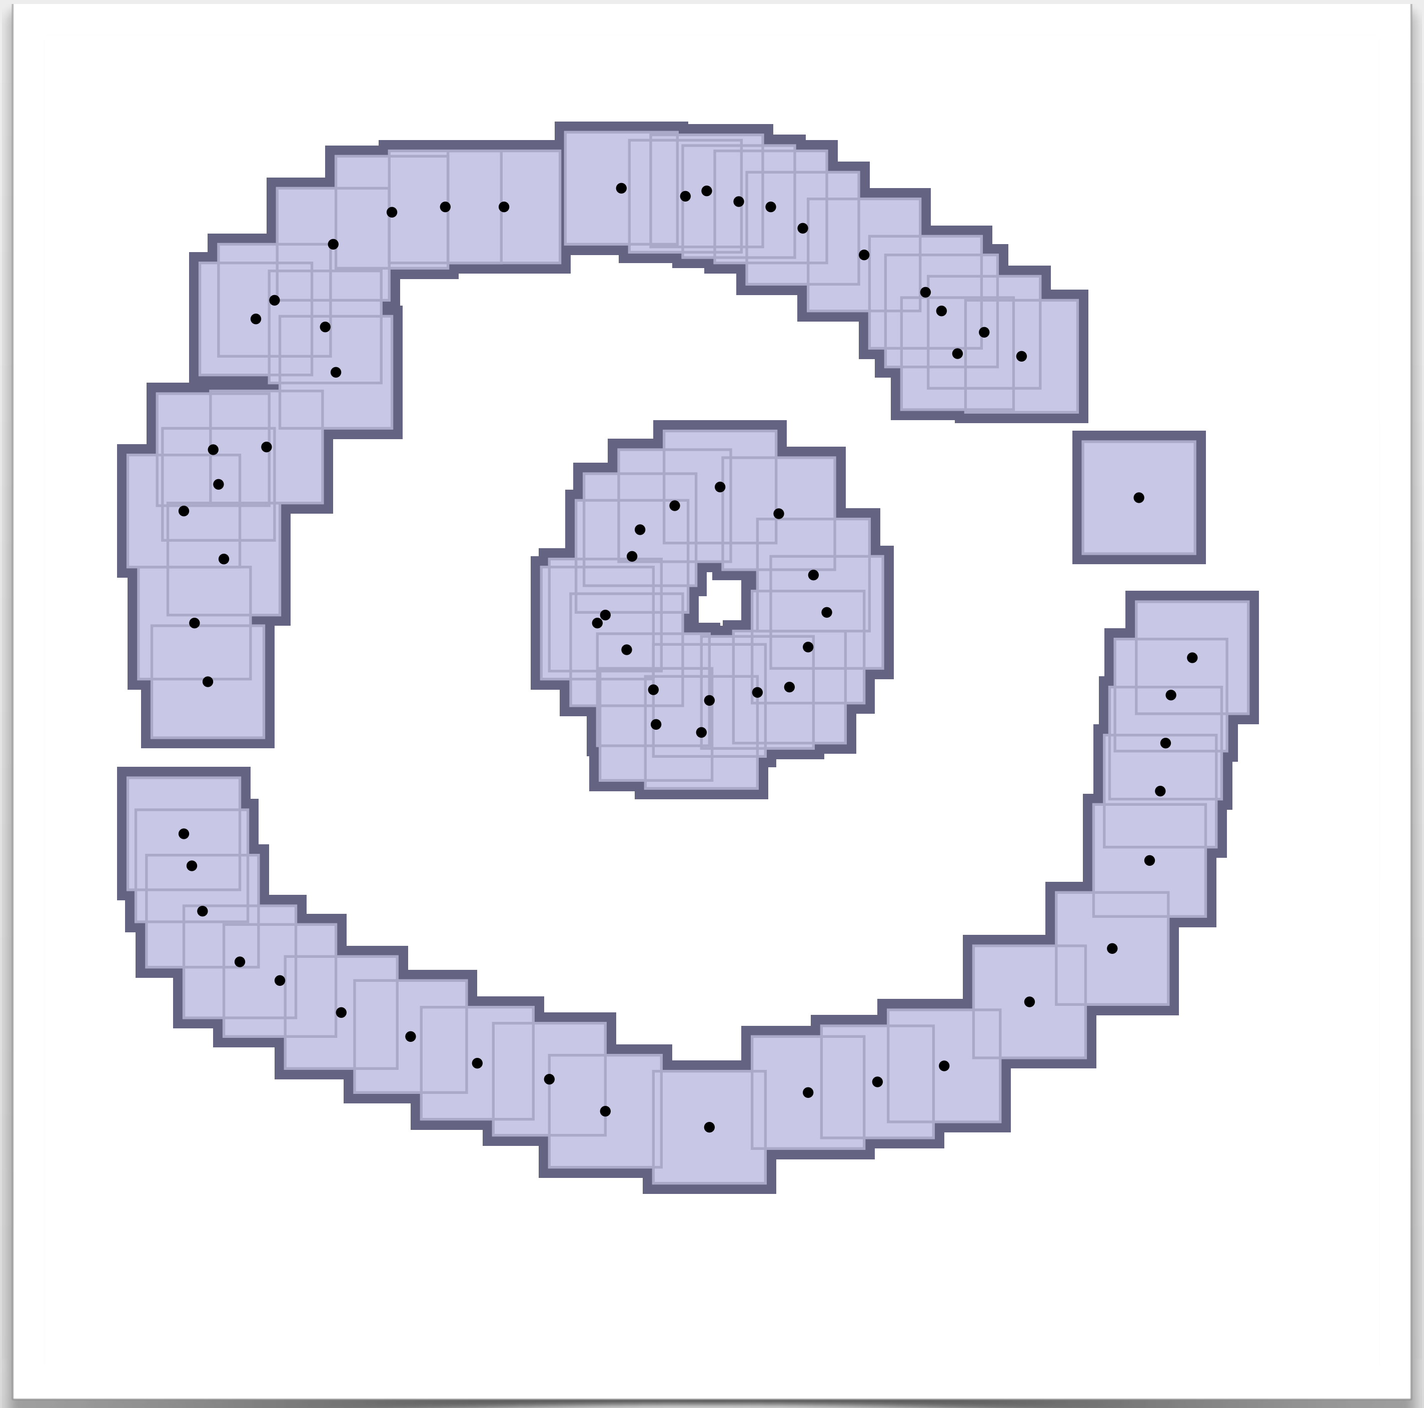

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

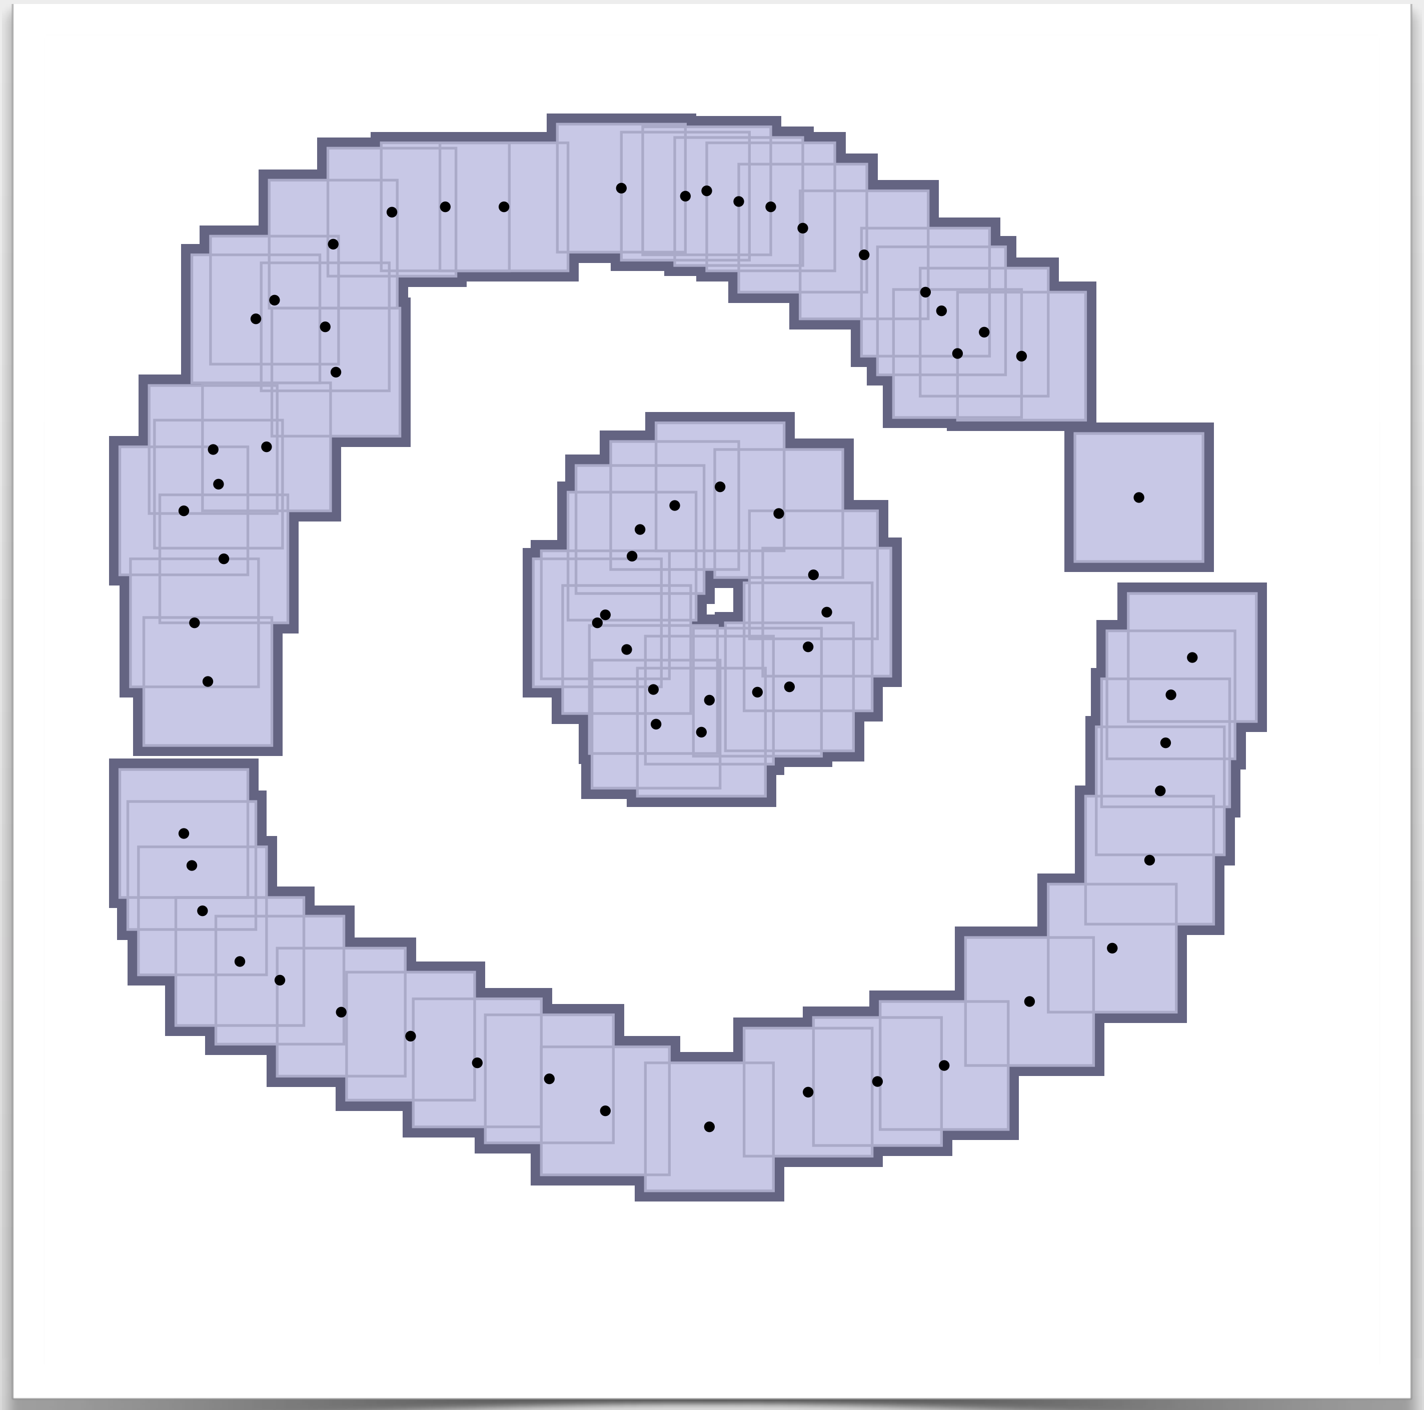

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

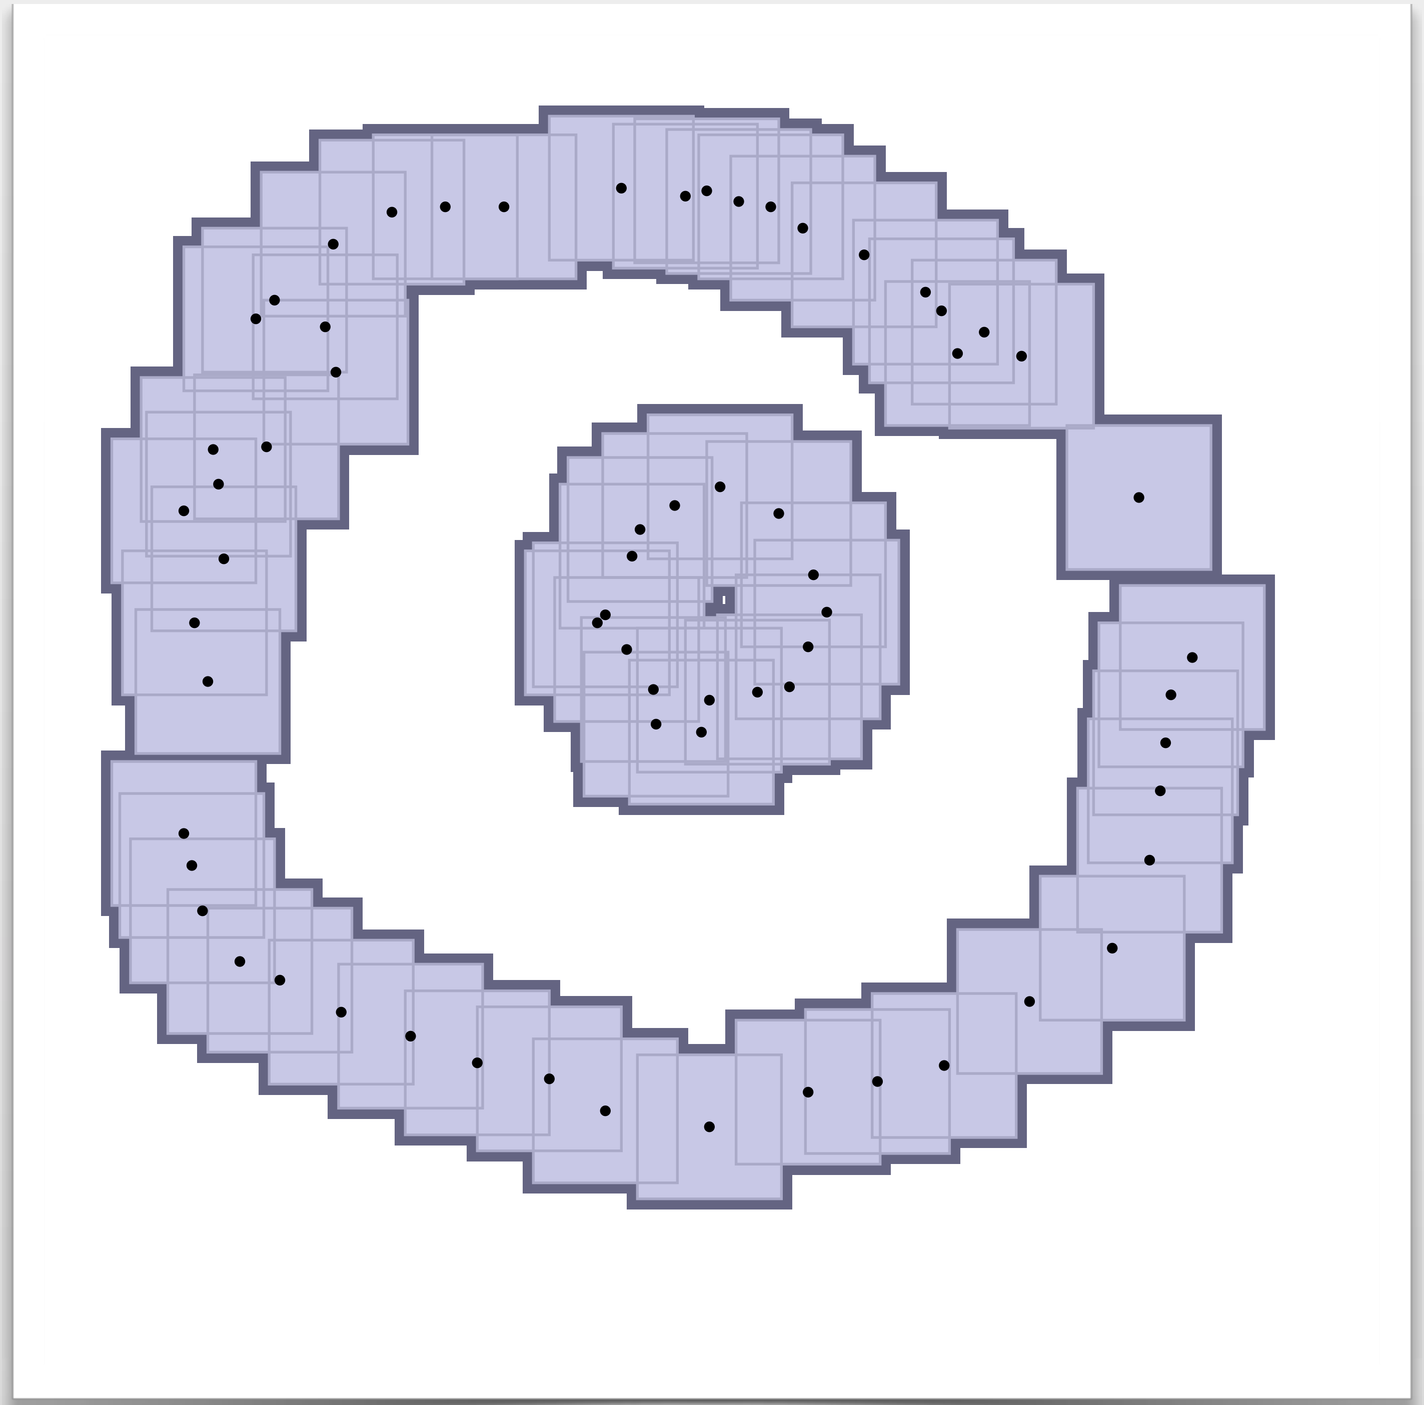

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

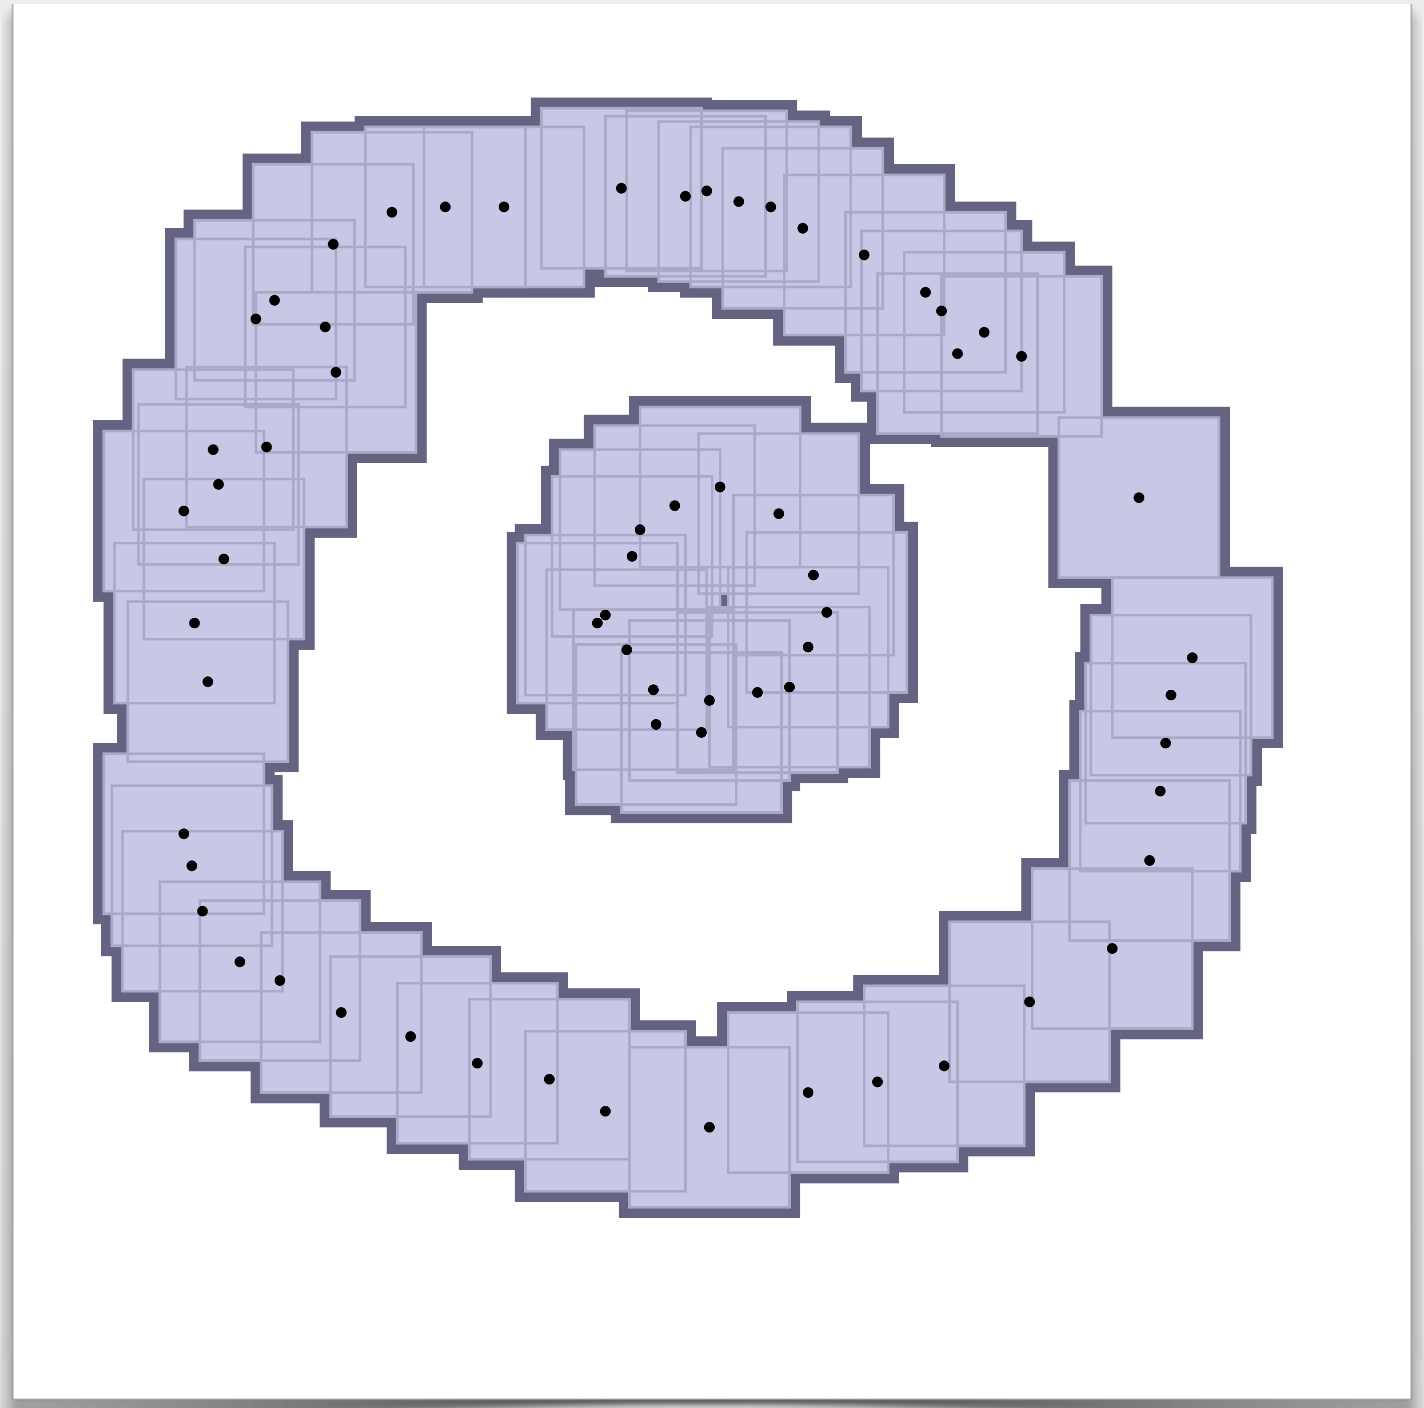

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

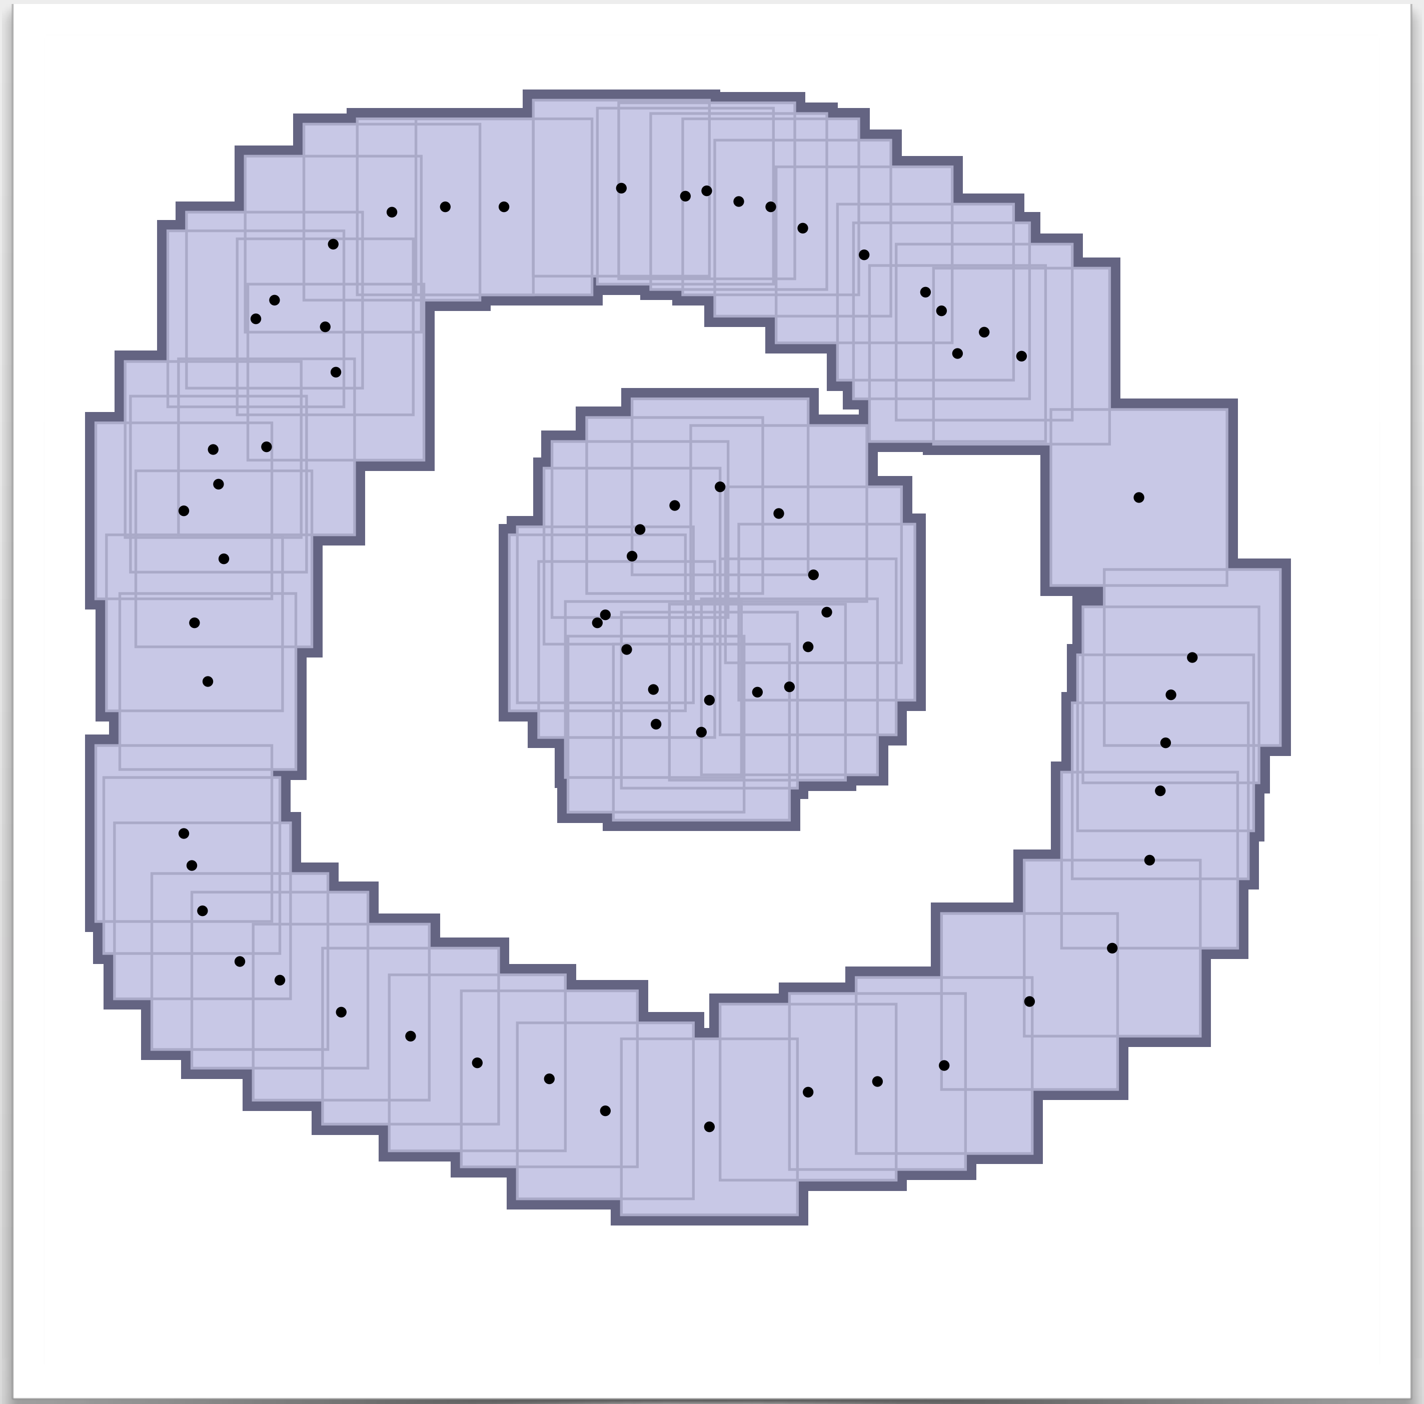

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

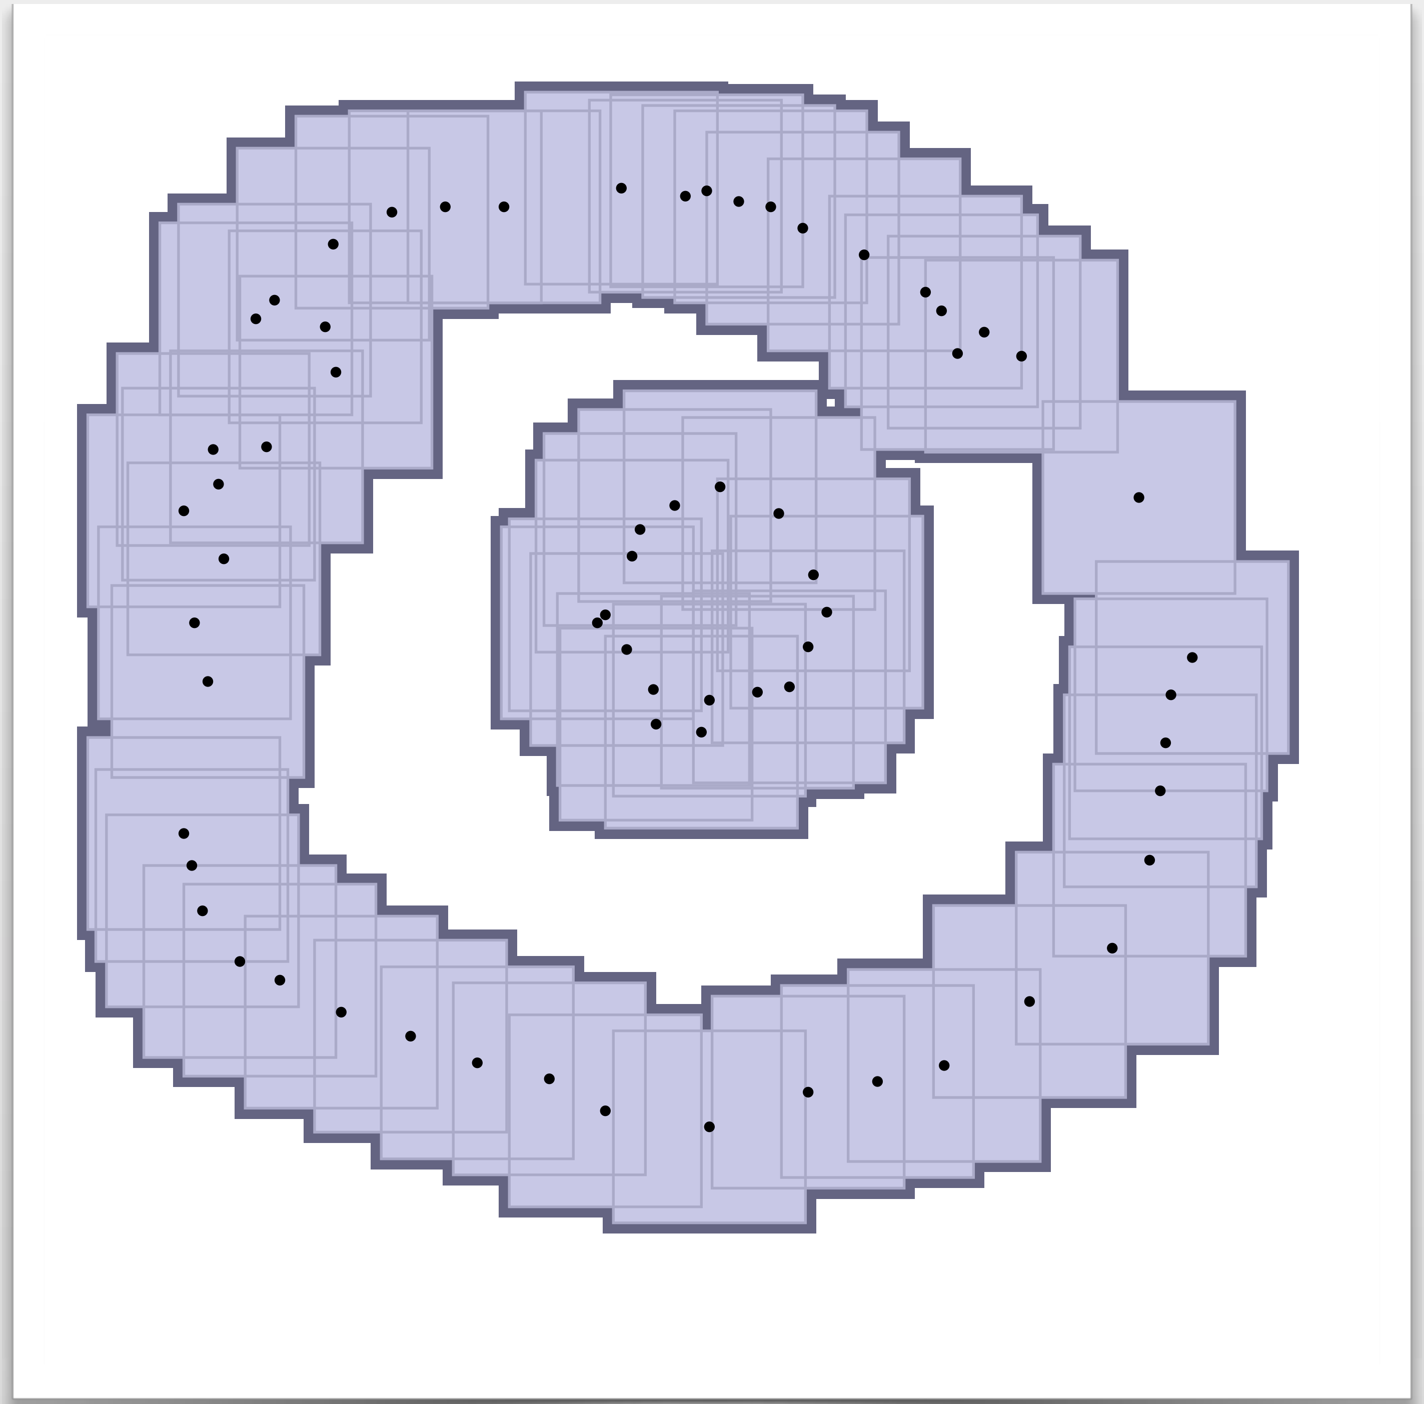

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

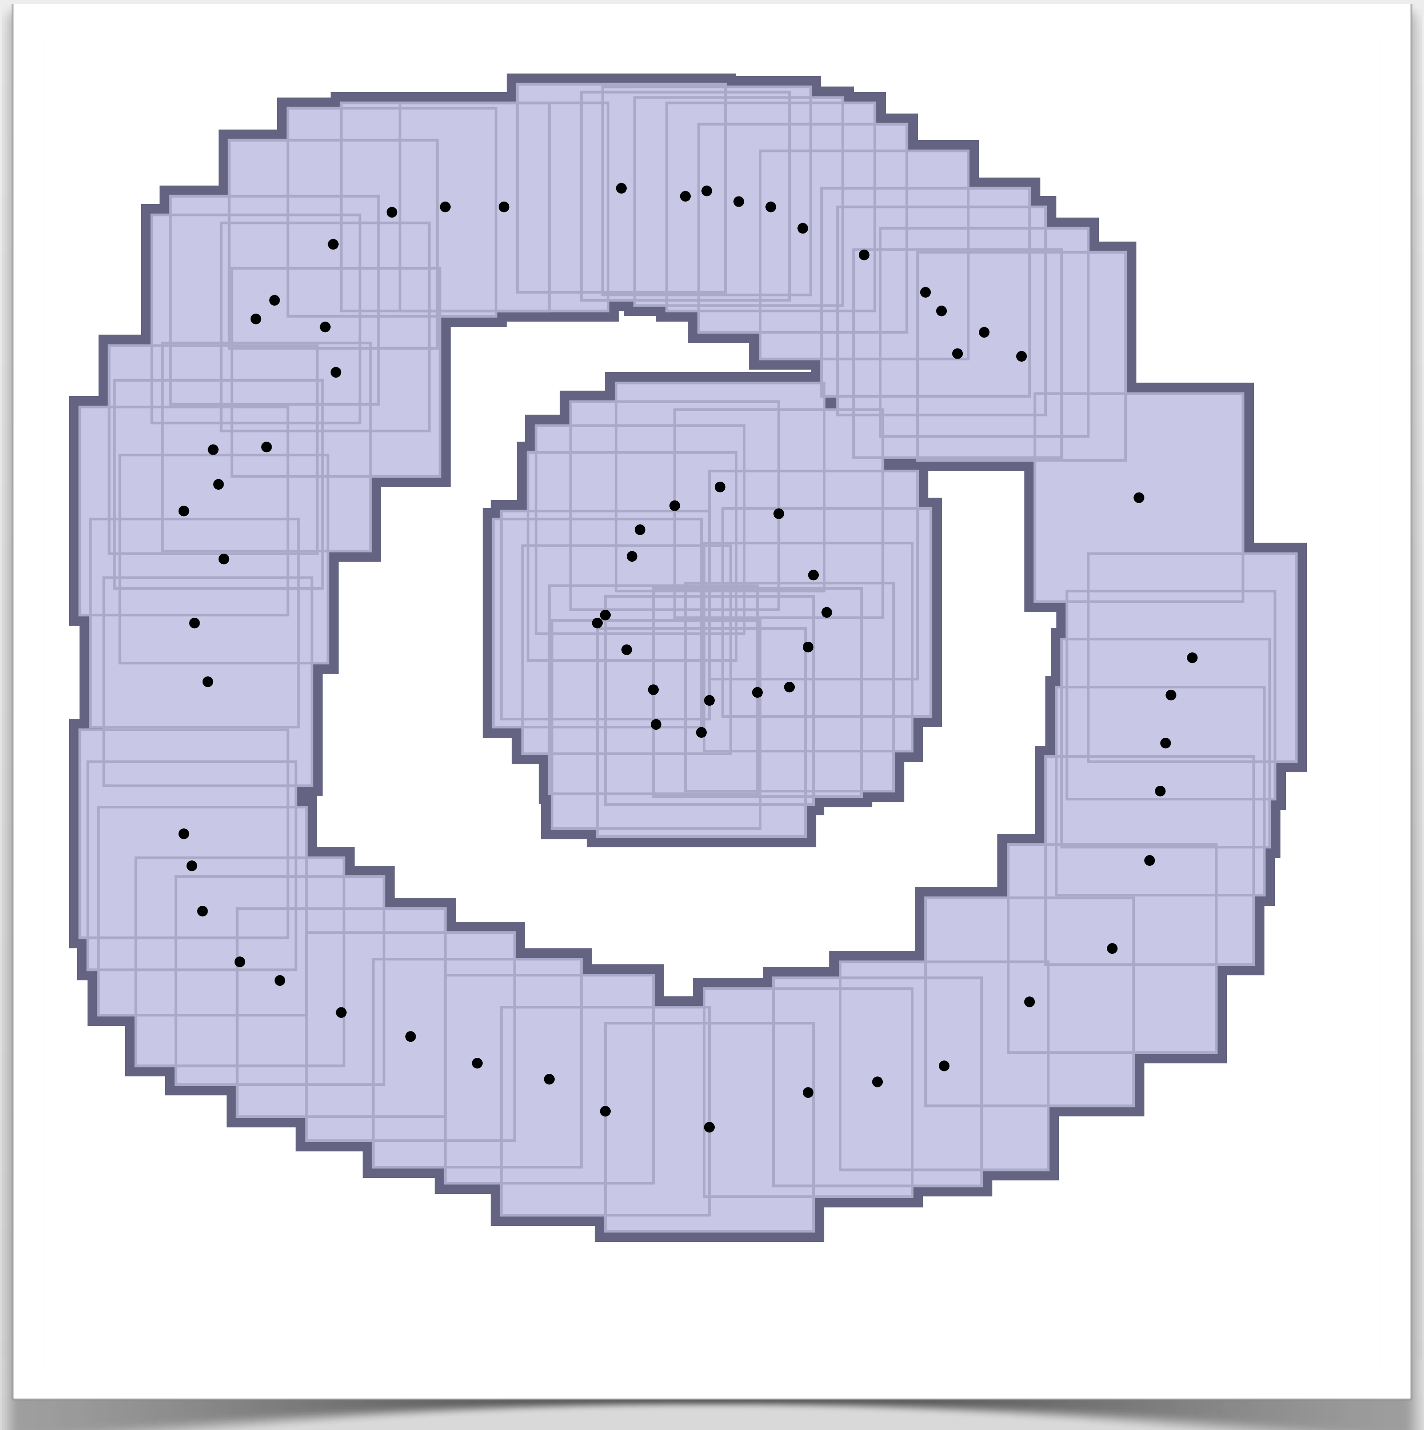

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

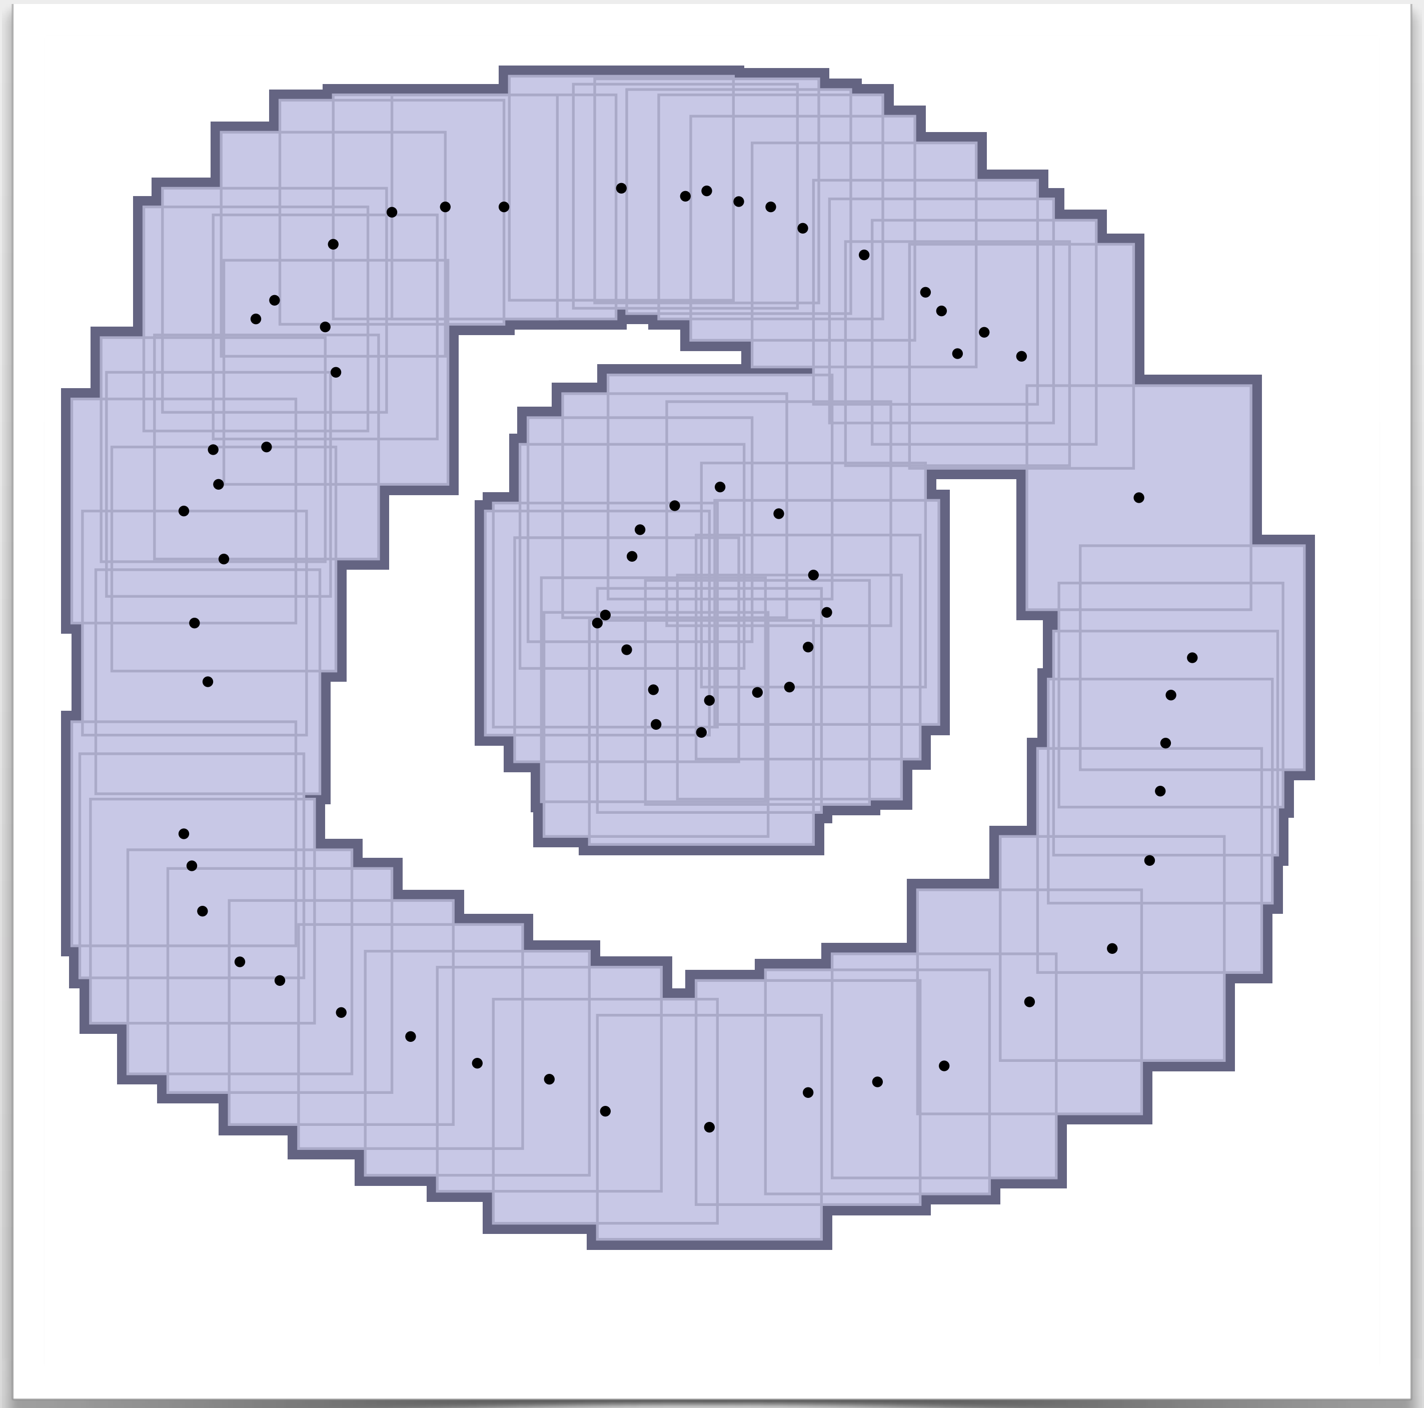

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

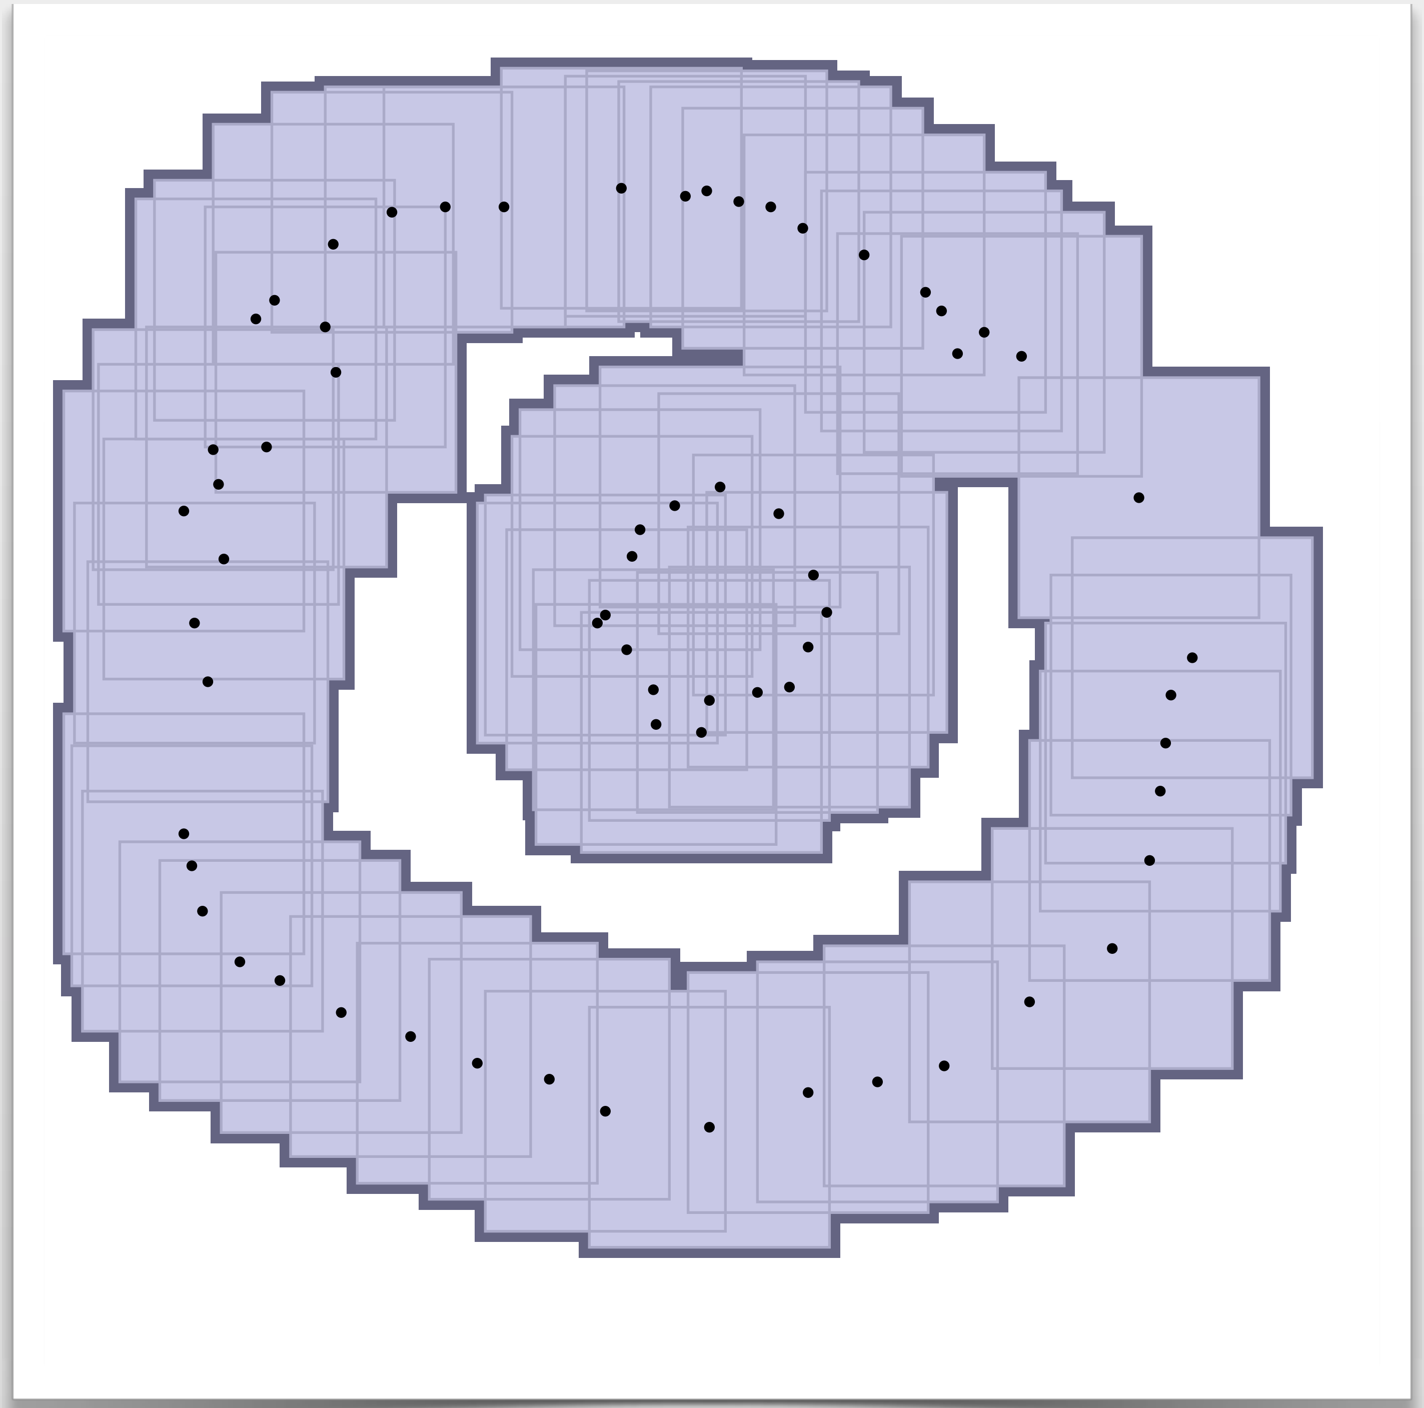

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

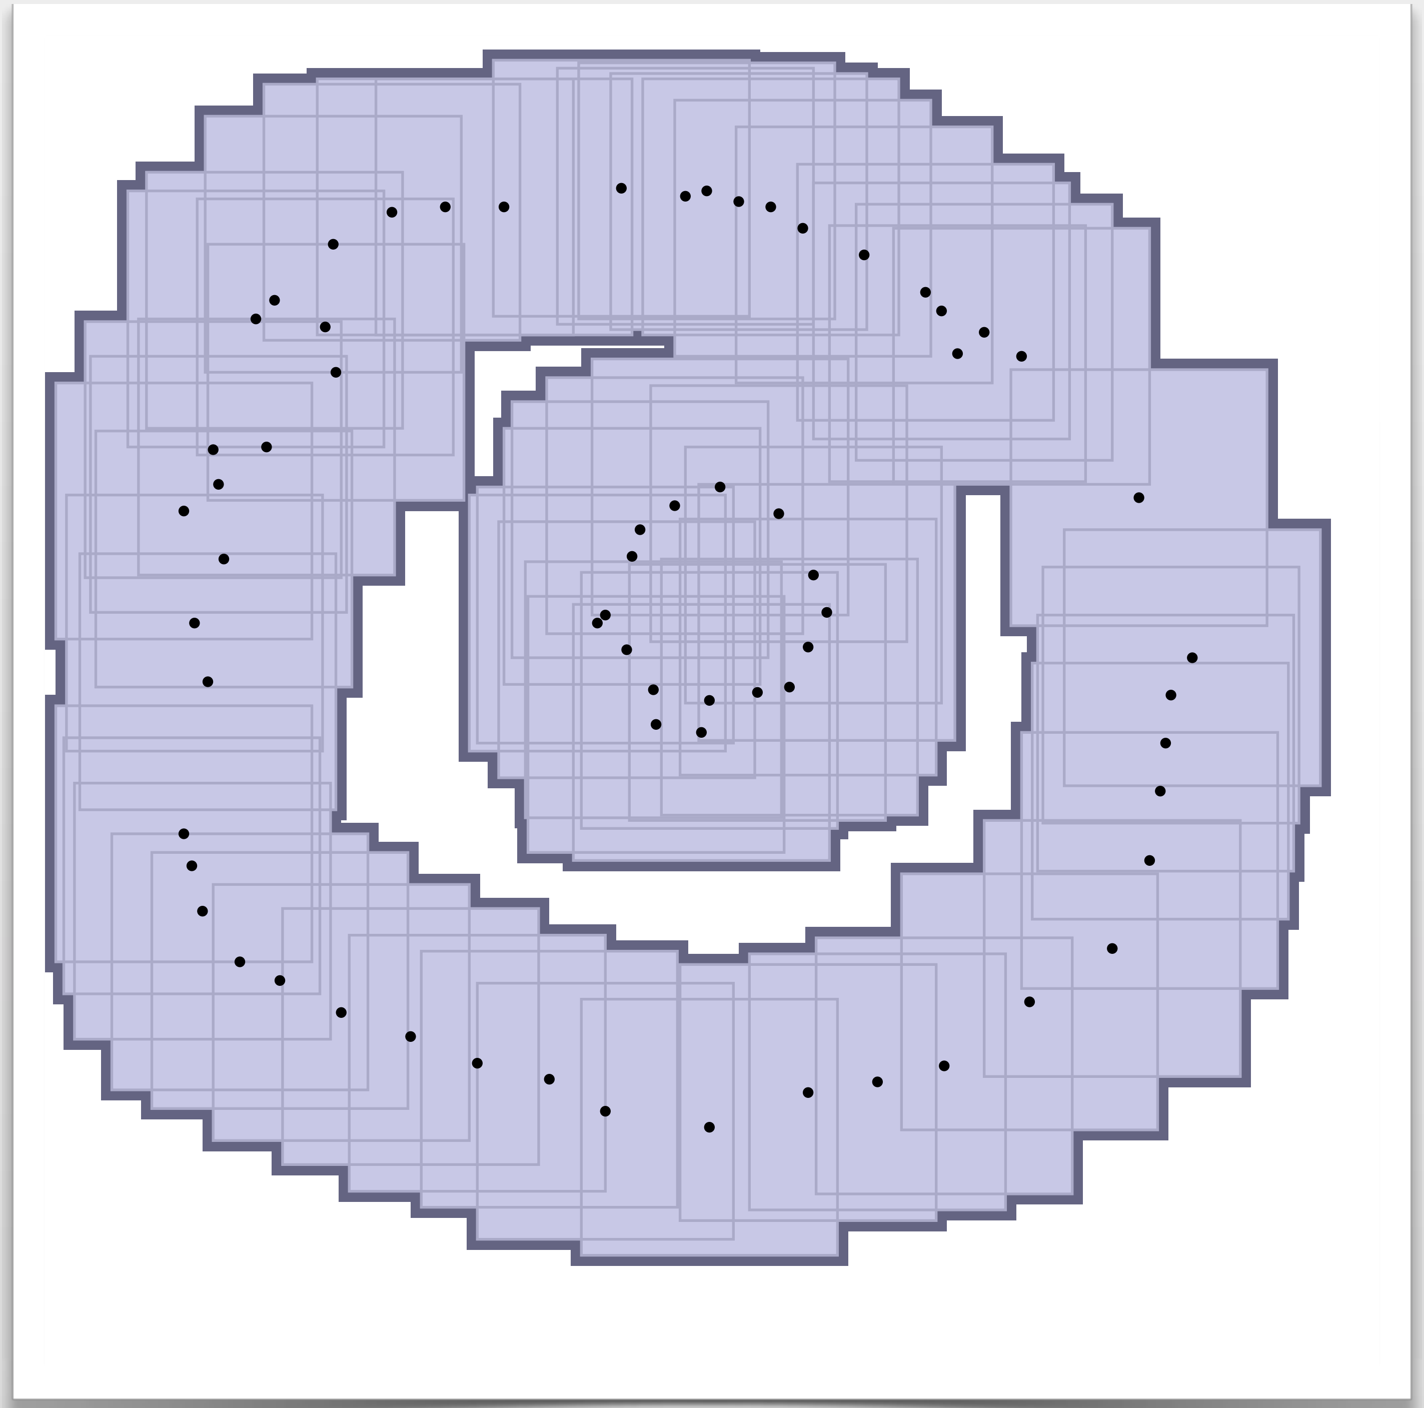

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

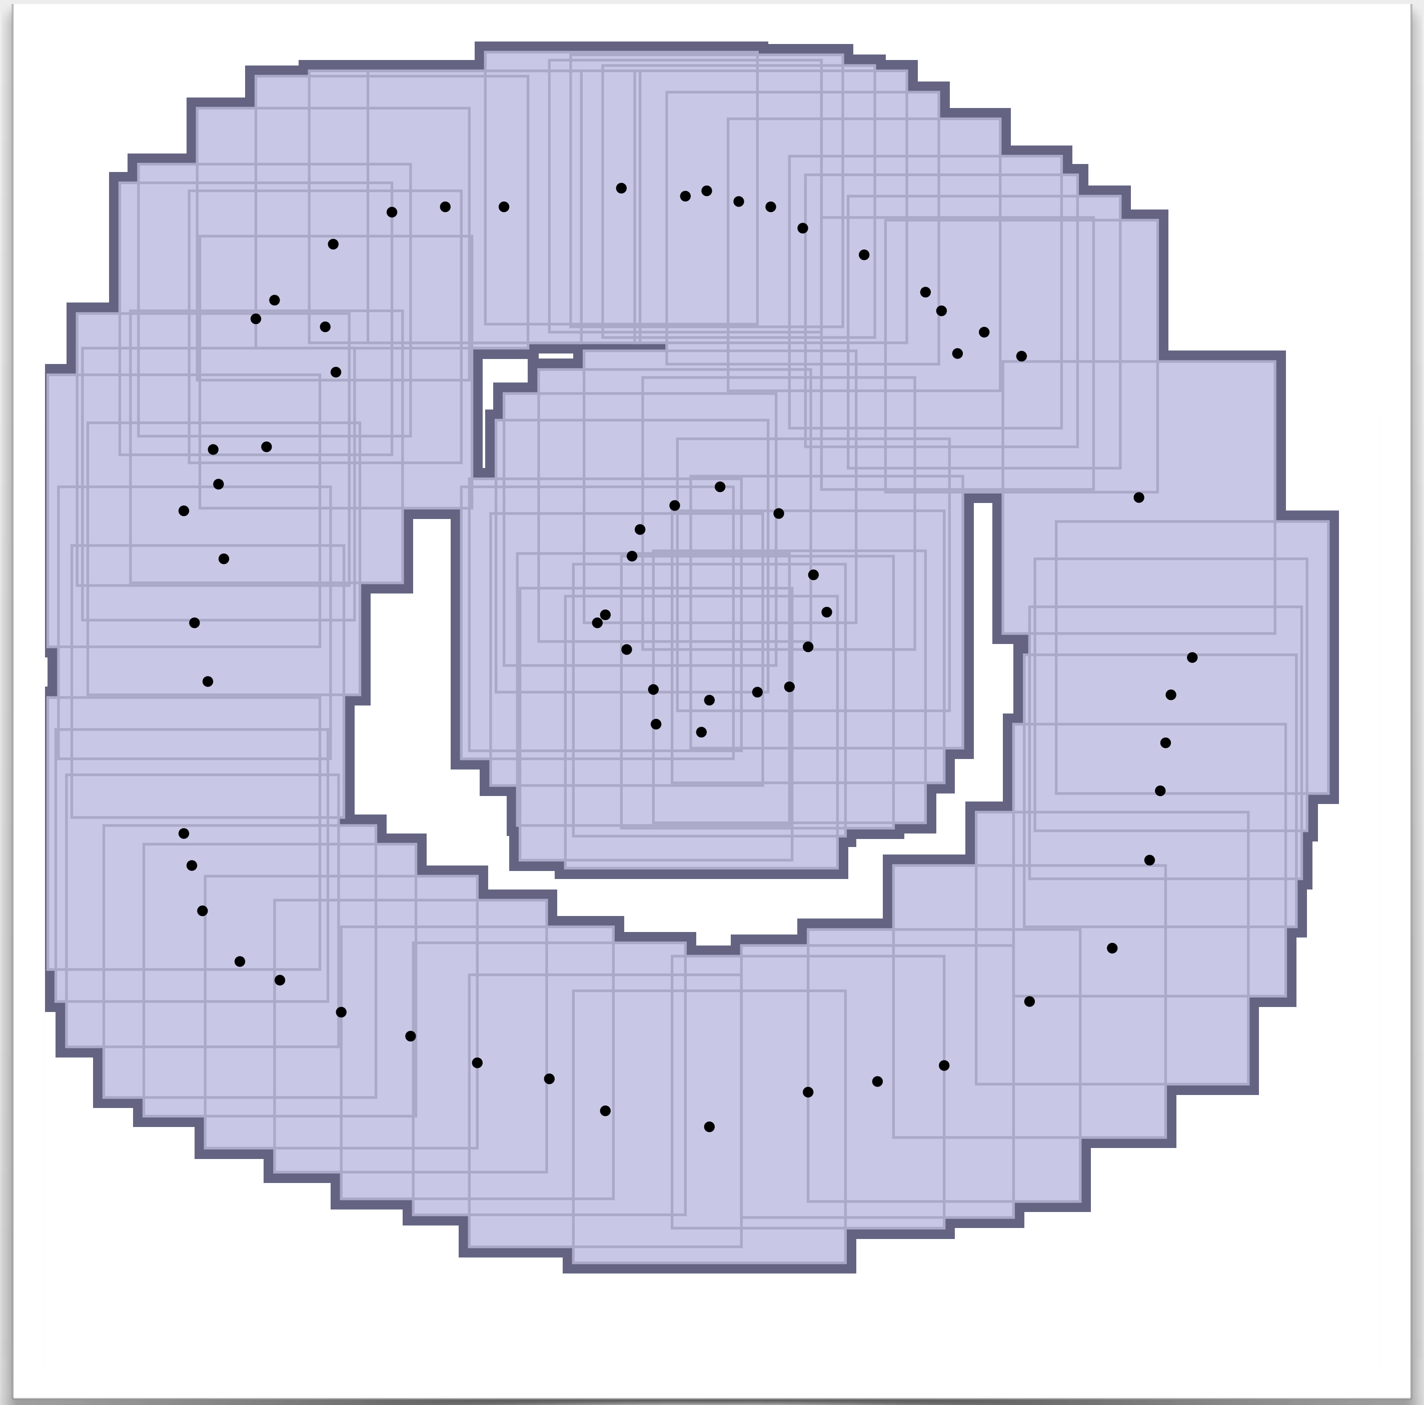

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

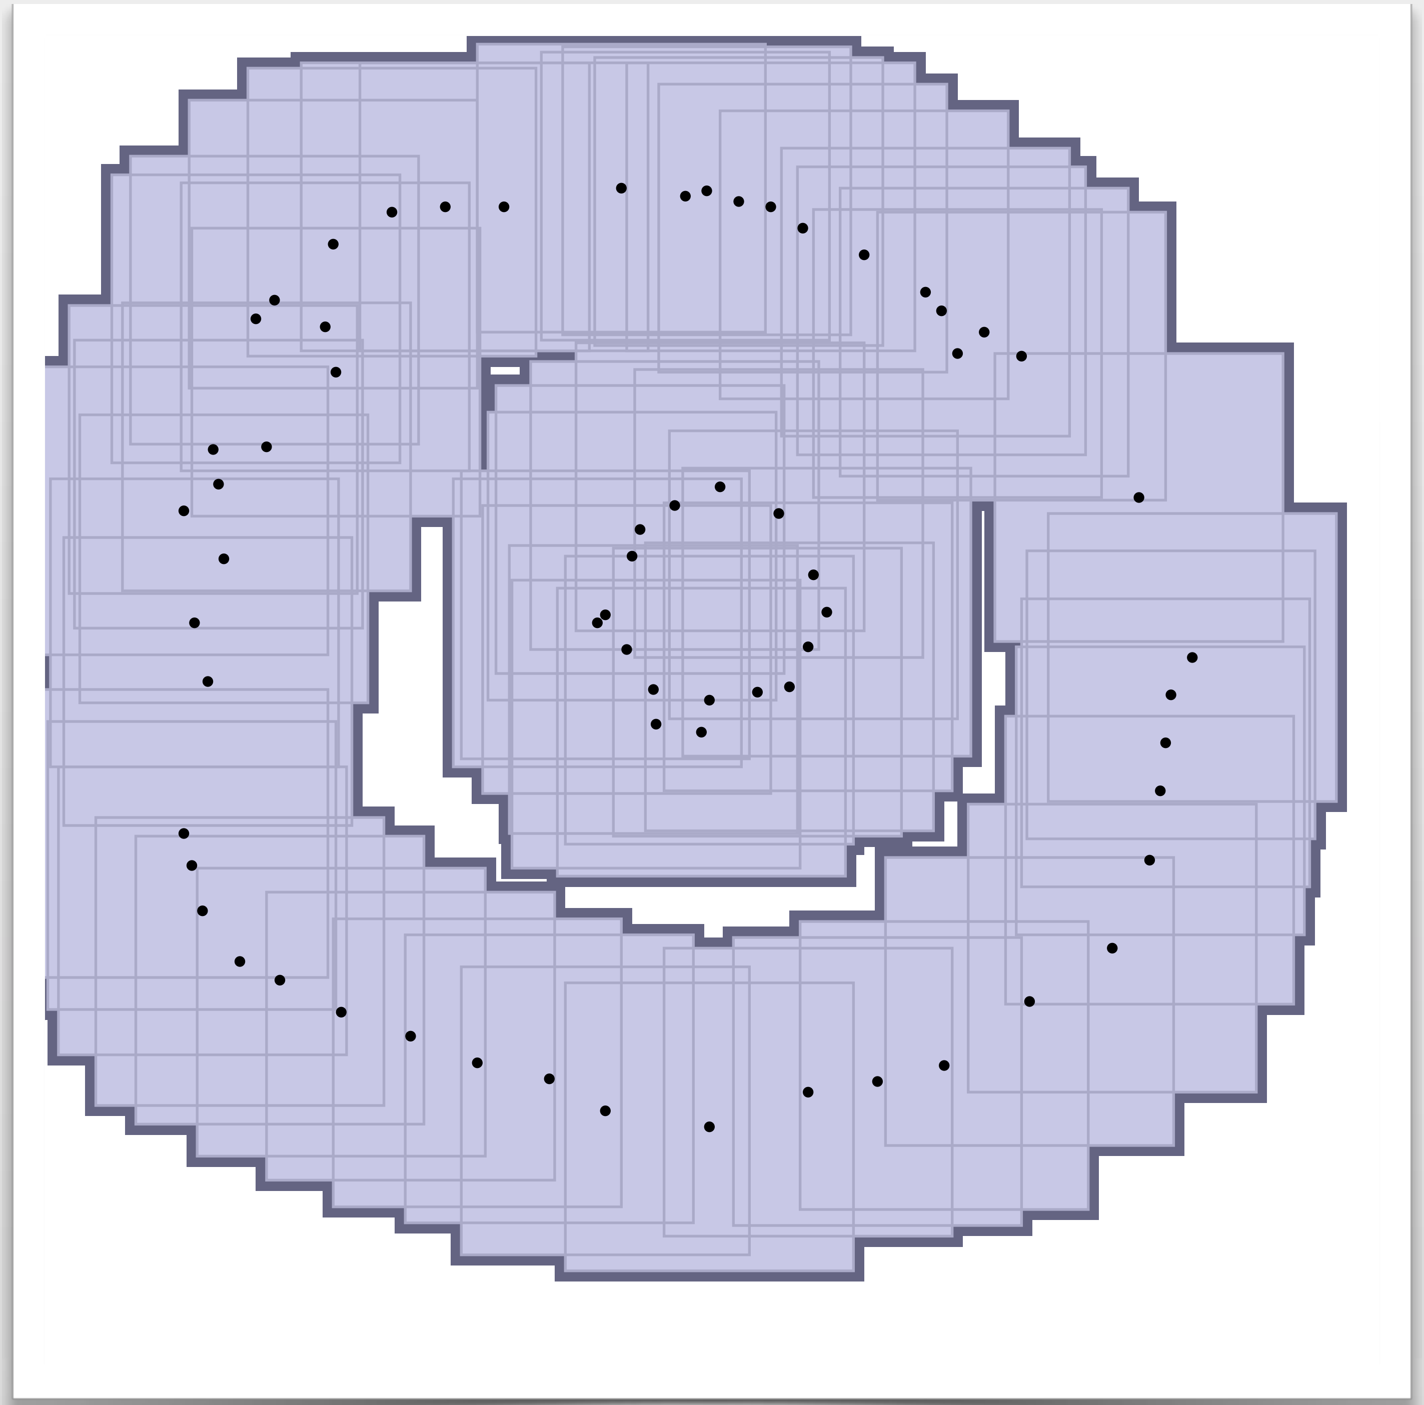

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

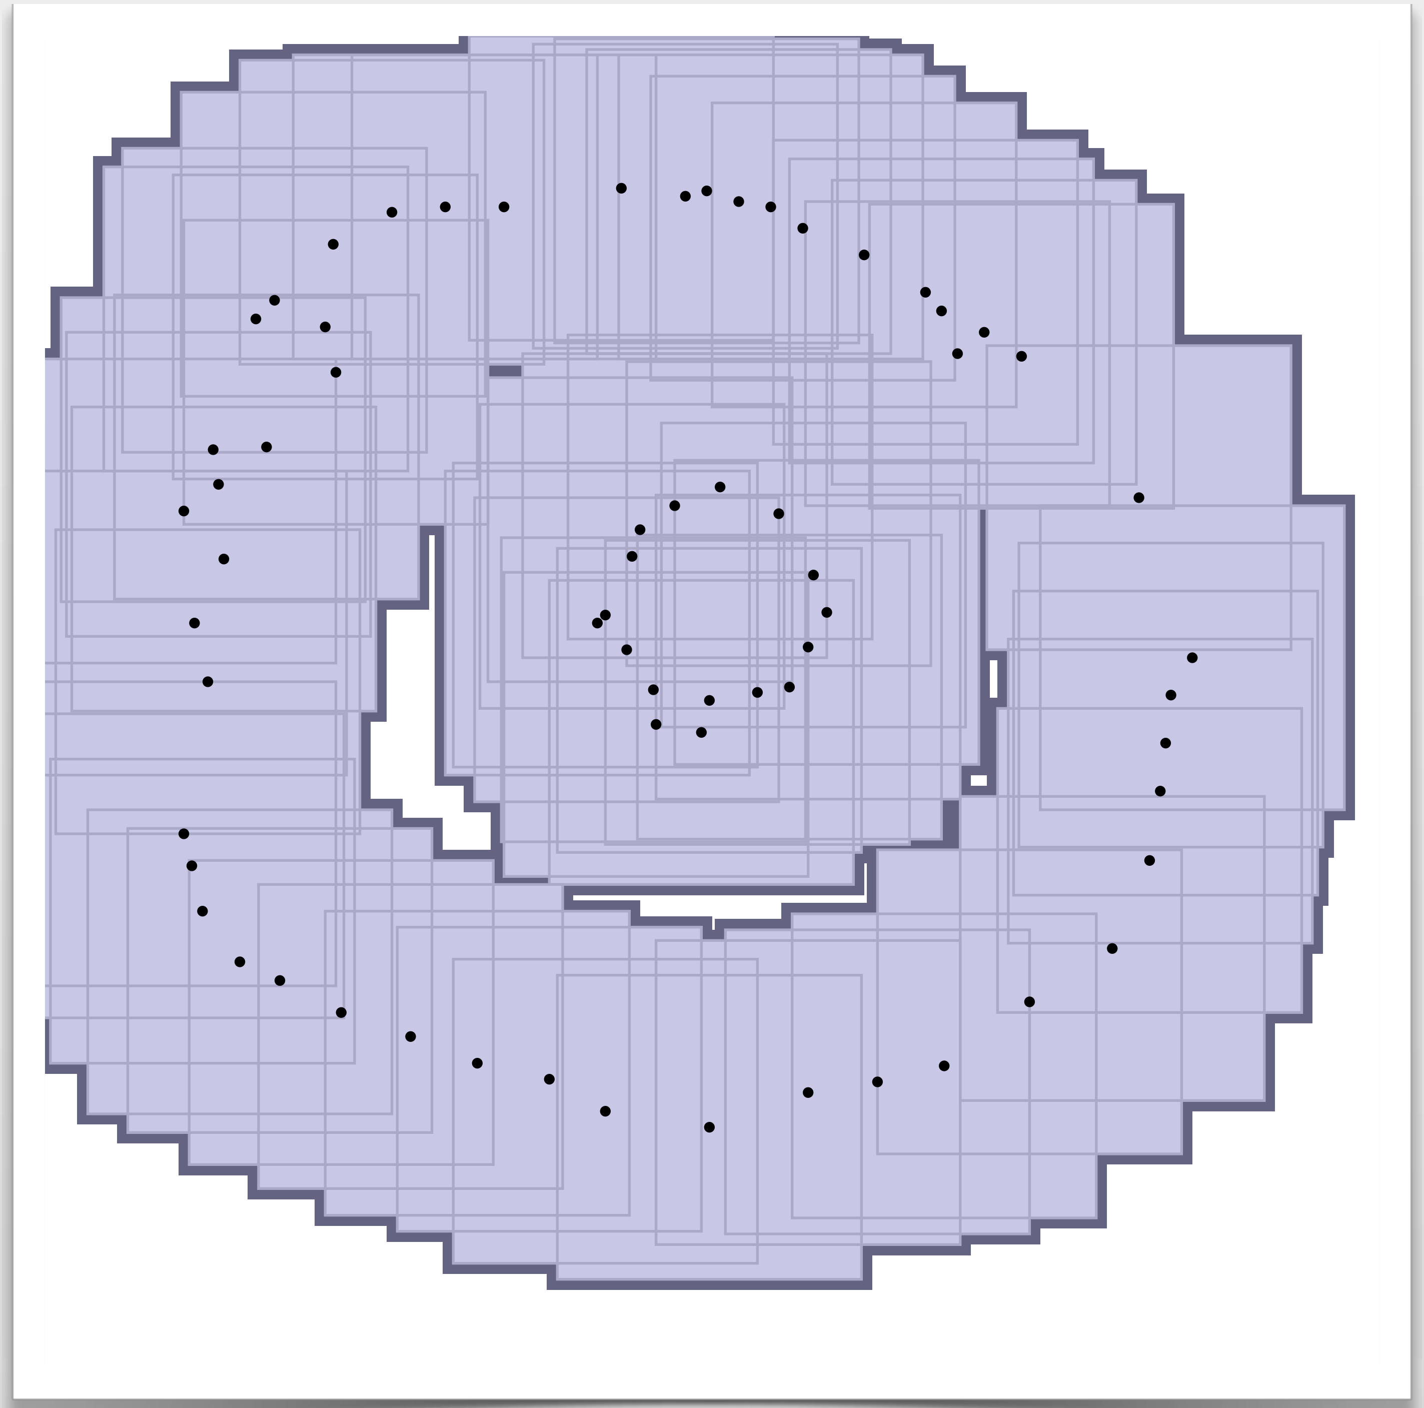

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

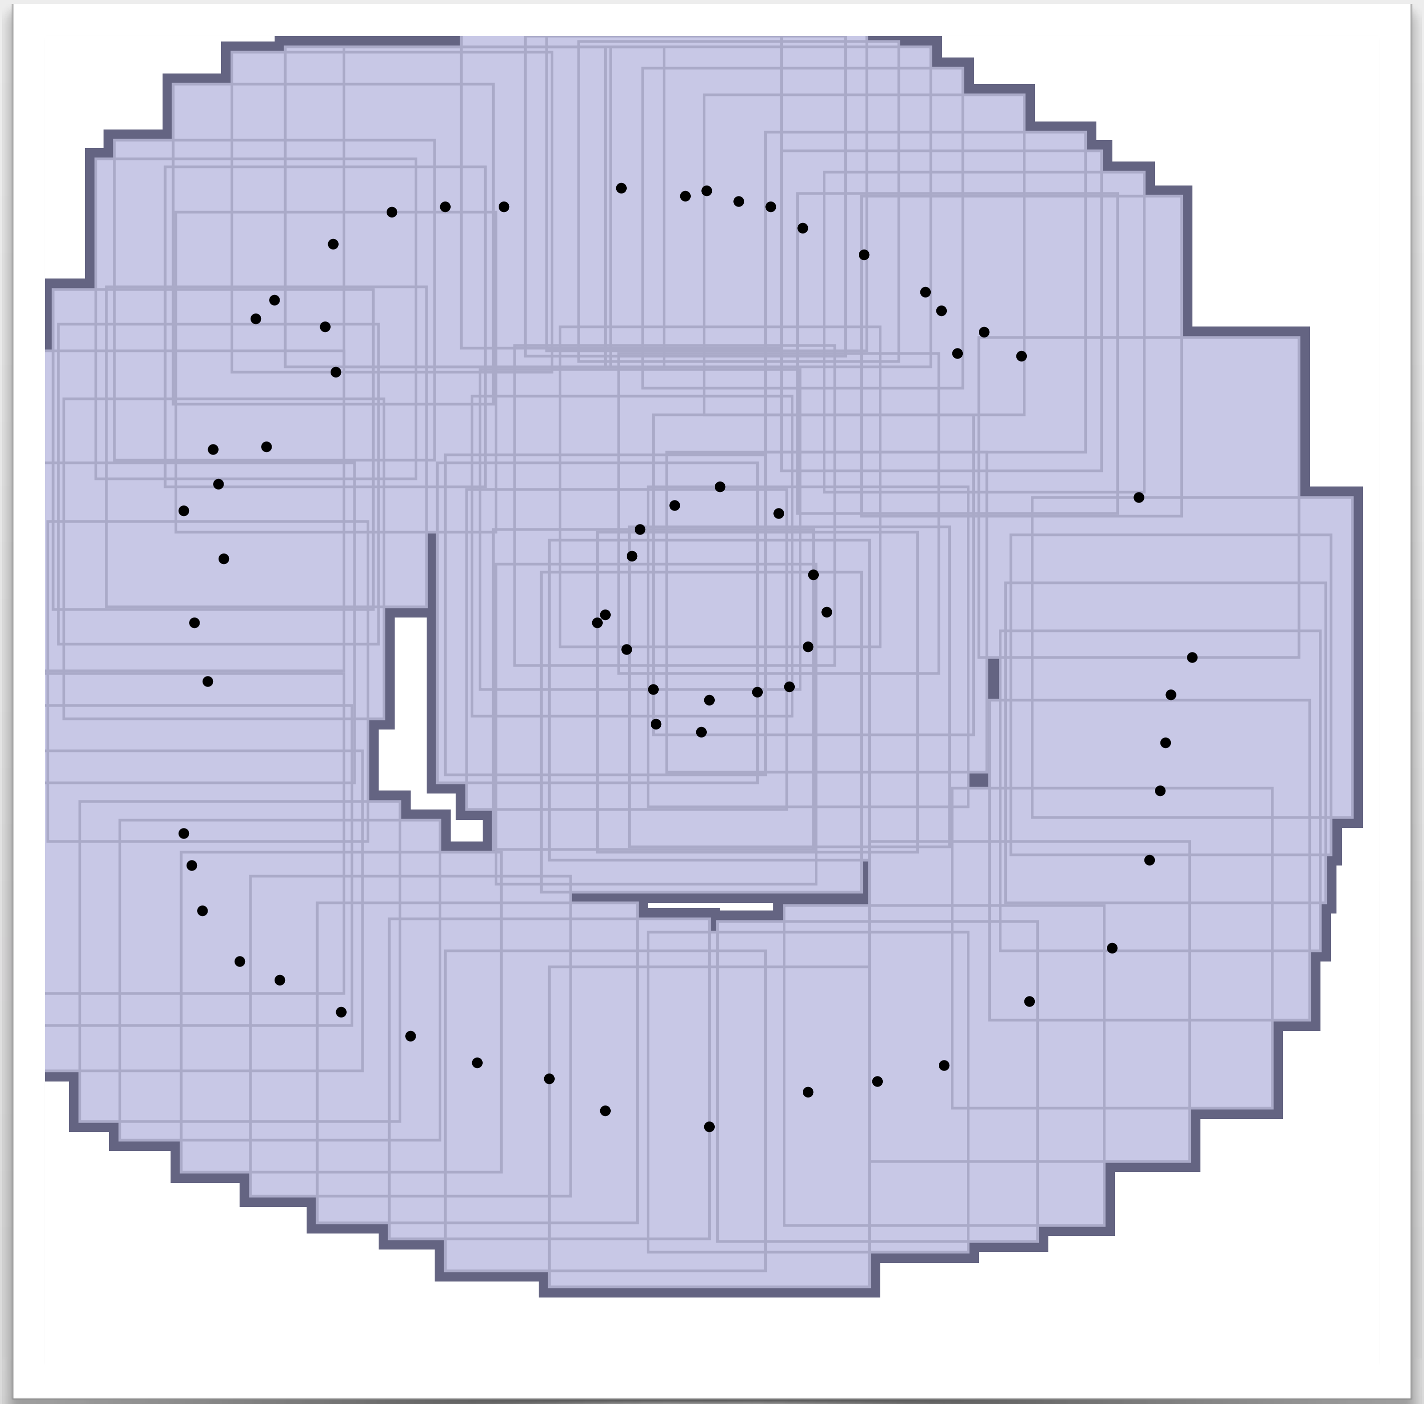

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

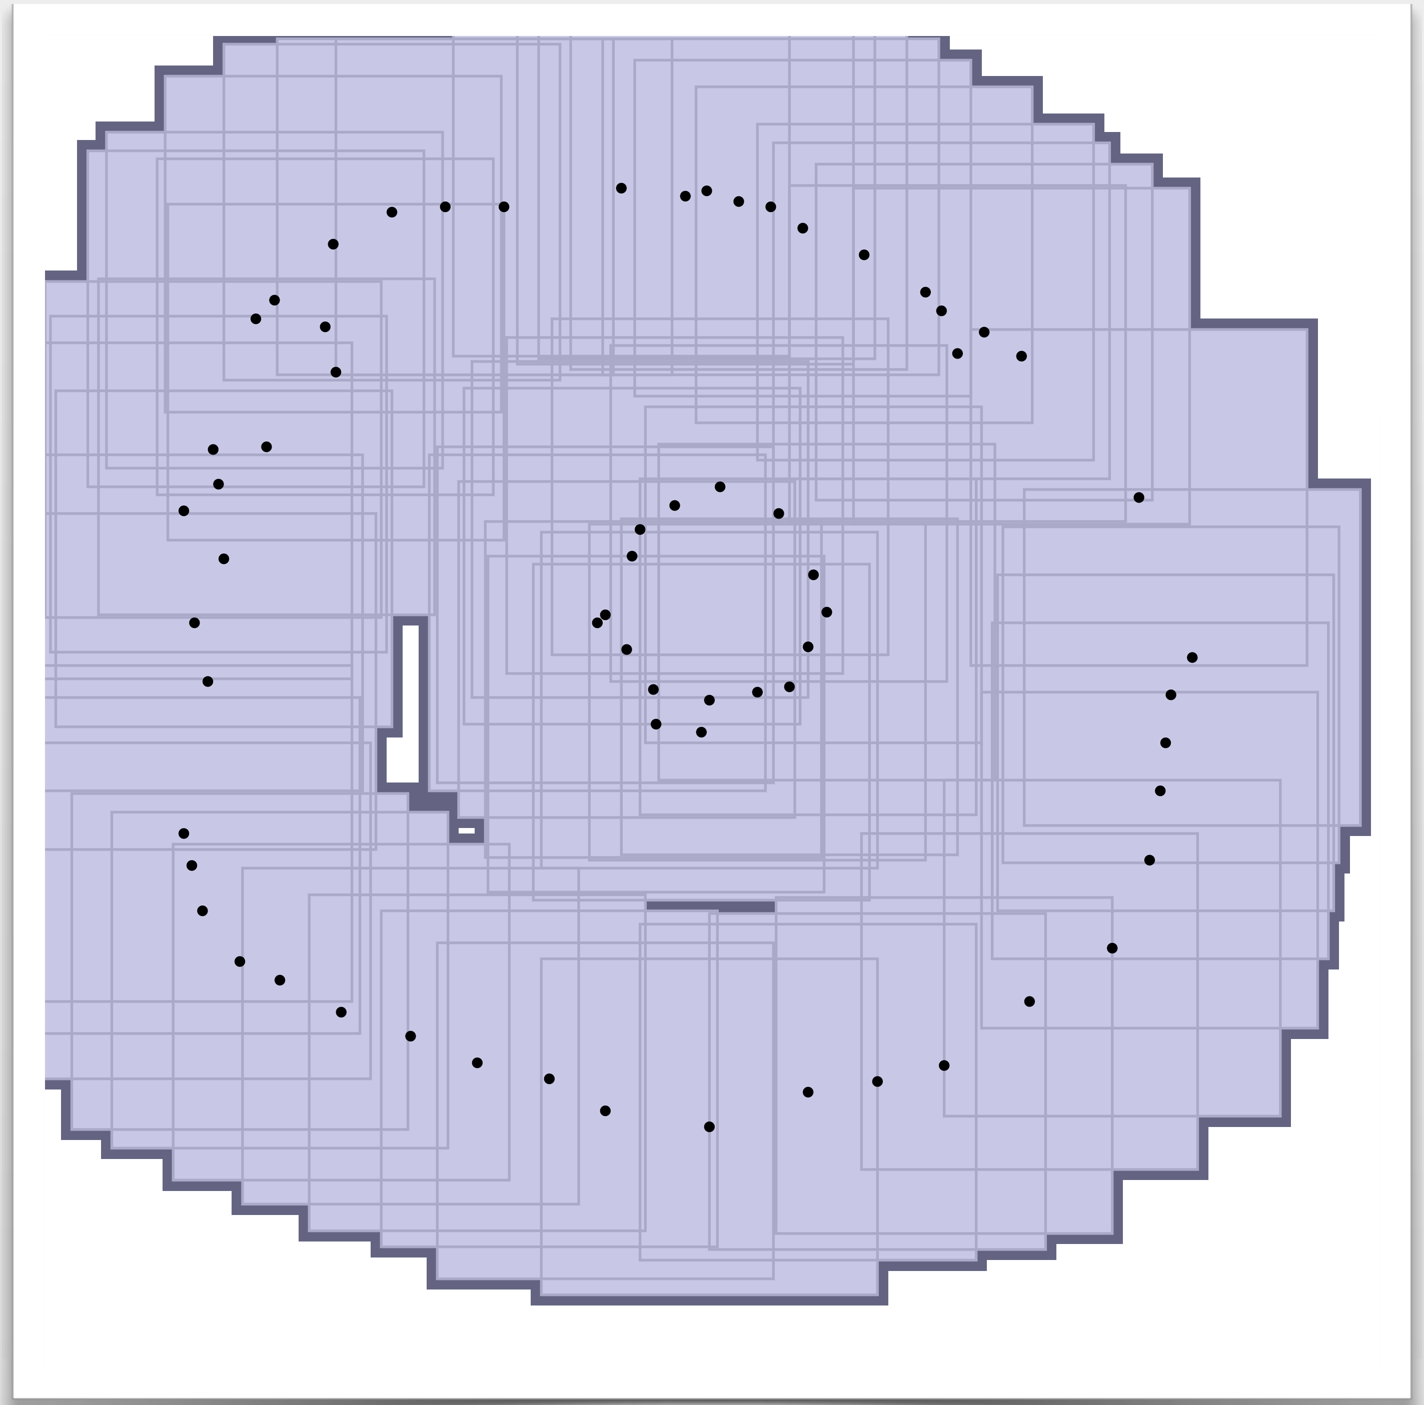

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

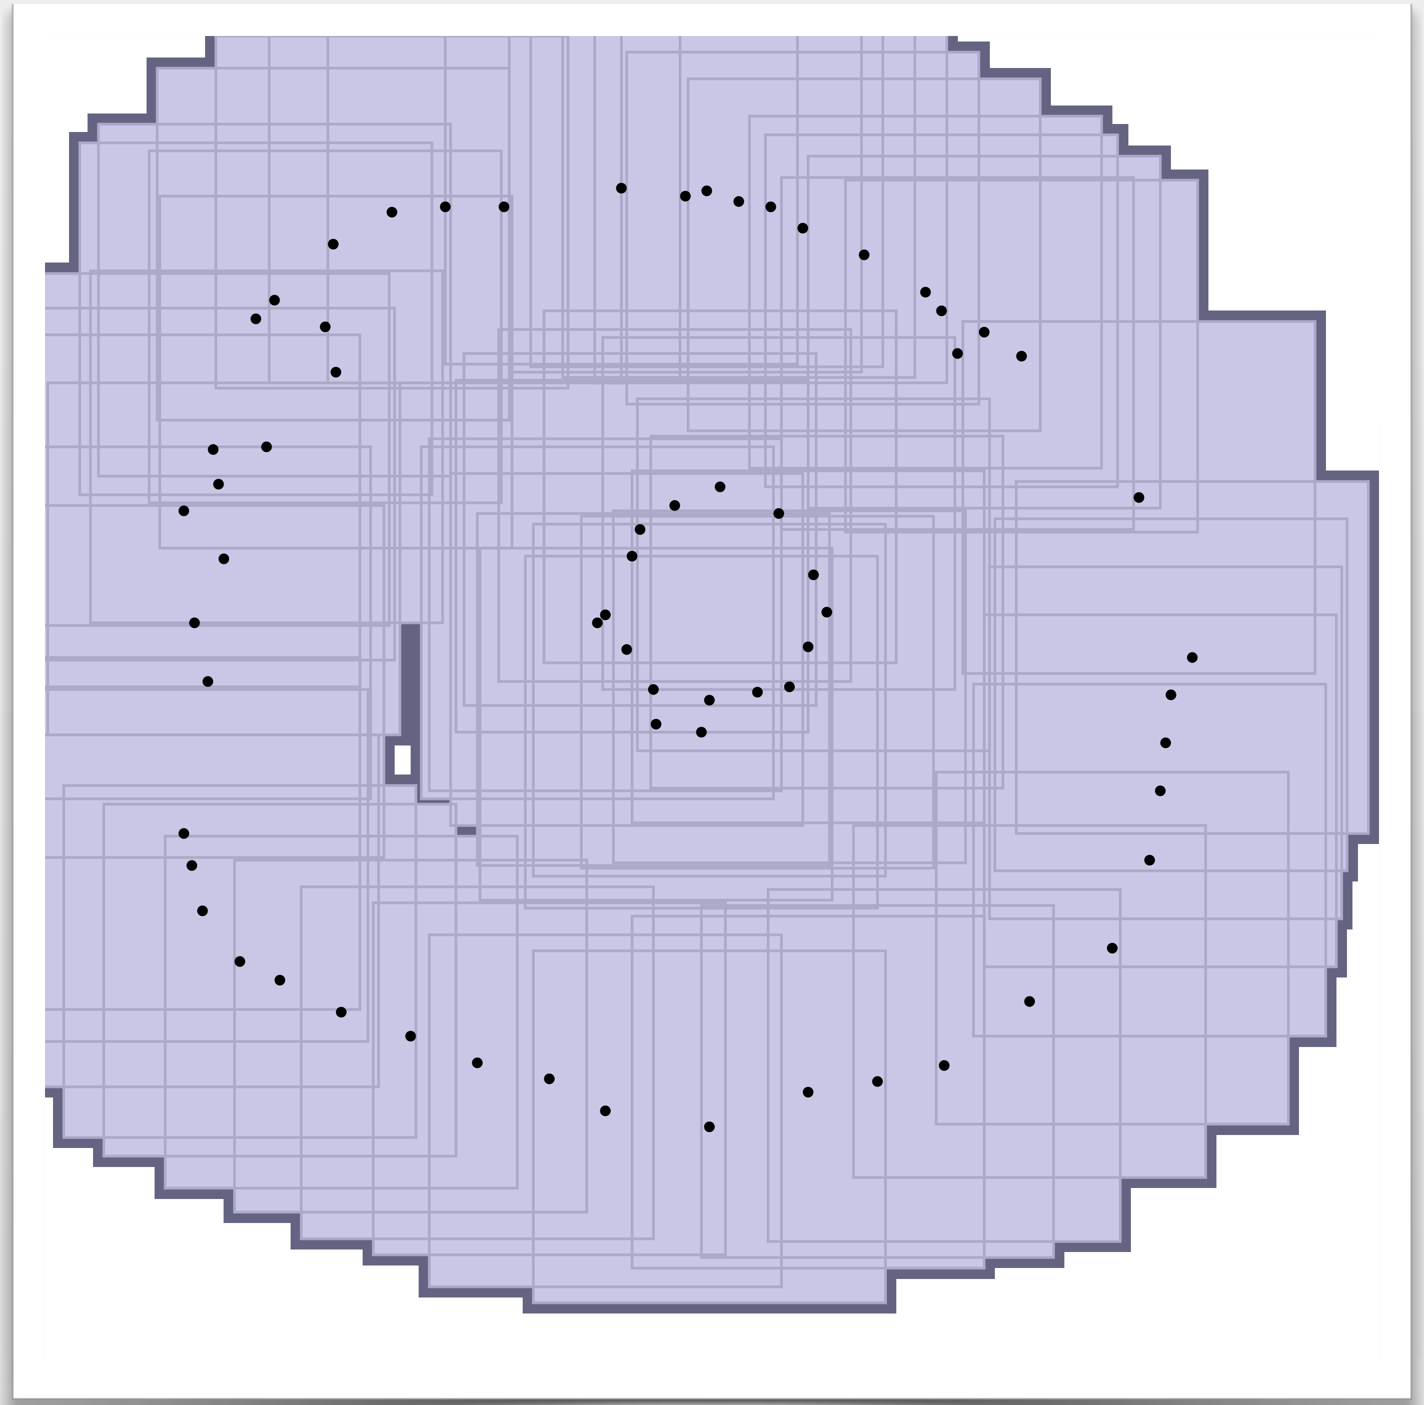

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

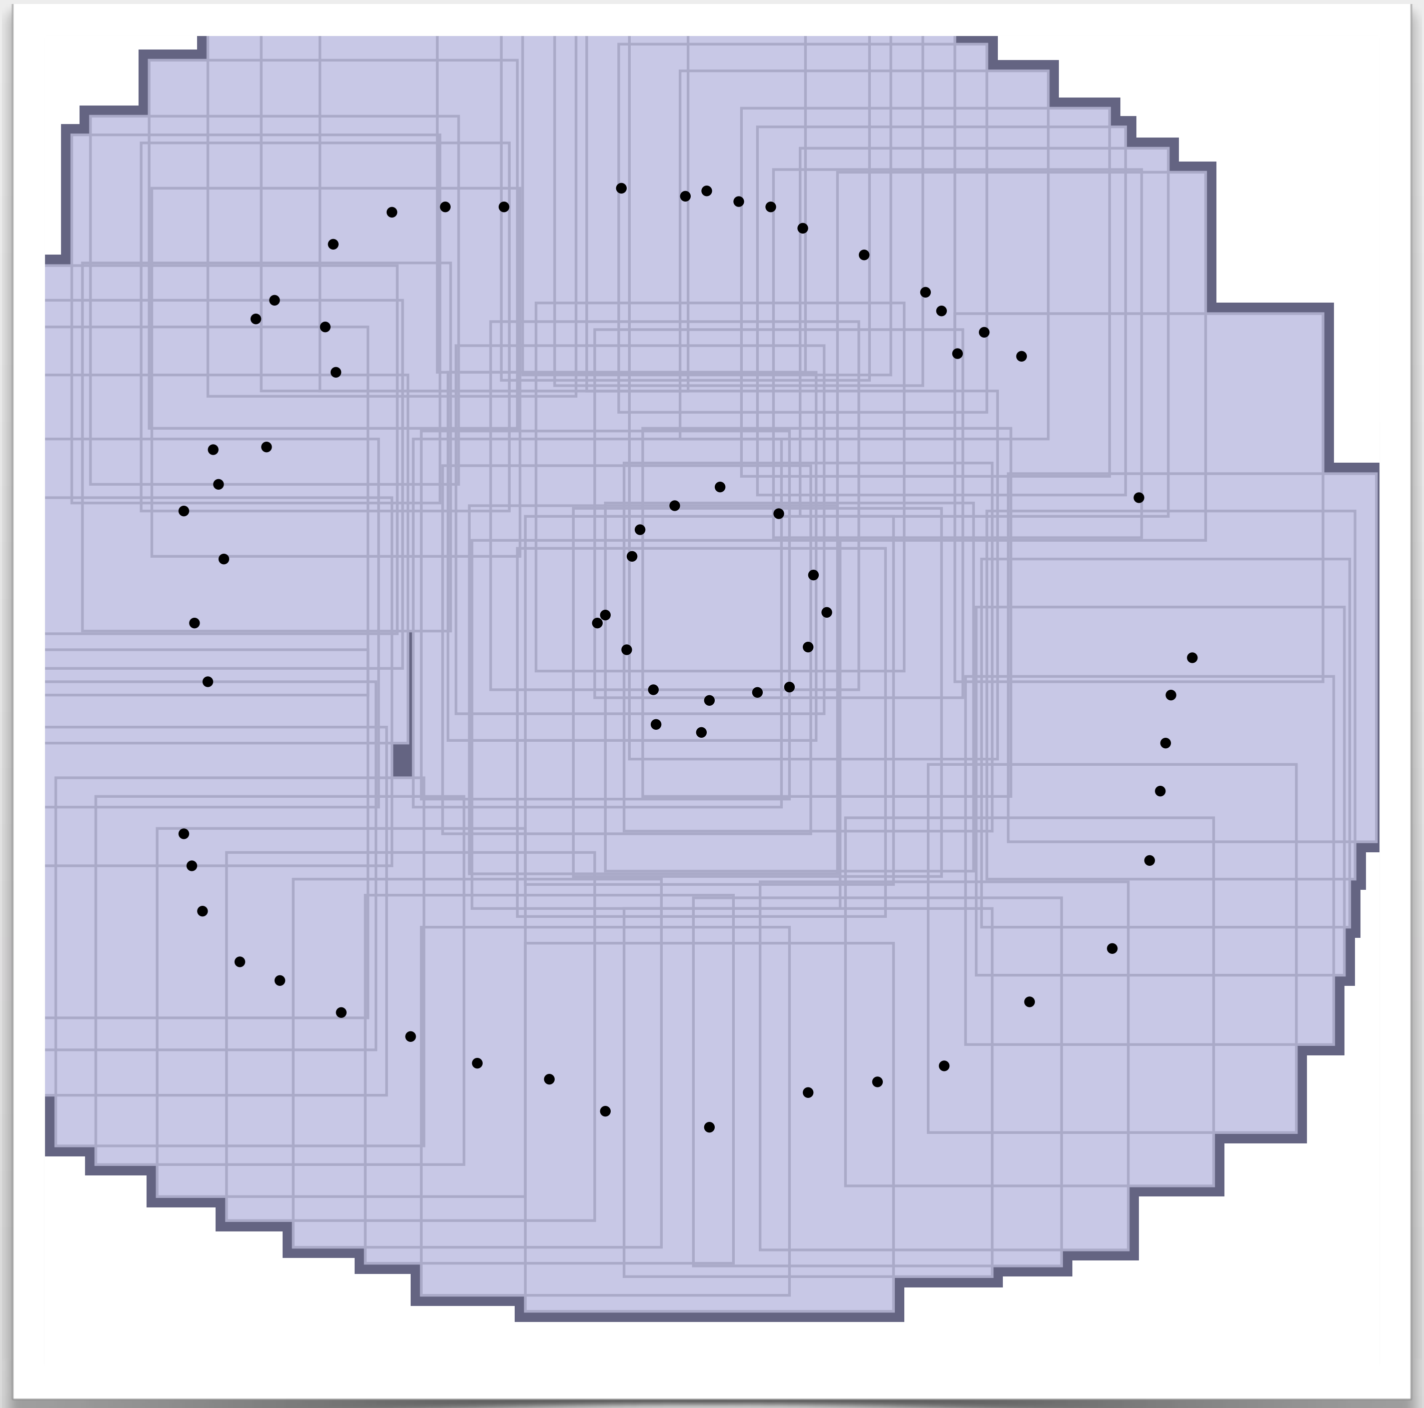

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

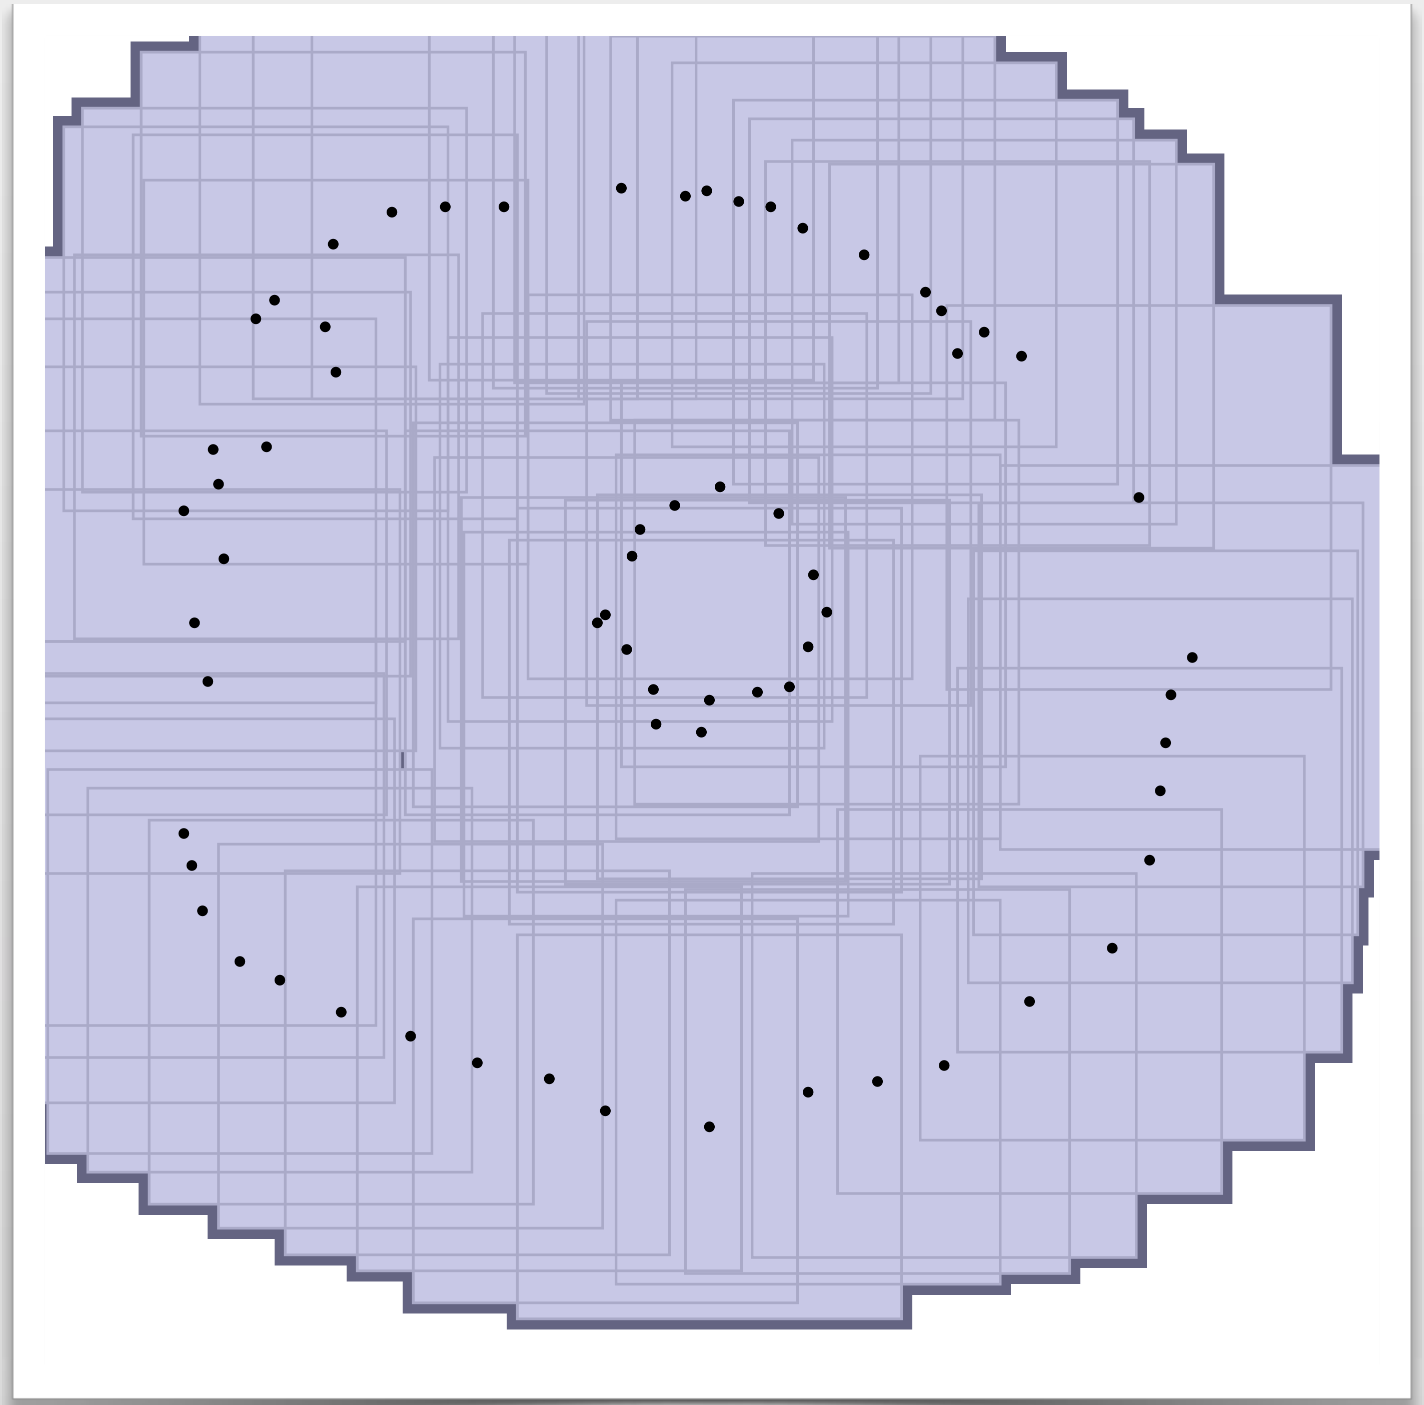

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

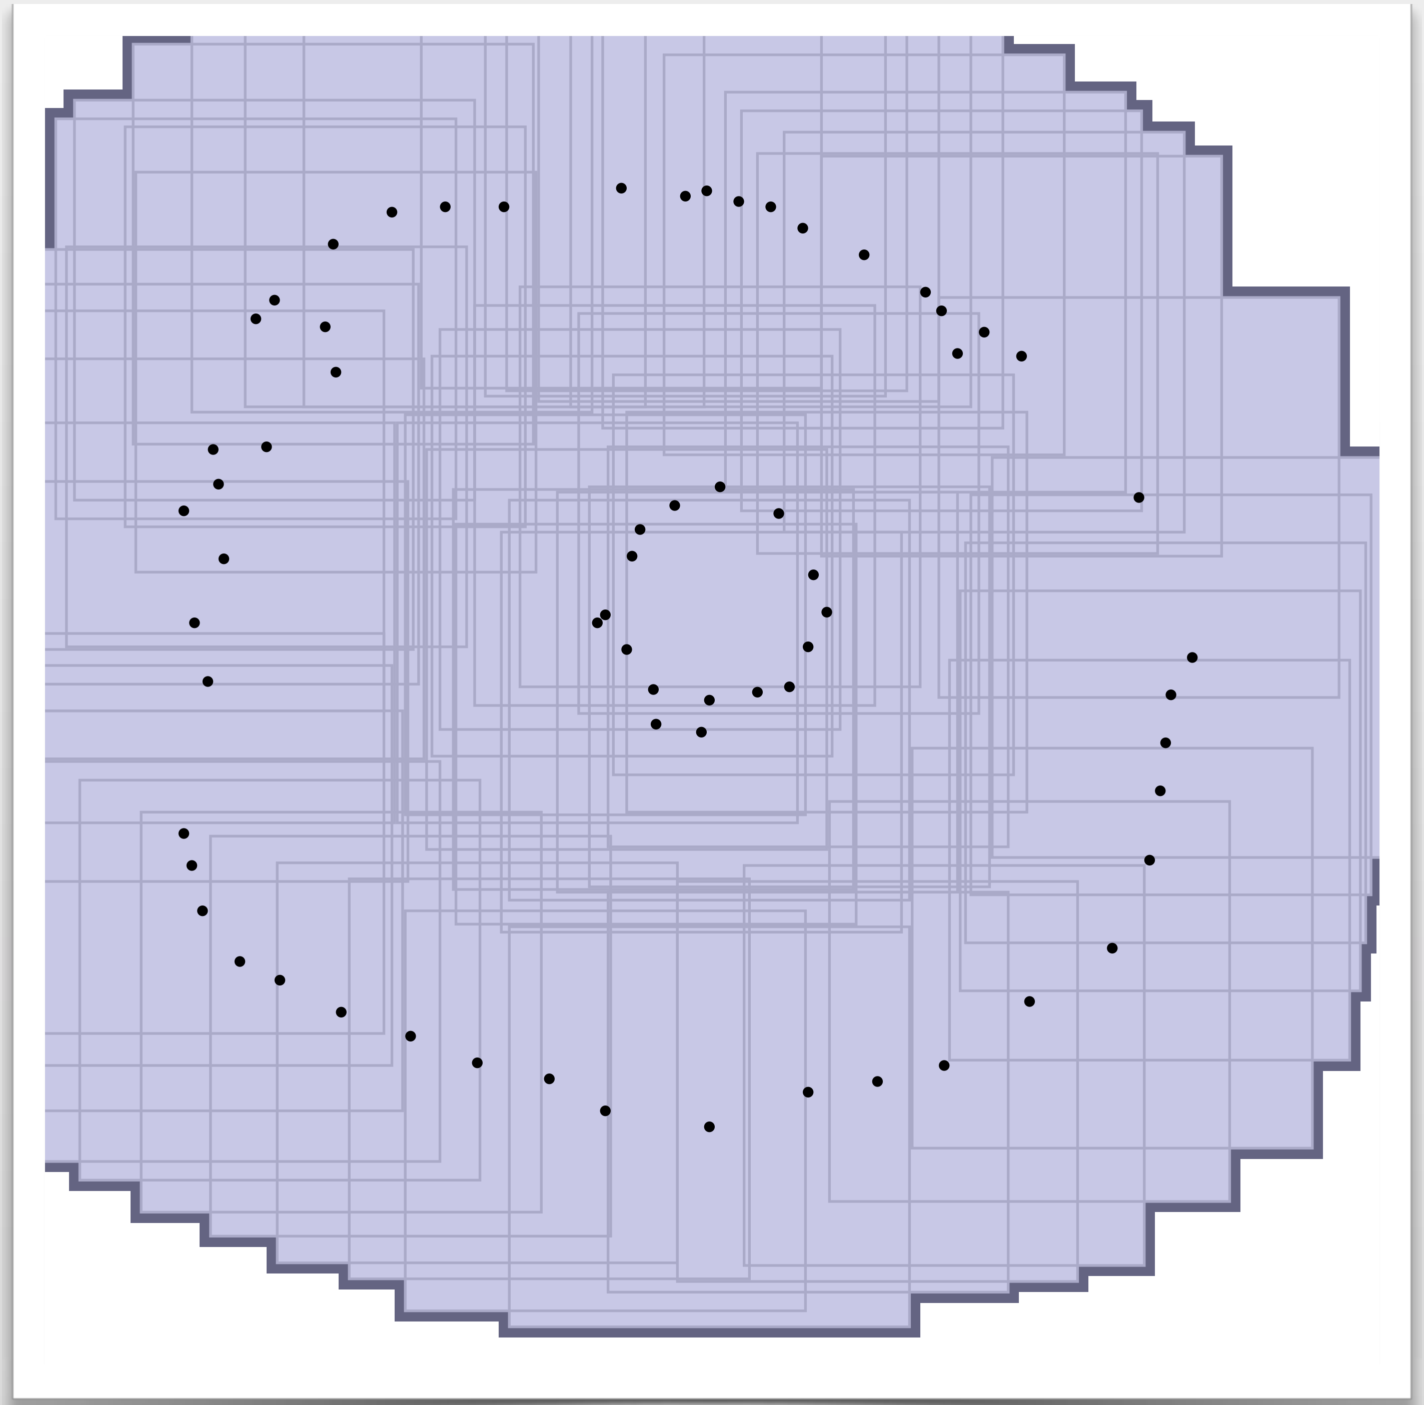

The Vietoris-Rips complex is a nerve of boxes.

Embed the input metric in $\mathbb{R}^n$ with the L_∞ norm.

In L_∞, the Rips complex is the same as the Čech complex

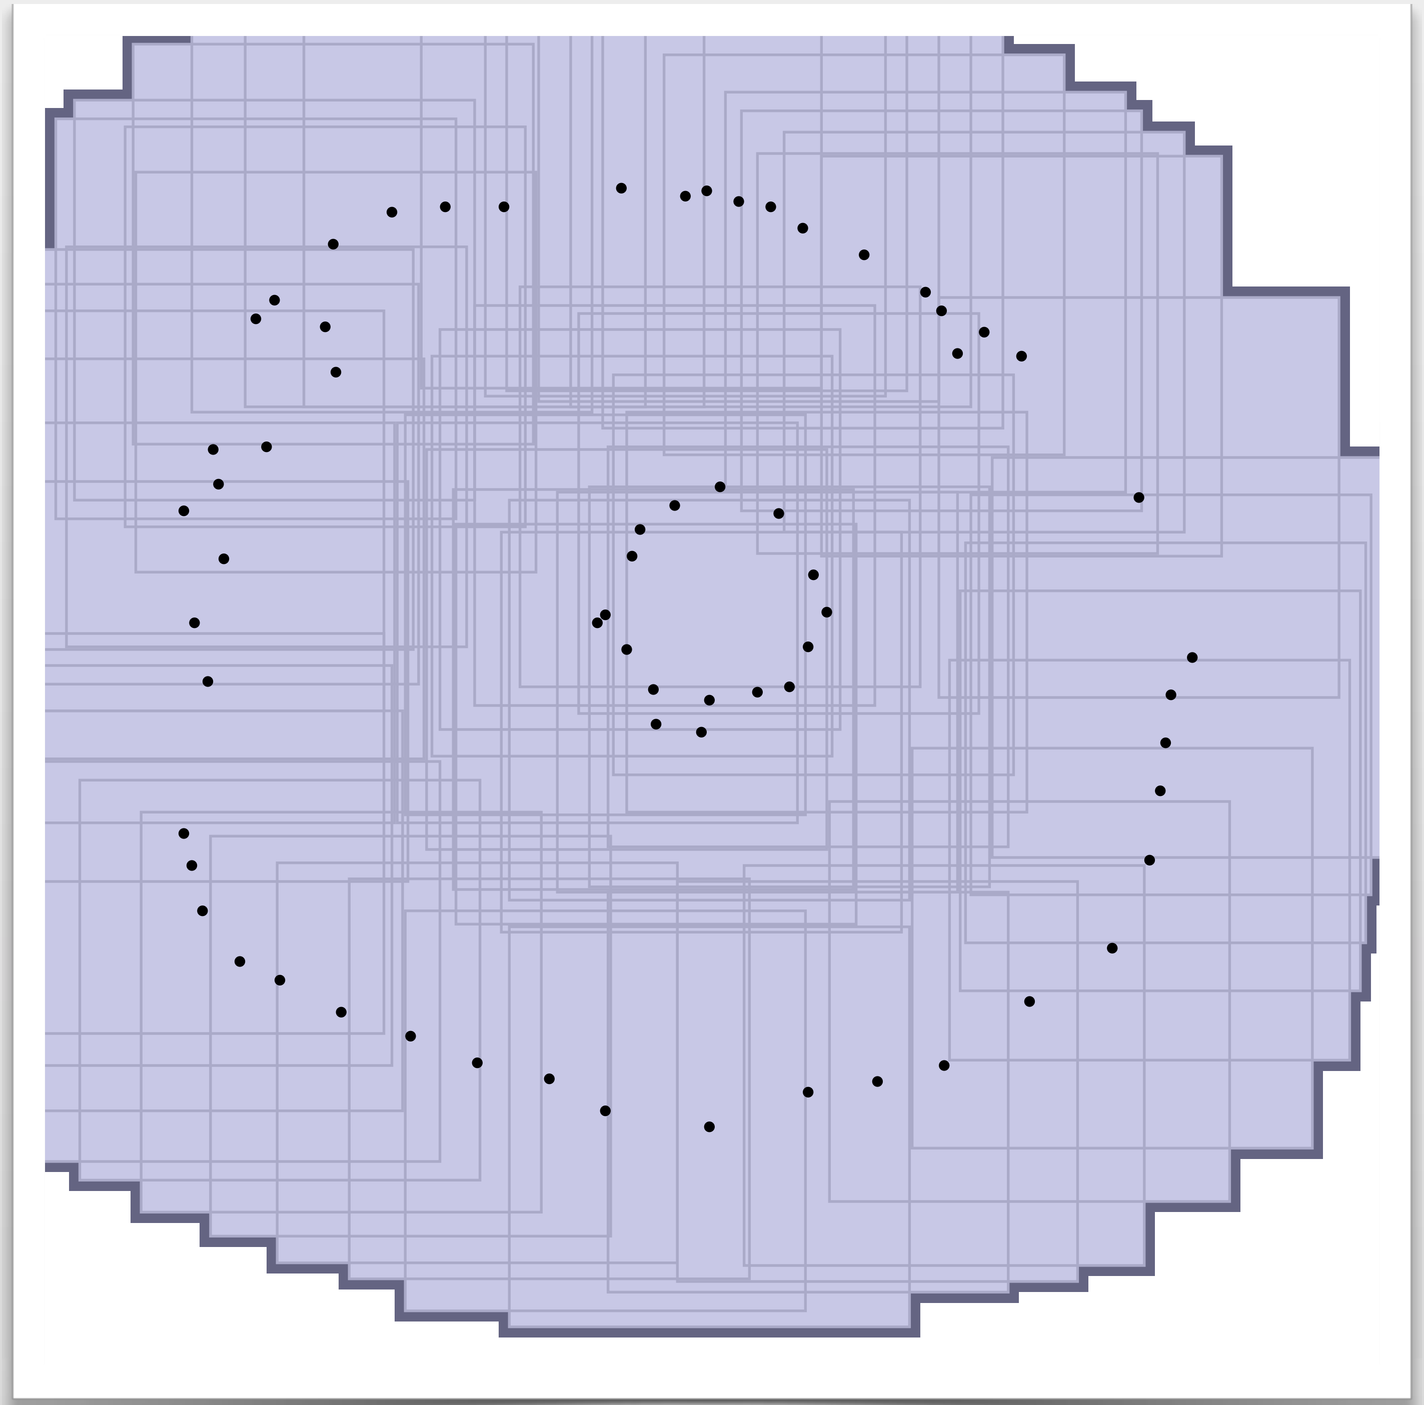

Some related work on sparse filtrations.

Some related work on sparse filtrations.

Previous Approaches at subsampling:

Some related work on sparse filtrations.

Previous Approaches at subsampling:
Witness Complexes [dSC04, BGO07]

Some related work on sparse filtrations.

Previous Approaches at subsampling:
- Witness Complexes [dSC04, BGO07]
- Persistence-based Reconstruction [CO08]

Some related work on sparse filtrations.

Previous Approaches at subsampling:
  Witness Complexes [dSC04, BGO07]
  Persistence-based Reconstruction [CO08]
Other methods:

Some related work on sparse filtrations.

Previous Approaches at subsampling:
- Witness Complexes [dSC04, BGO07]
- Persistence-based Reconstruction [CO08]

Other methods:
- Meshes in Euclidean Space [HMSO10]

Some related work on sparse filtrations.

Previous Approaches at subsampling:
- Witness Complexes [dSC04, BGO07]
- Persistence-based Reconstruction [CO08]

Other methods:
- Meshes in Euclidean Space [HMSO10]
- Topological simplification [Z10, ALS11]

Key idea: Treat many close points as one point.

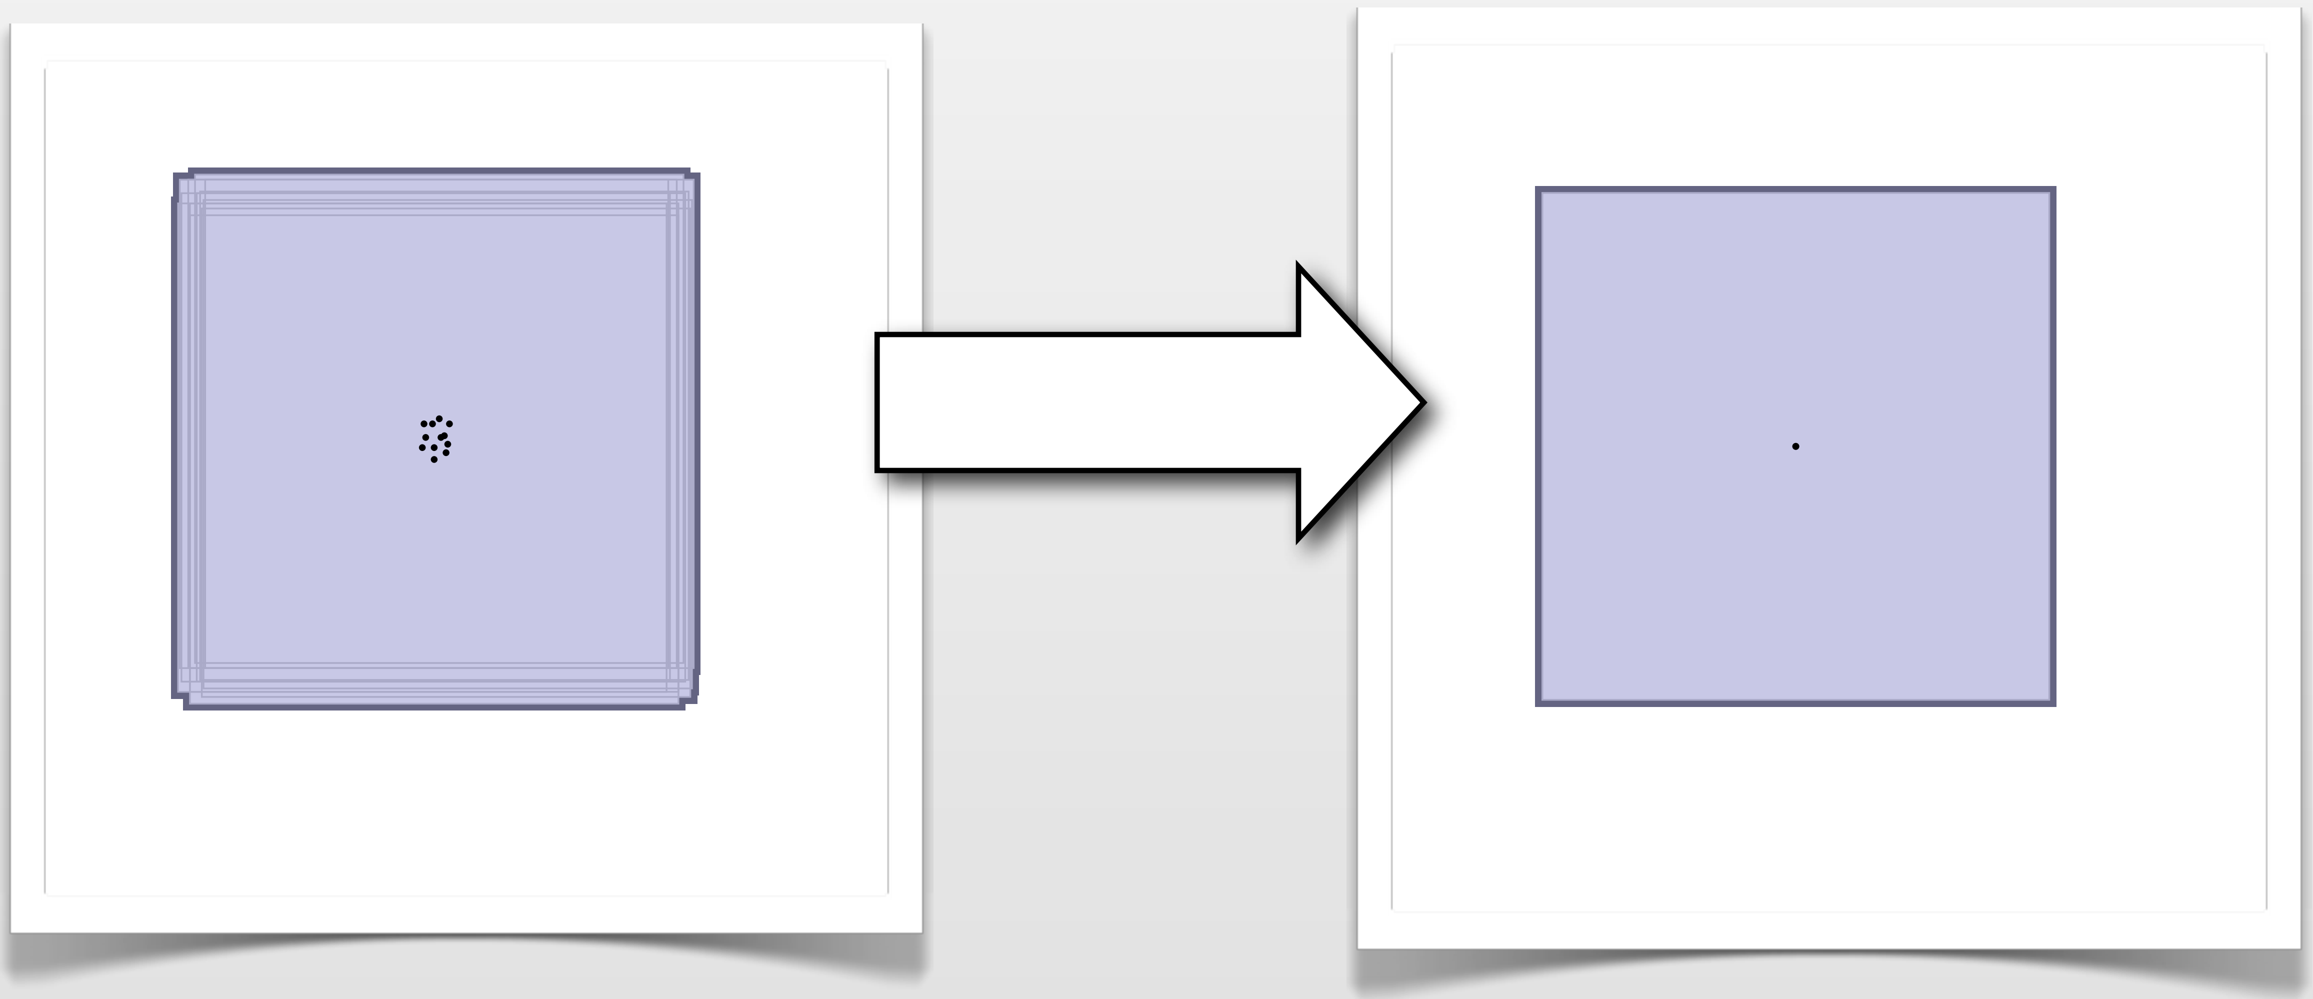

Key idea: Treat many close points as one point.

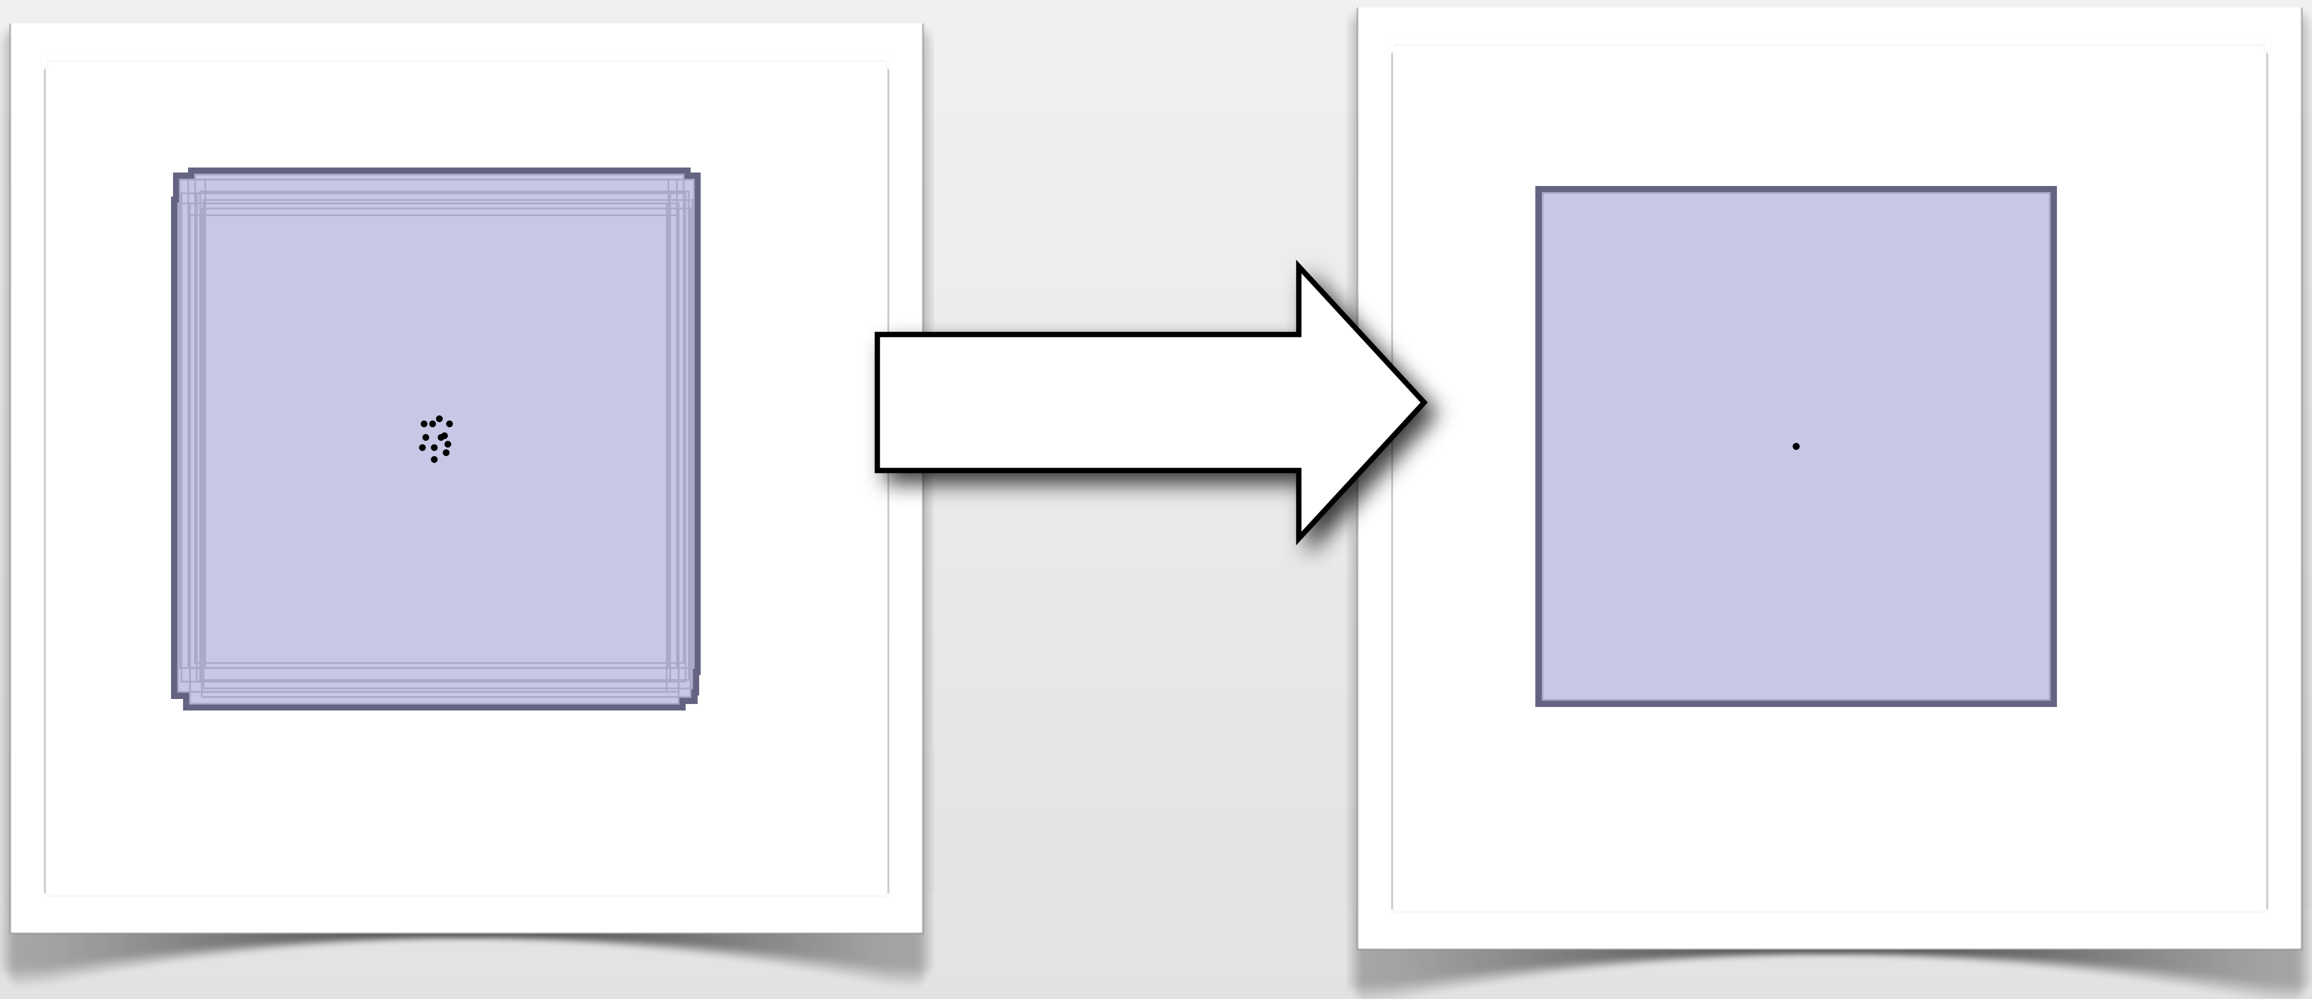

This idea is ubiquitous in computational geometry.

Key idea: Treat many close points as one point.

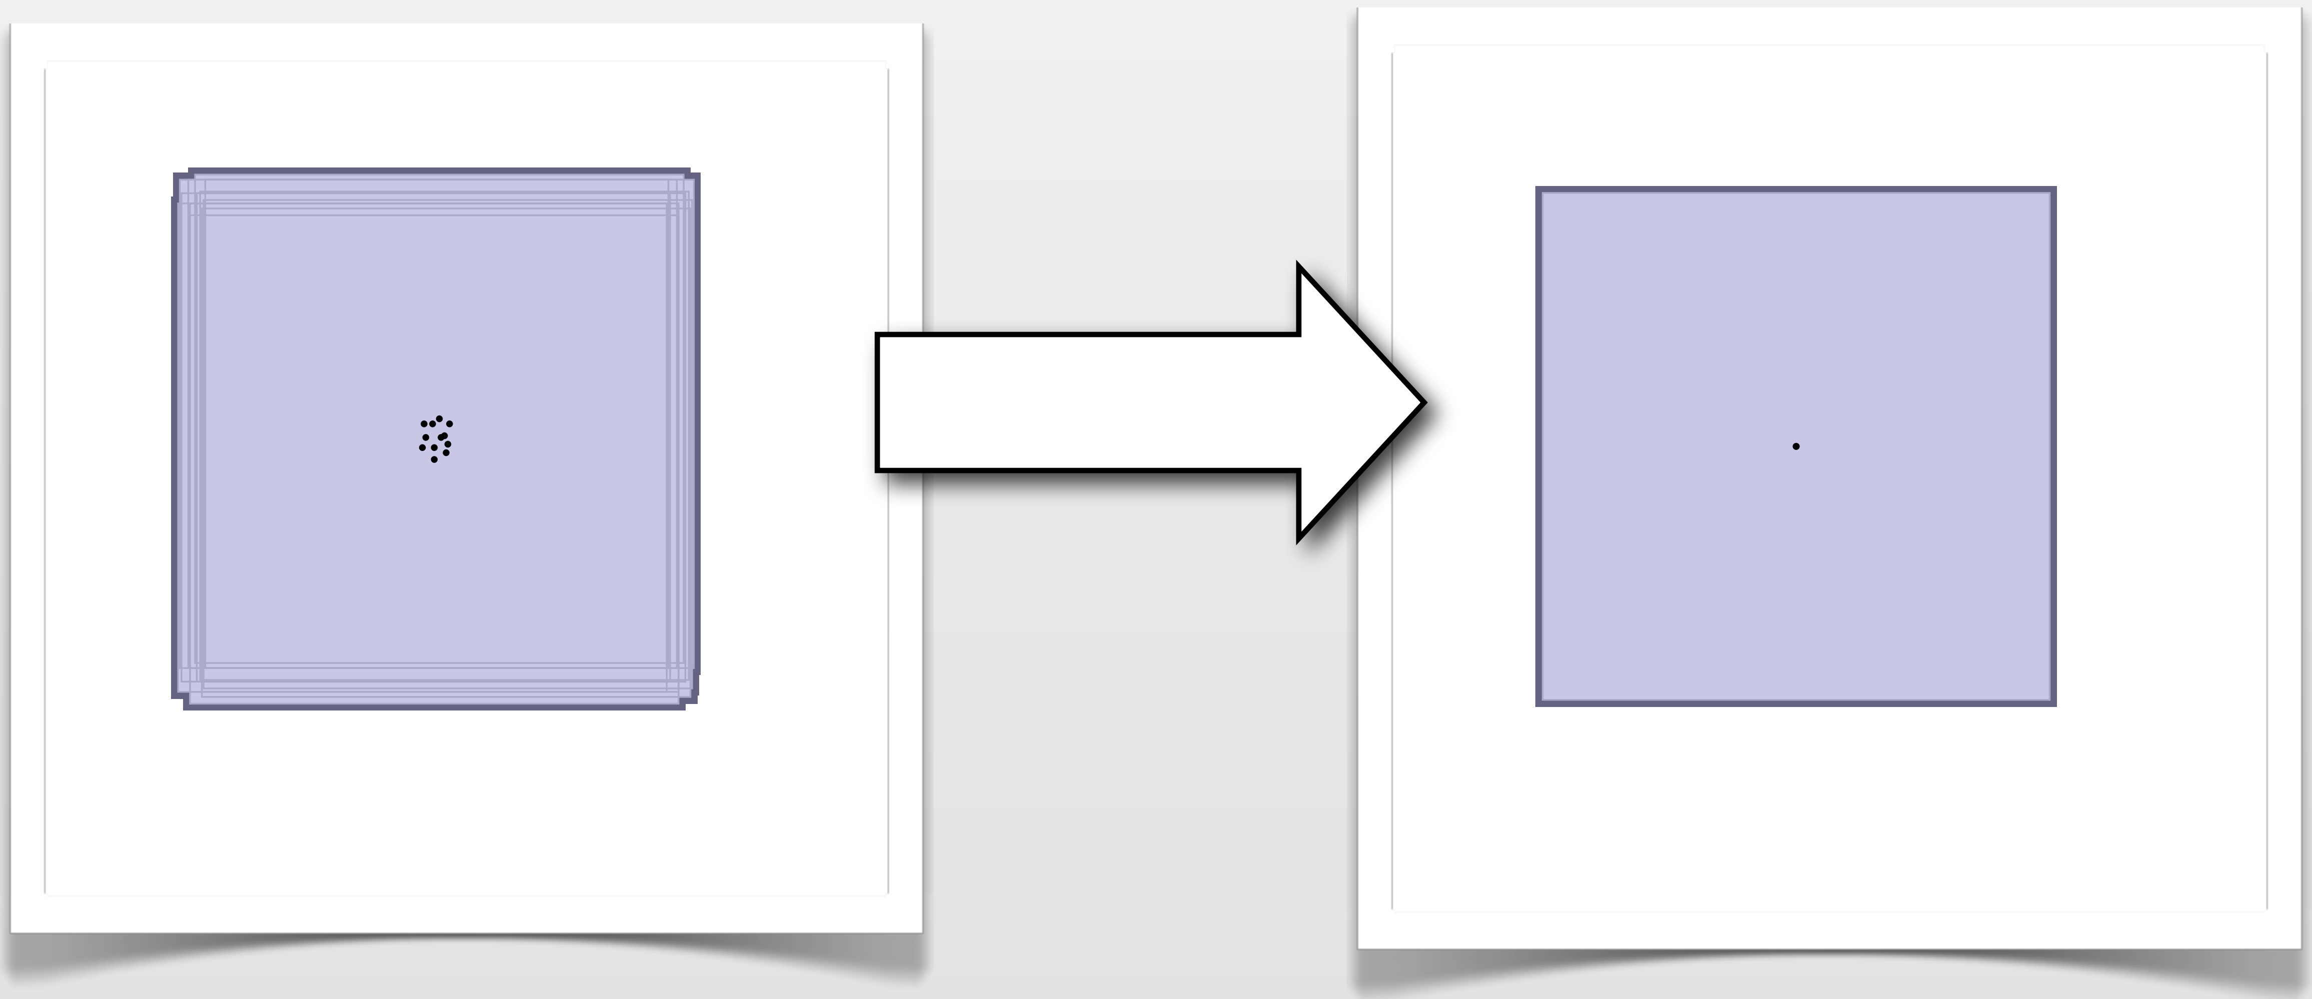

This idea is ubiquitous in computational geometry.

n-body simulation, approximate nearest neighbor search, spanners, well-separated pair decomposition,...

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

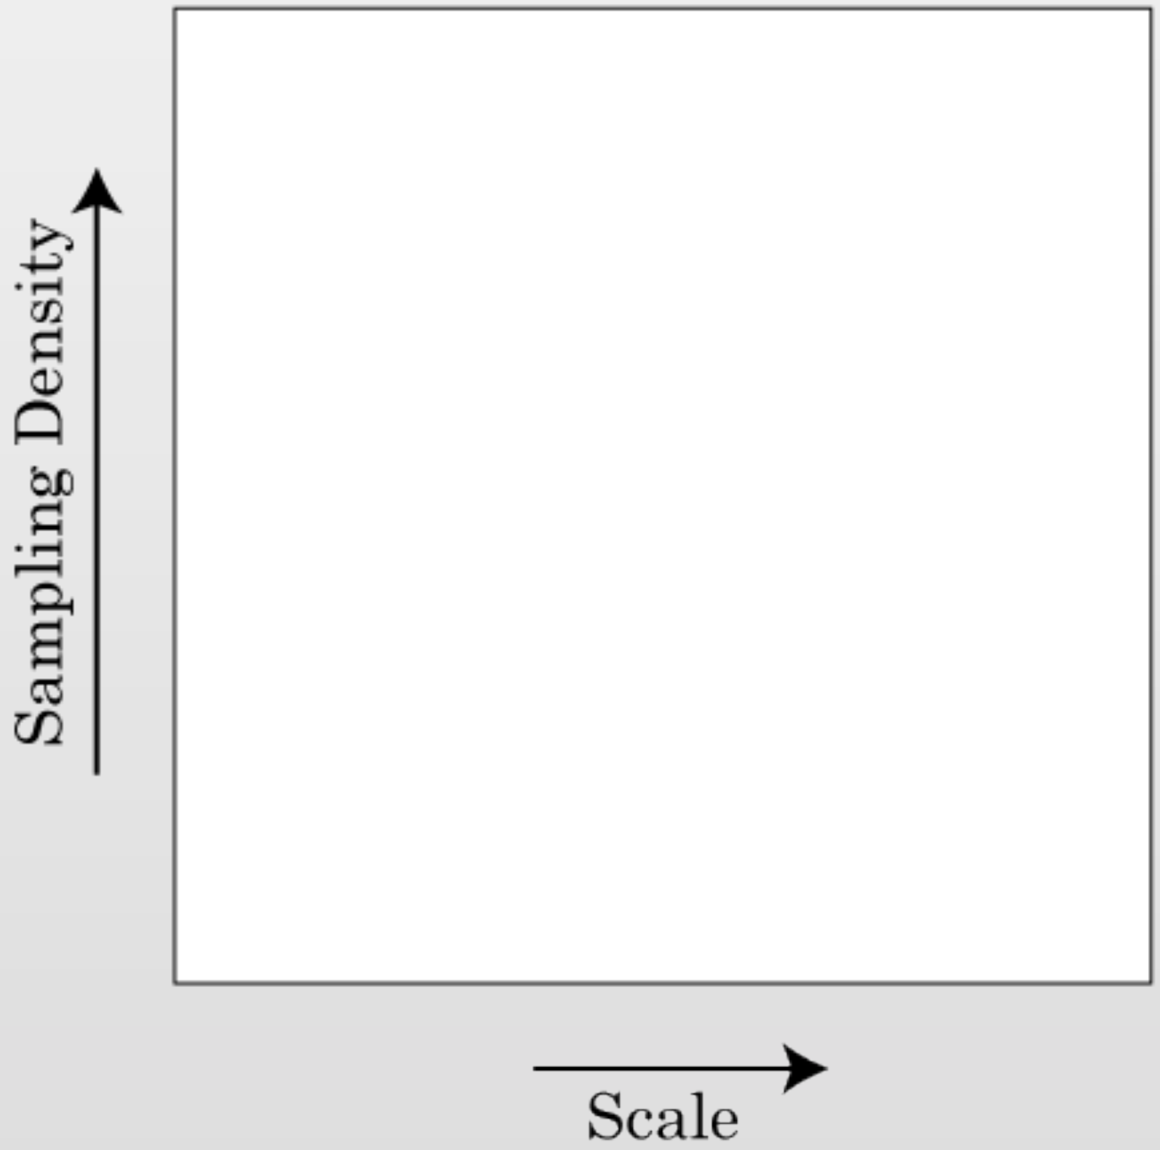

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

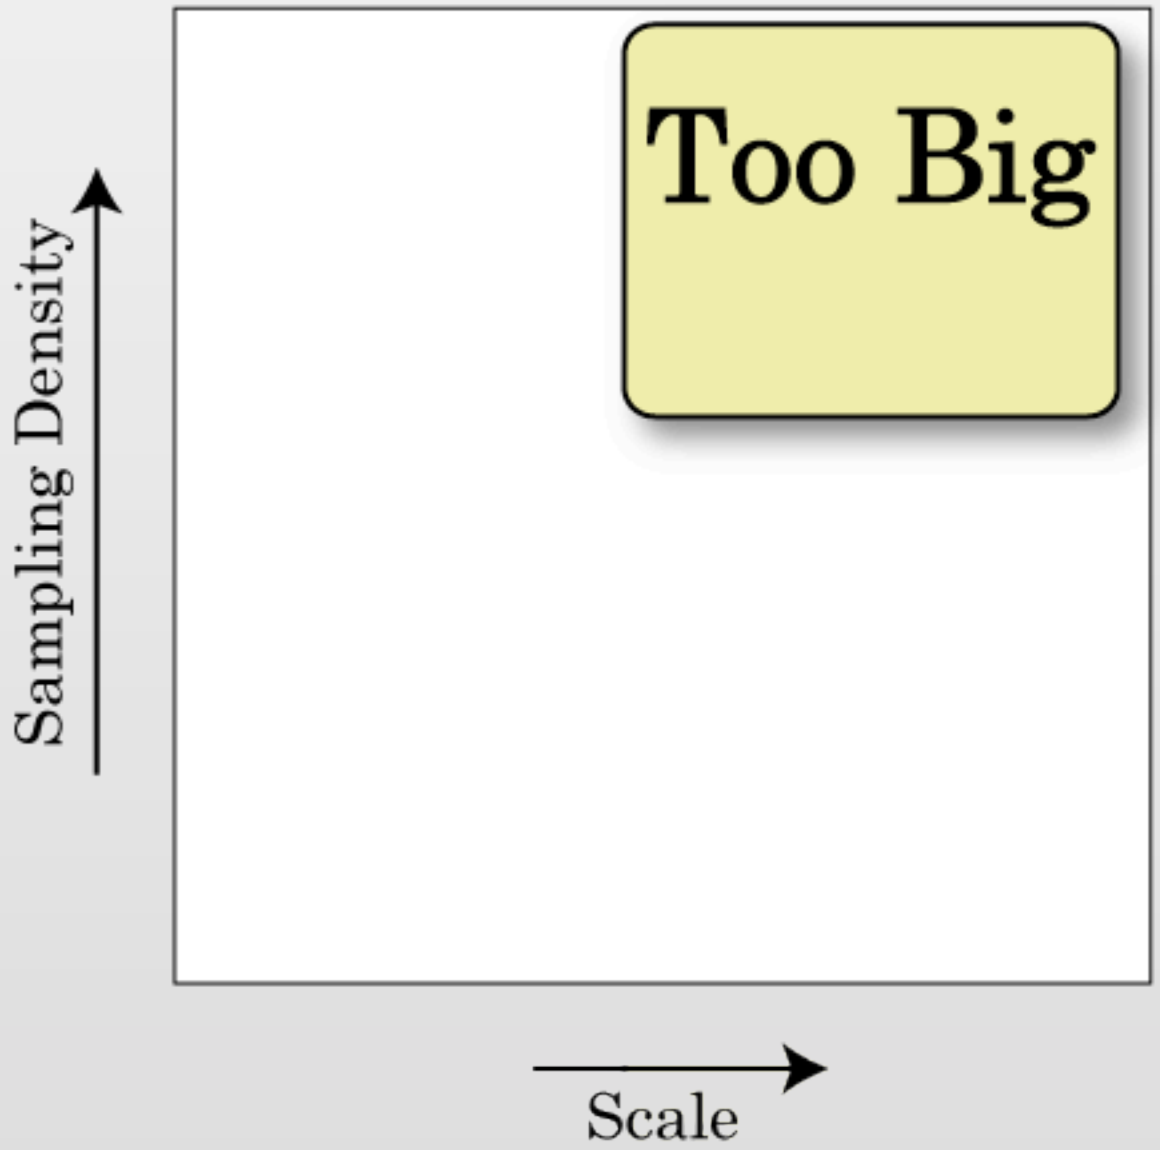

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

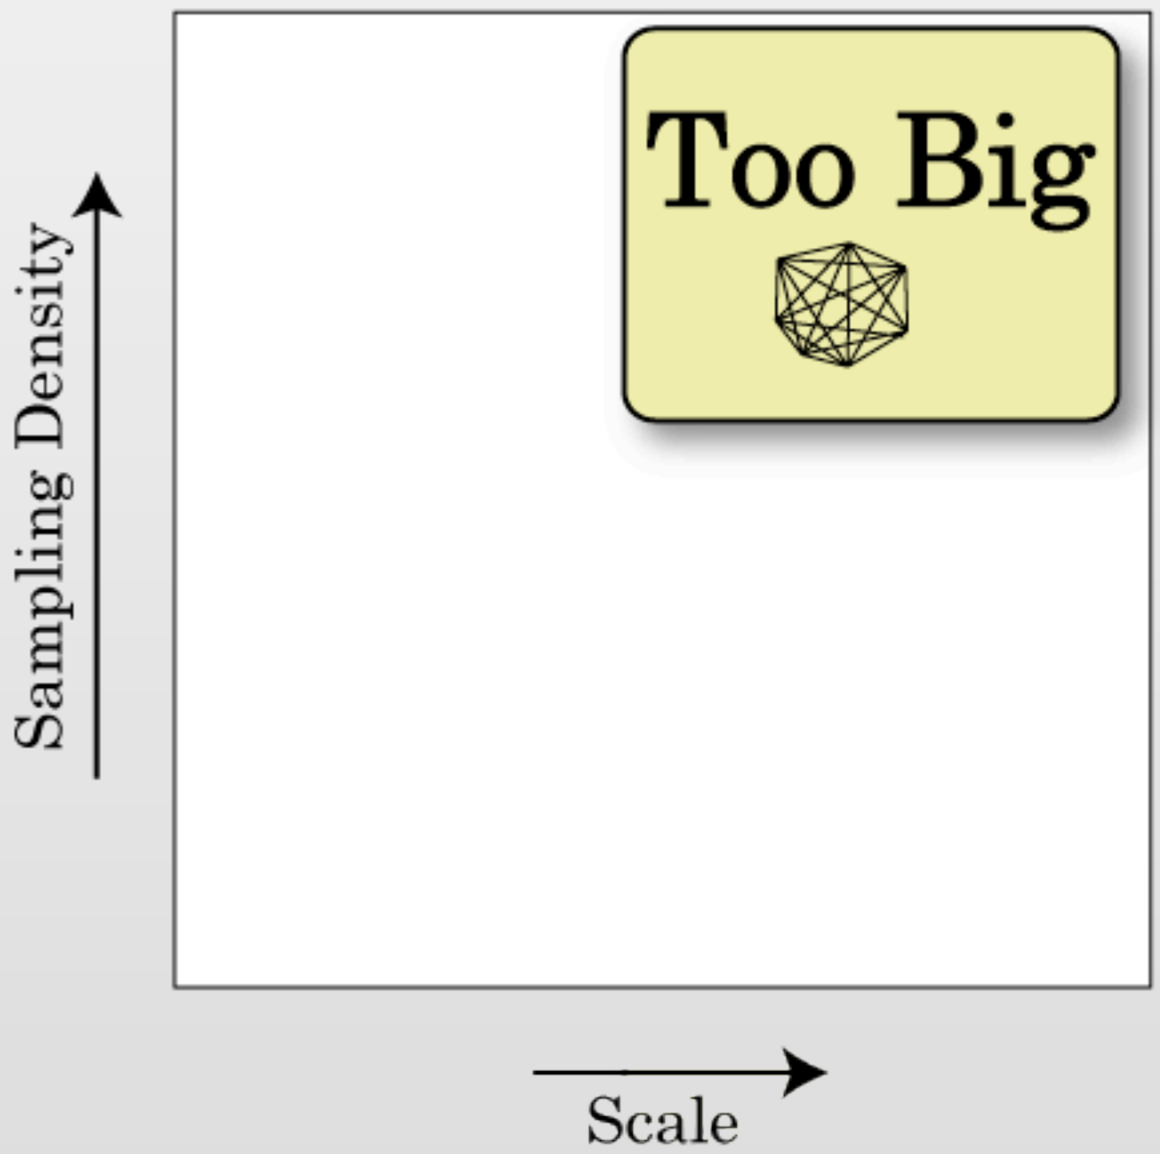

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

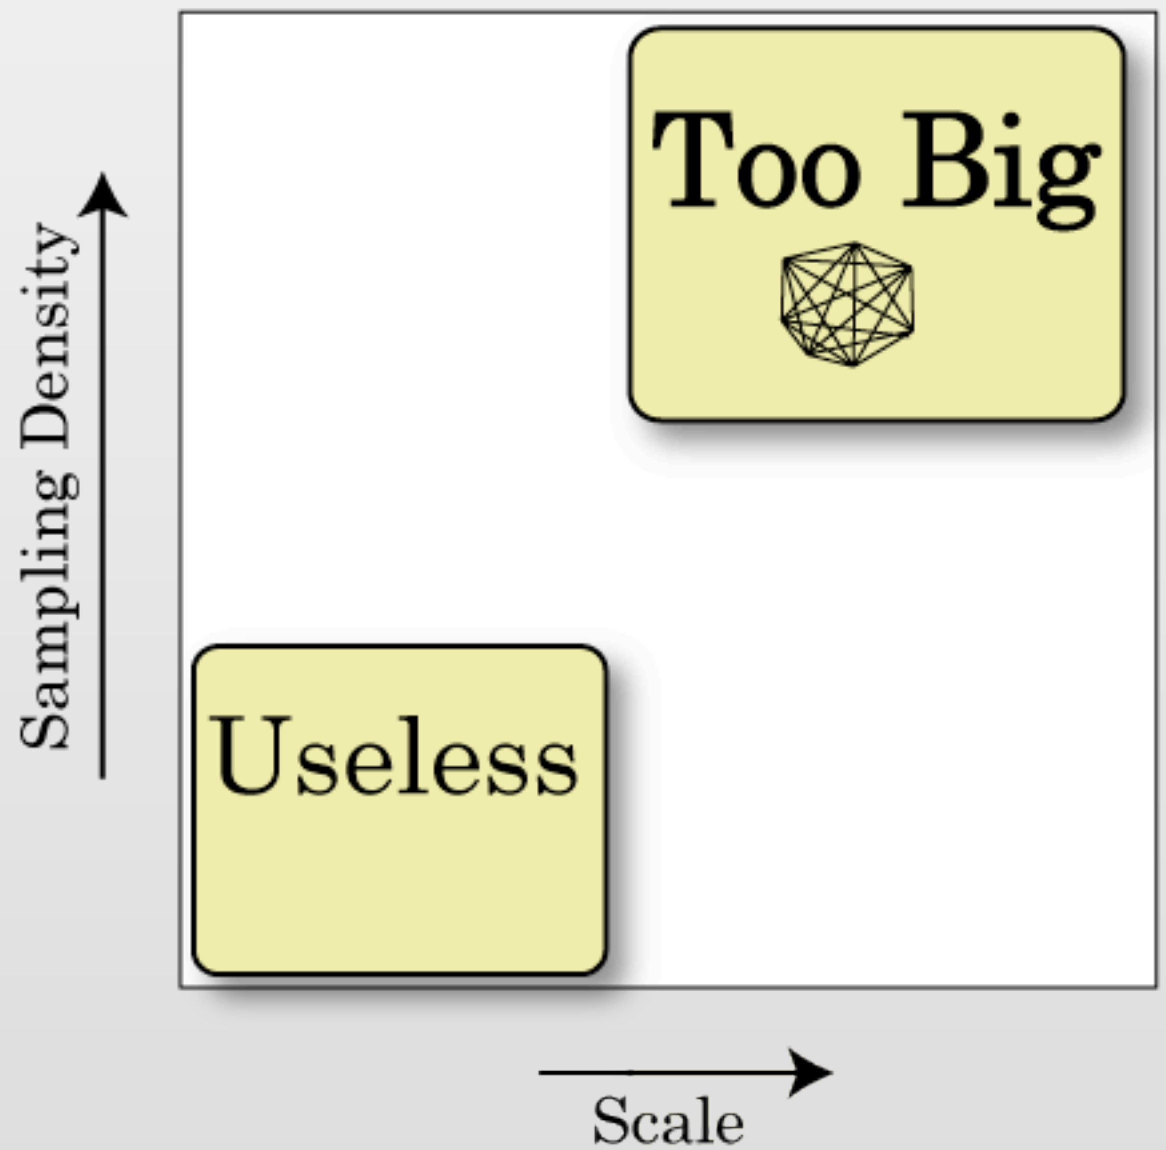

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

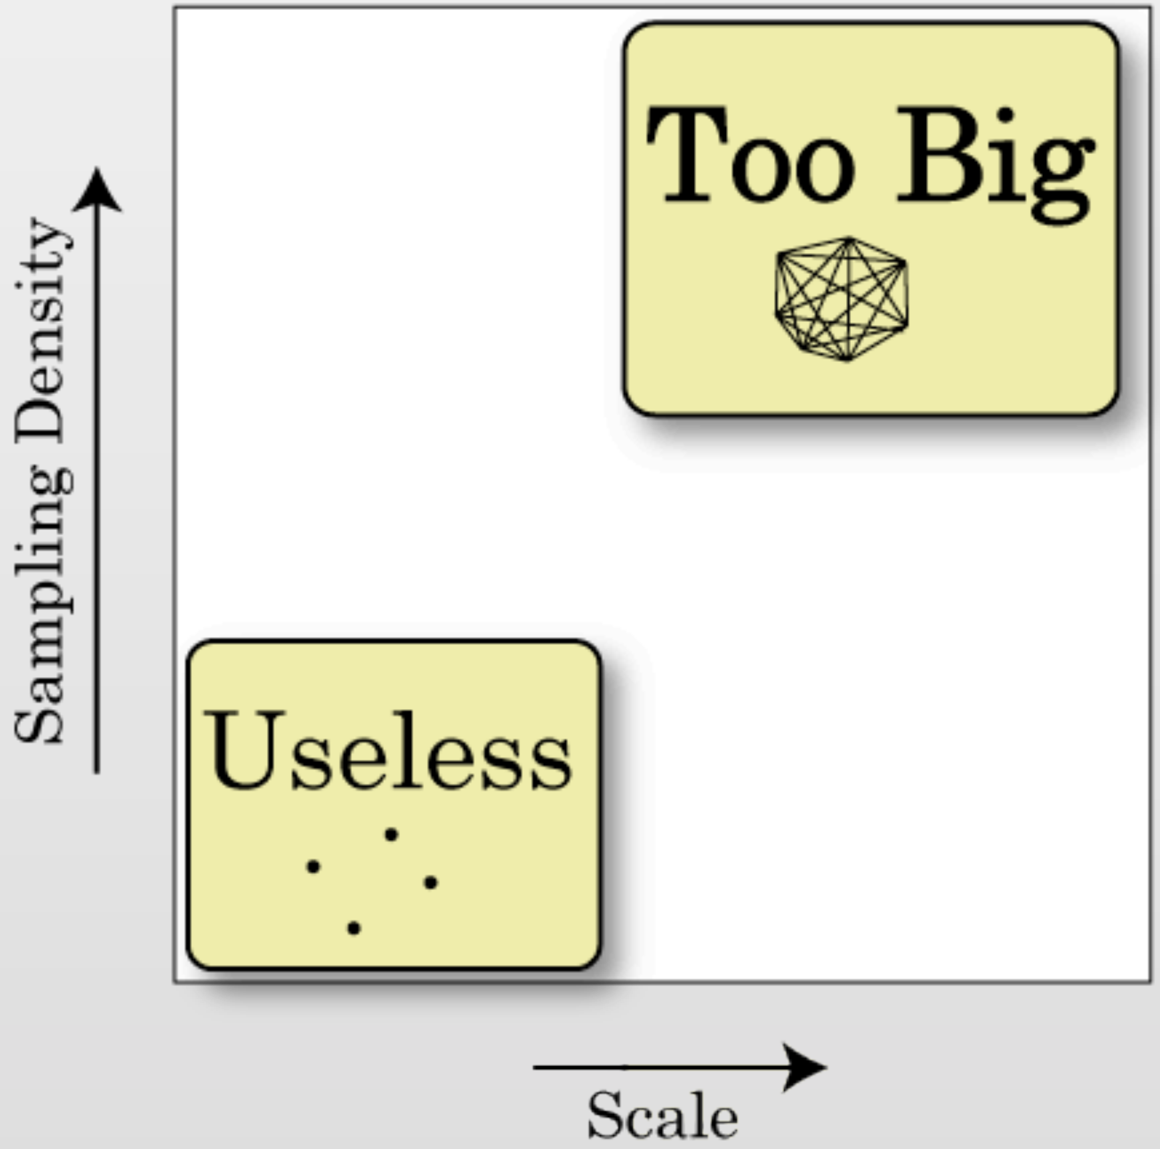

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

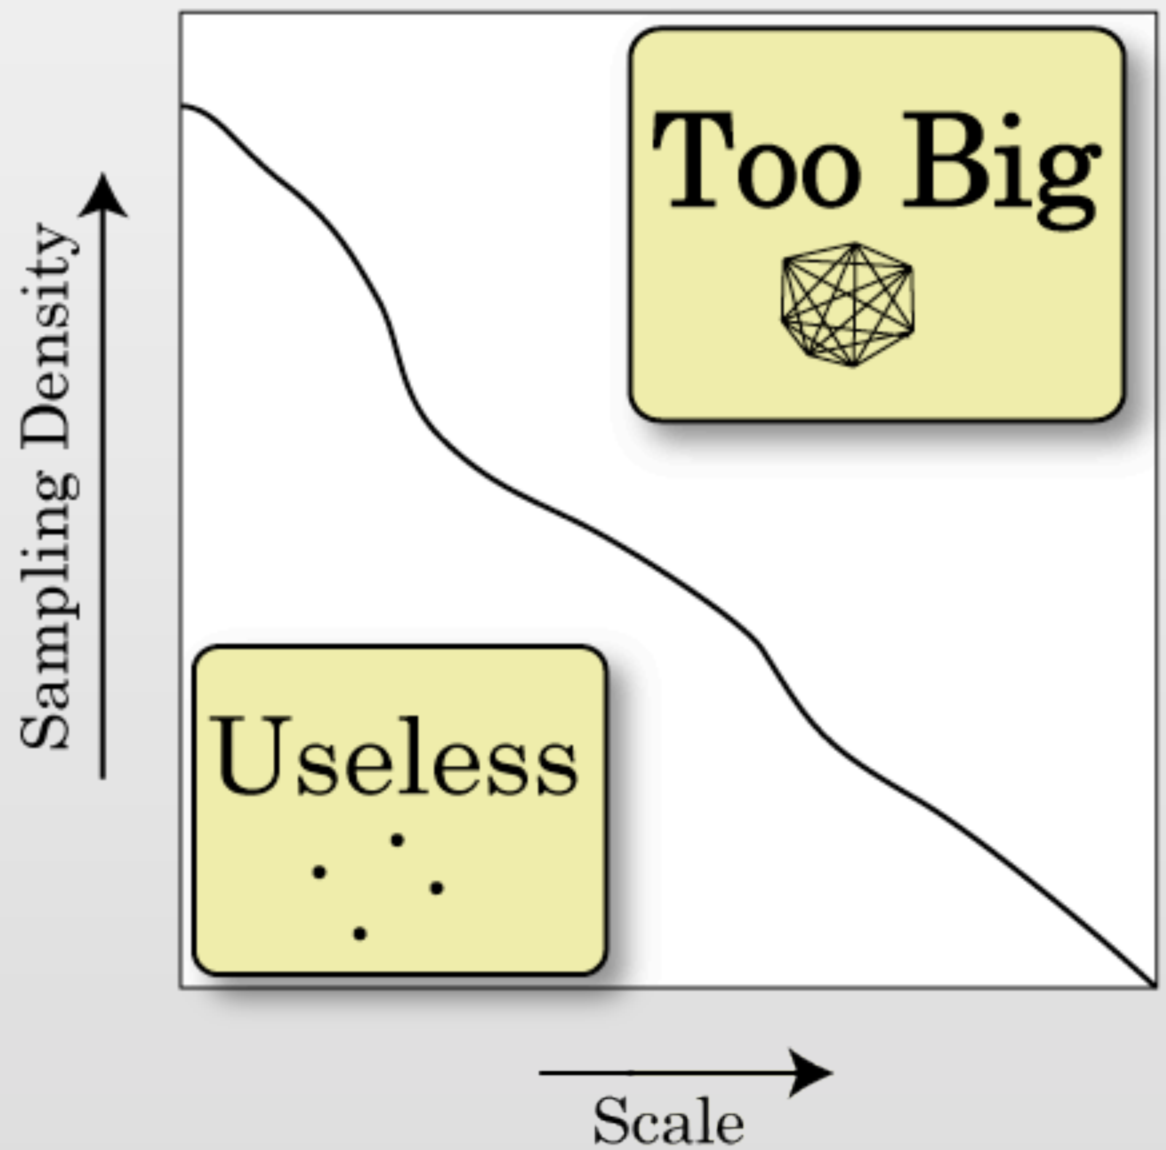

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

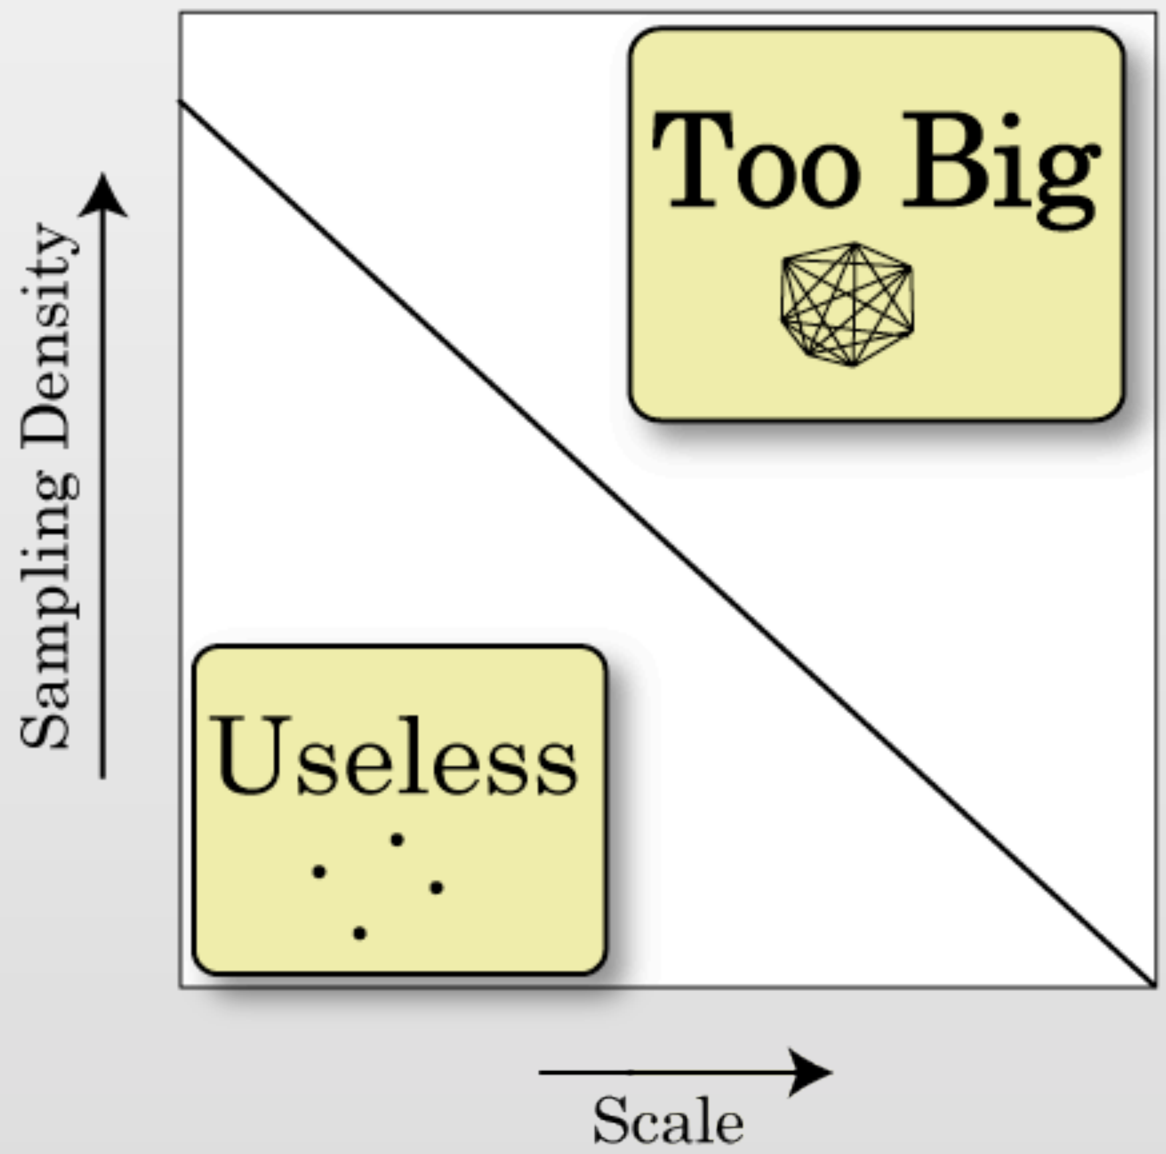

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

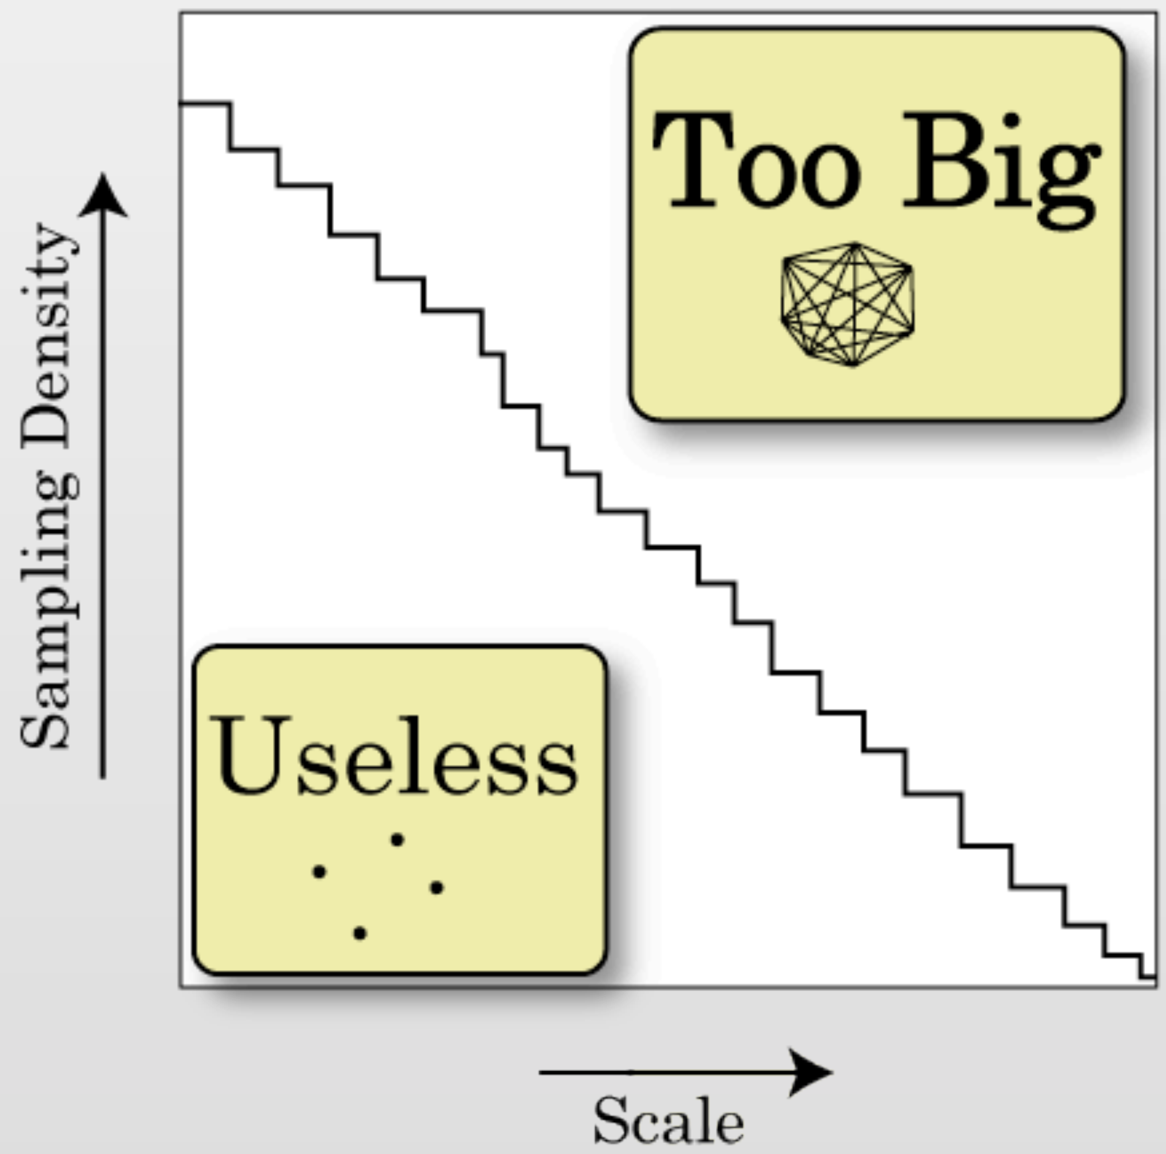

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

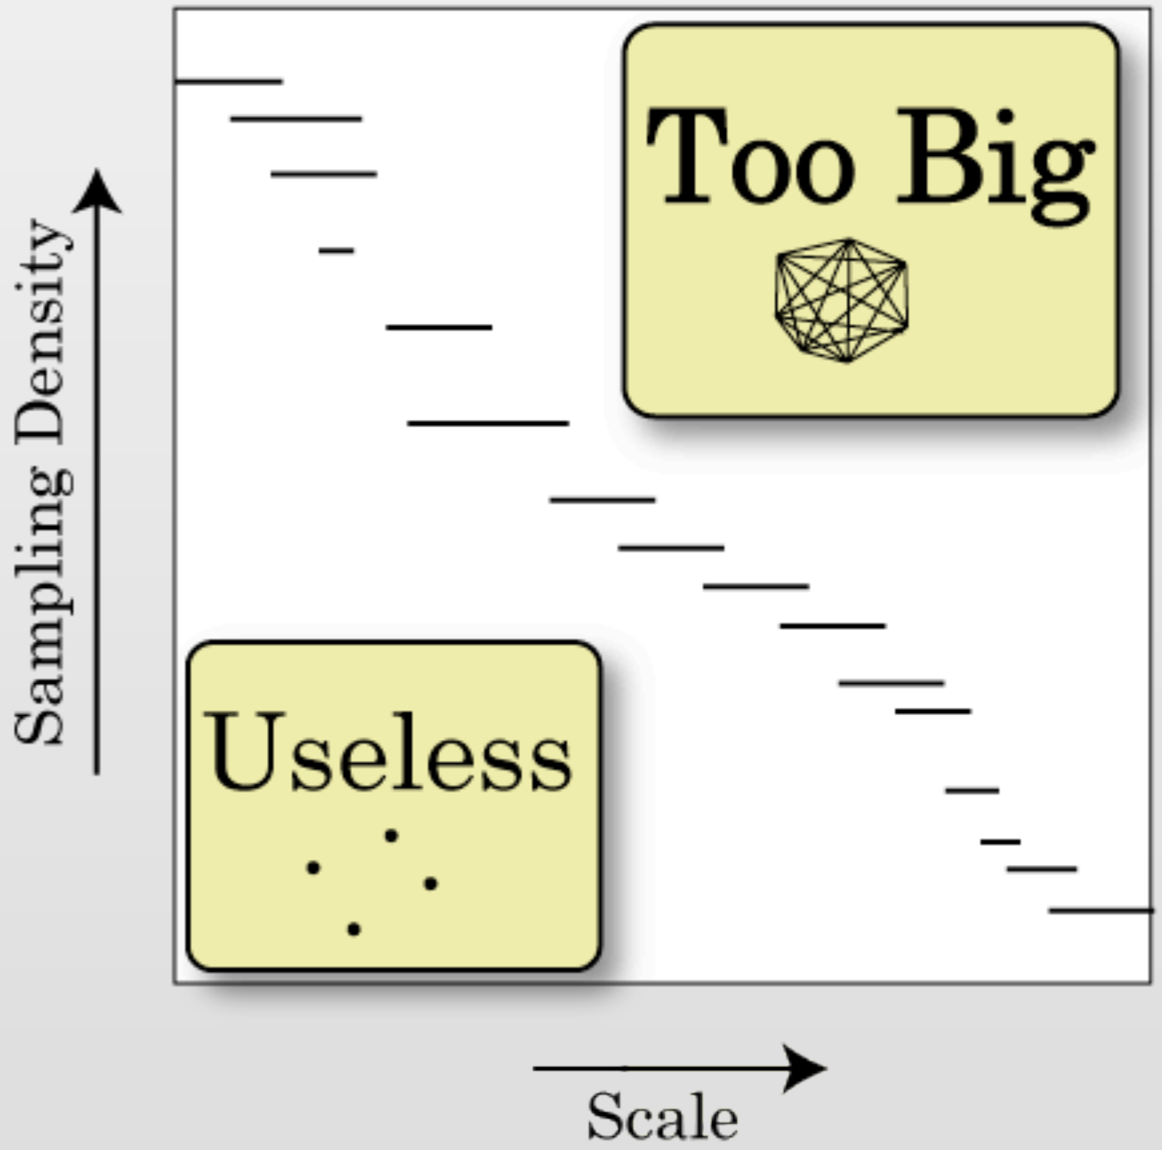

Consider a 2-dimensional filtration parameterized by both scale and sampling density.

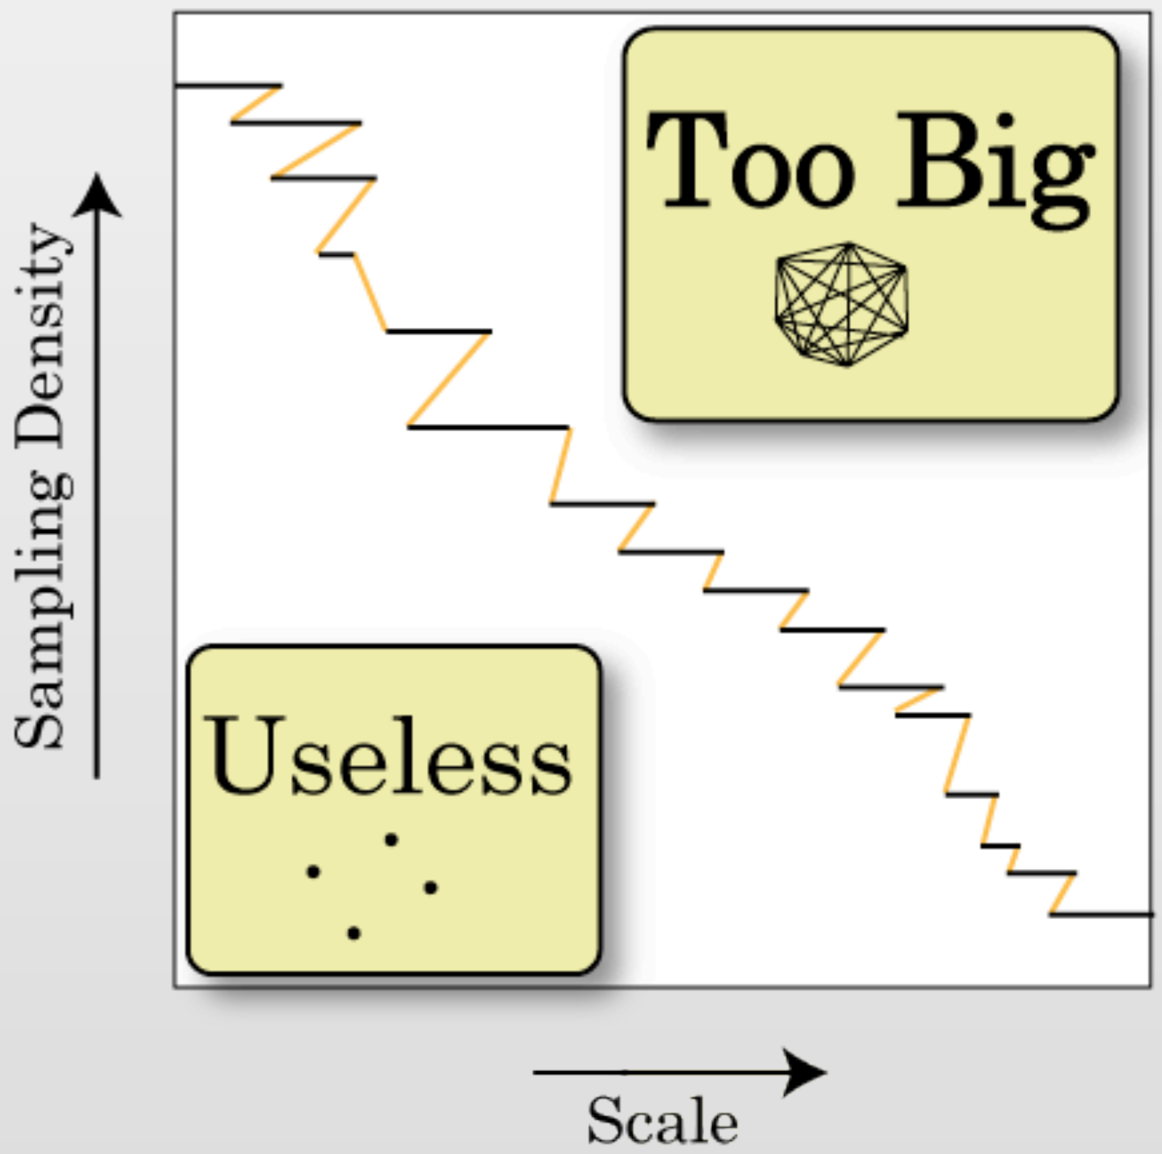

Intuition: Remove points that are covered by their neighbors.

Intuition: Remove points that are covered by their neighbors.

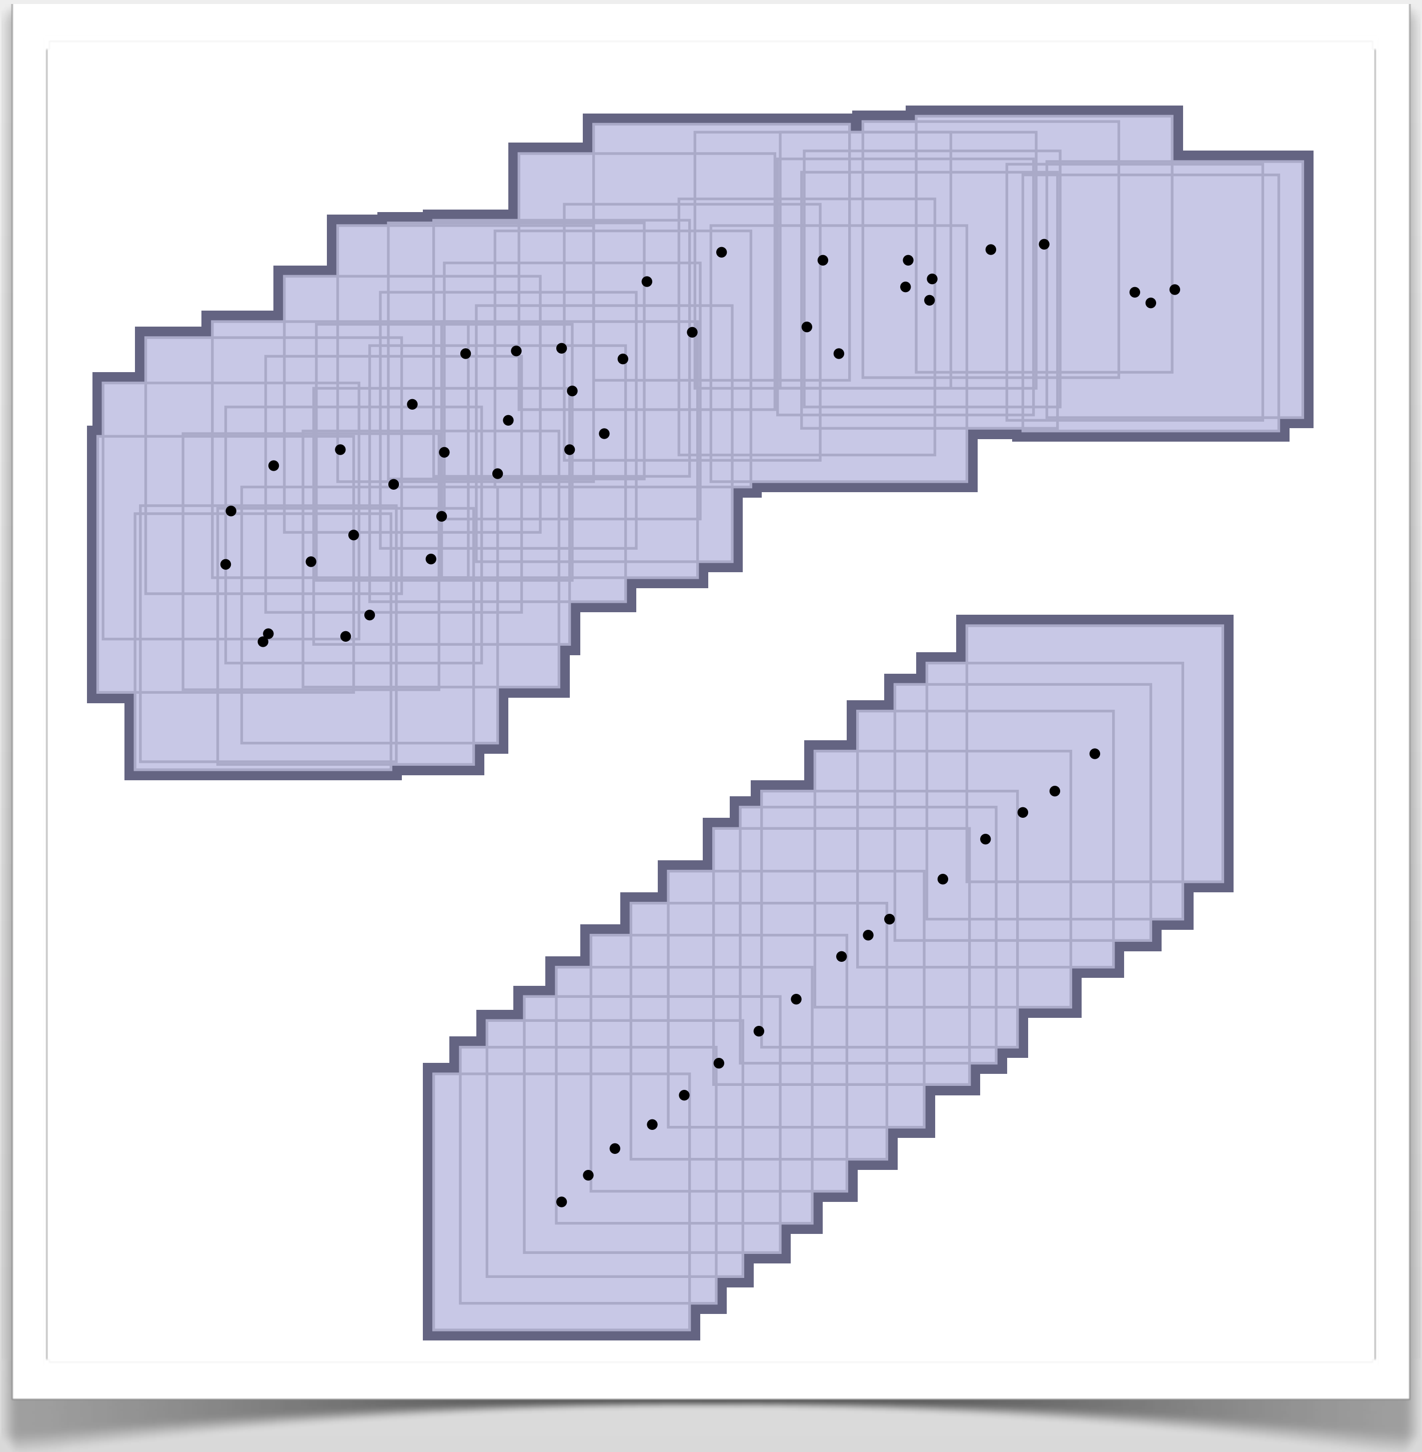

Intuition: Remove points that are covered by their neighbors.

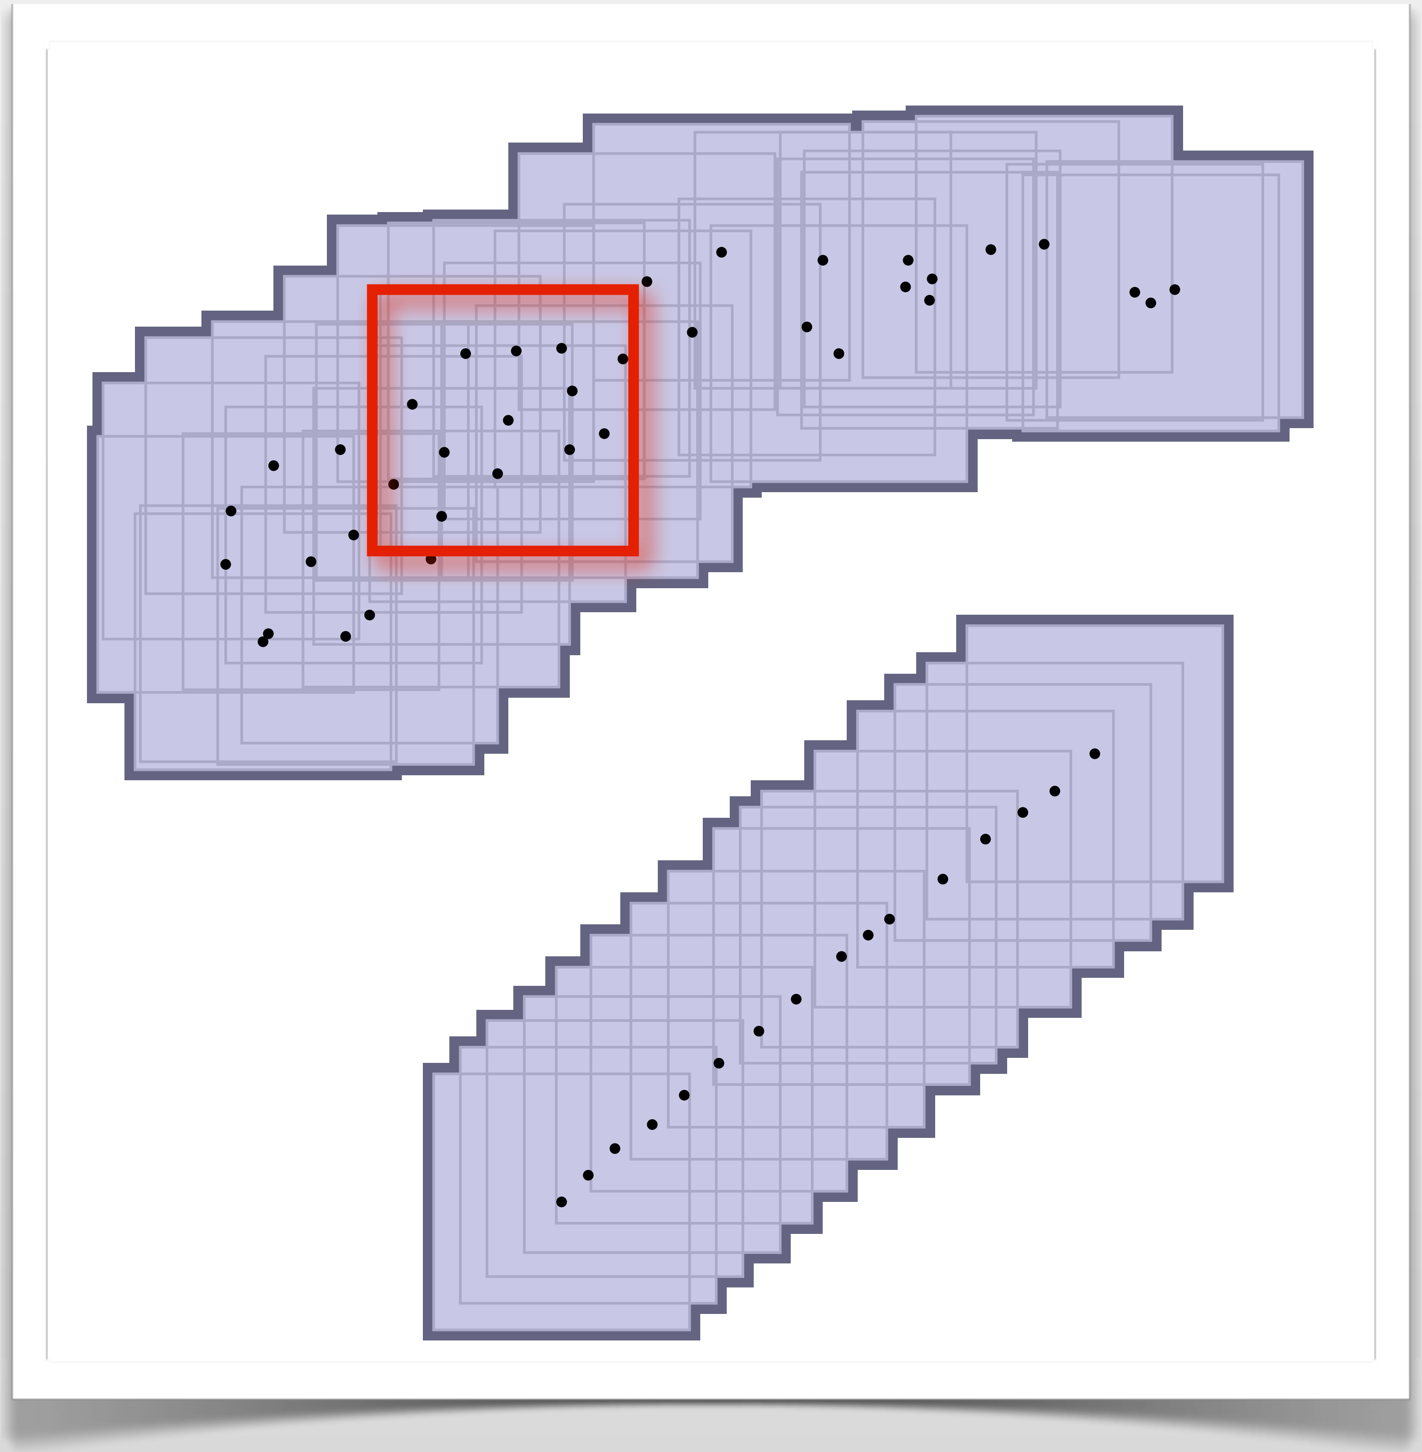

Intuition: Remove points that are covered by their neighbors.

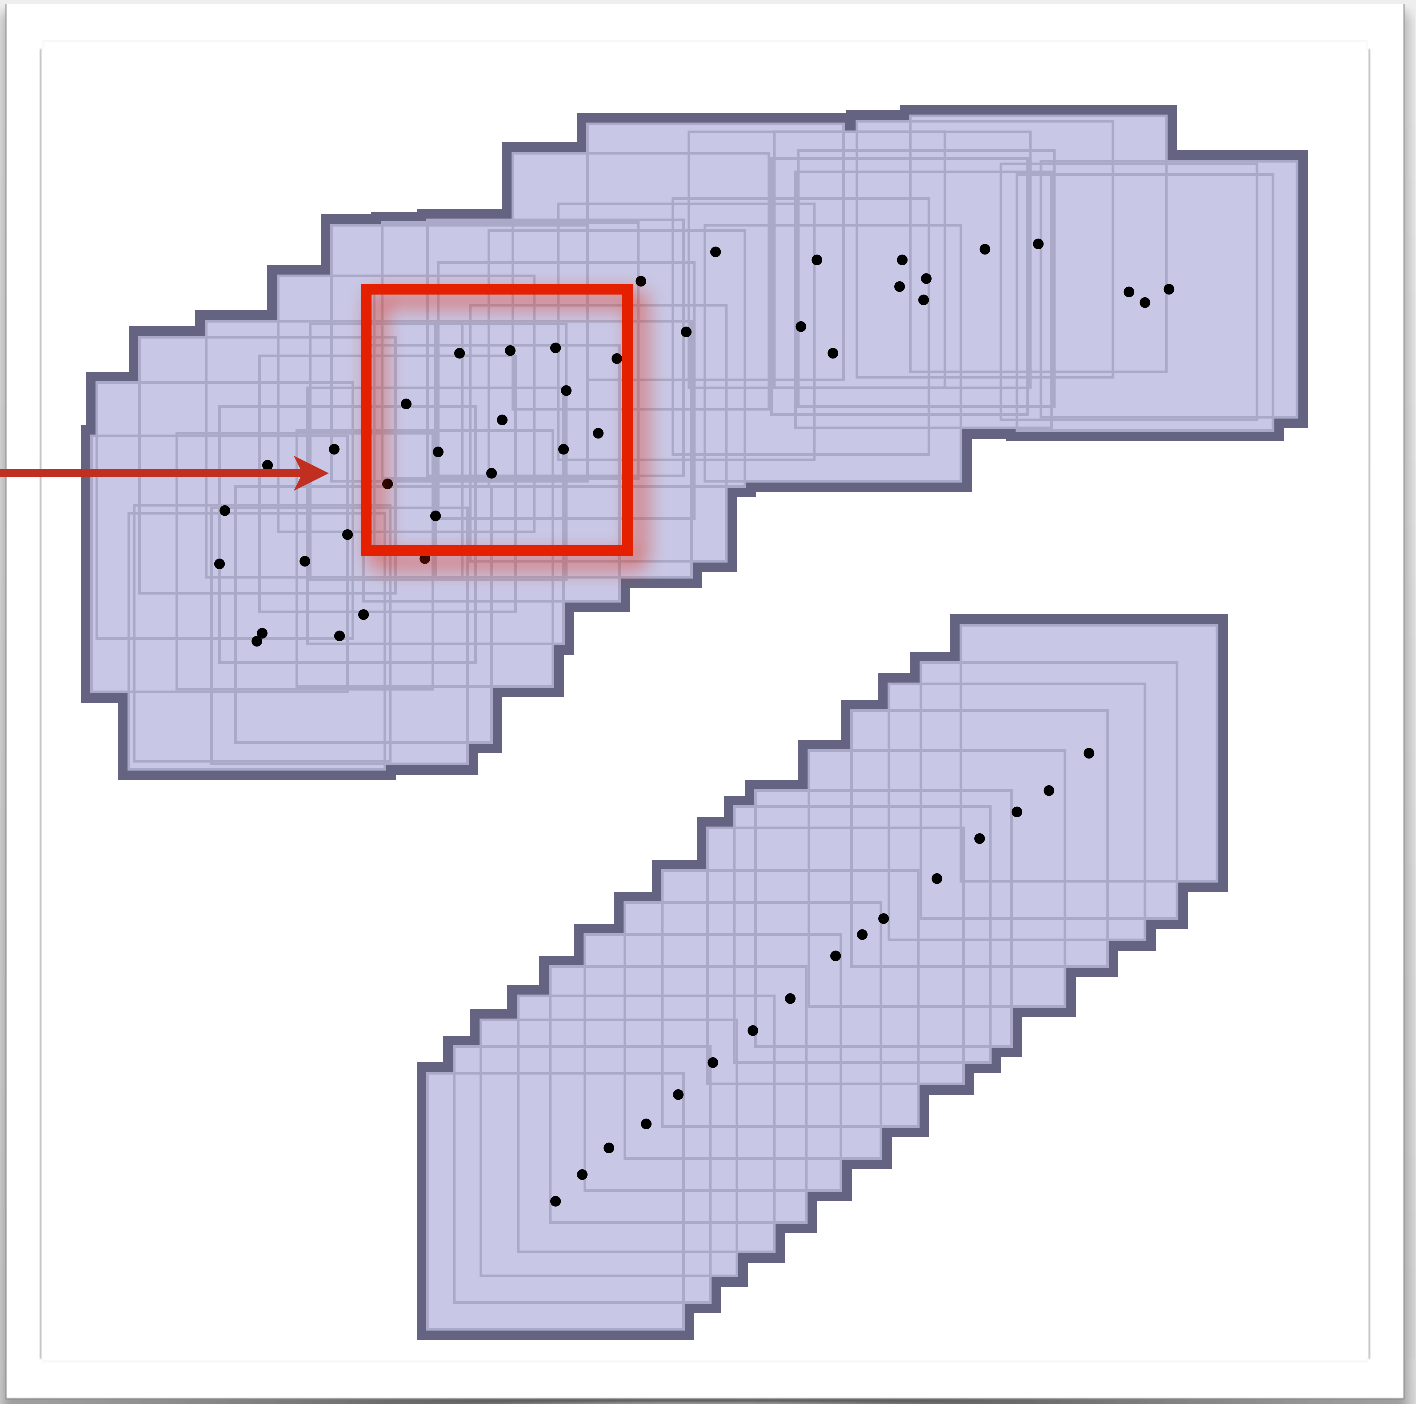

No change in topology

Intuition: Remove points that are covered by their neighbors.

No change in topology

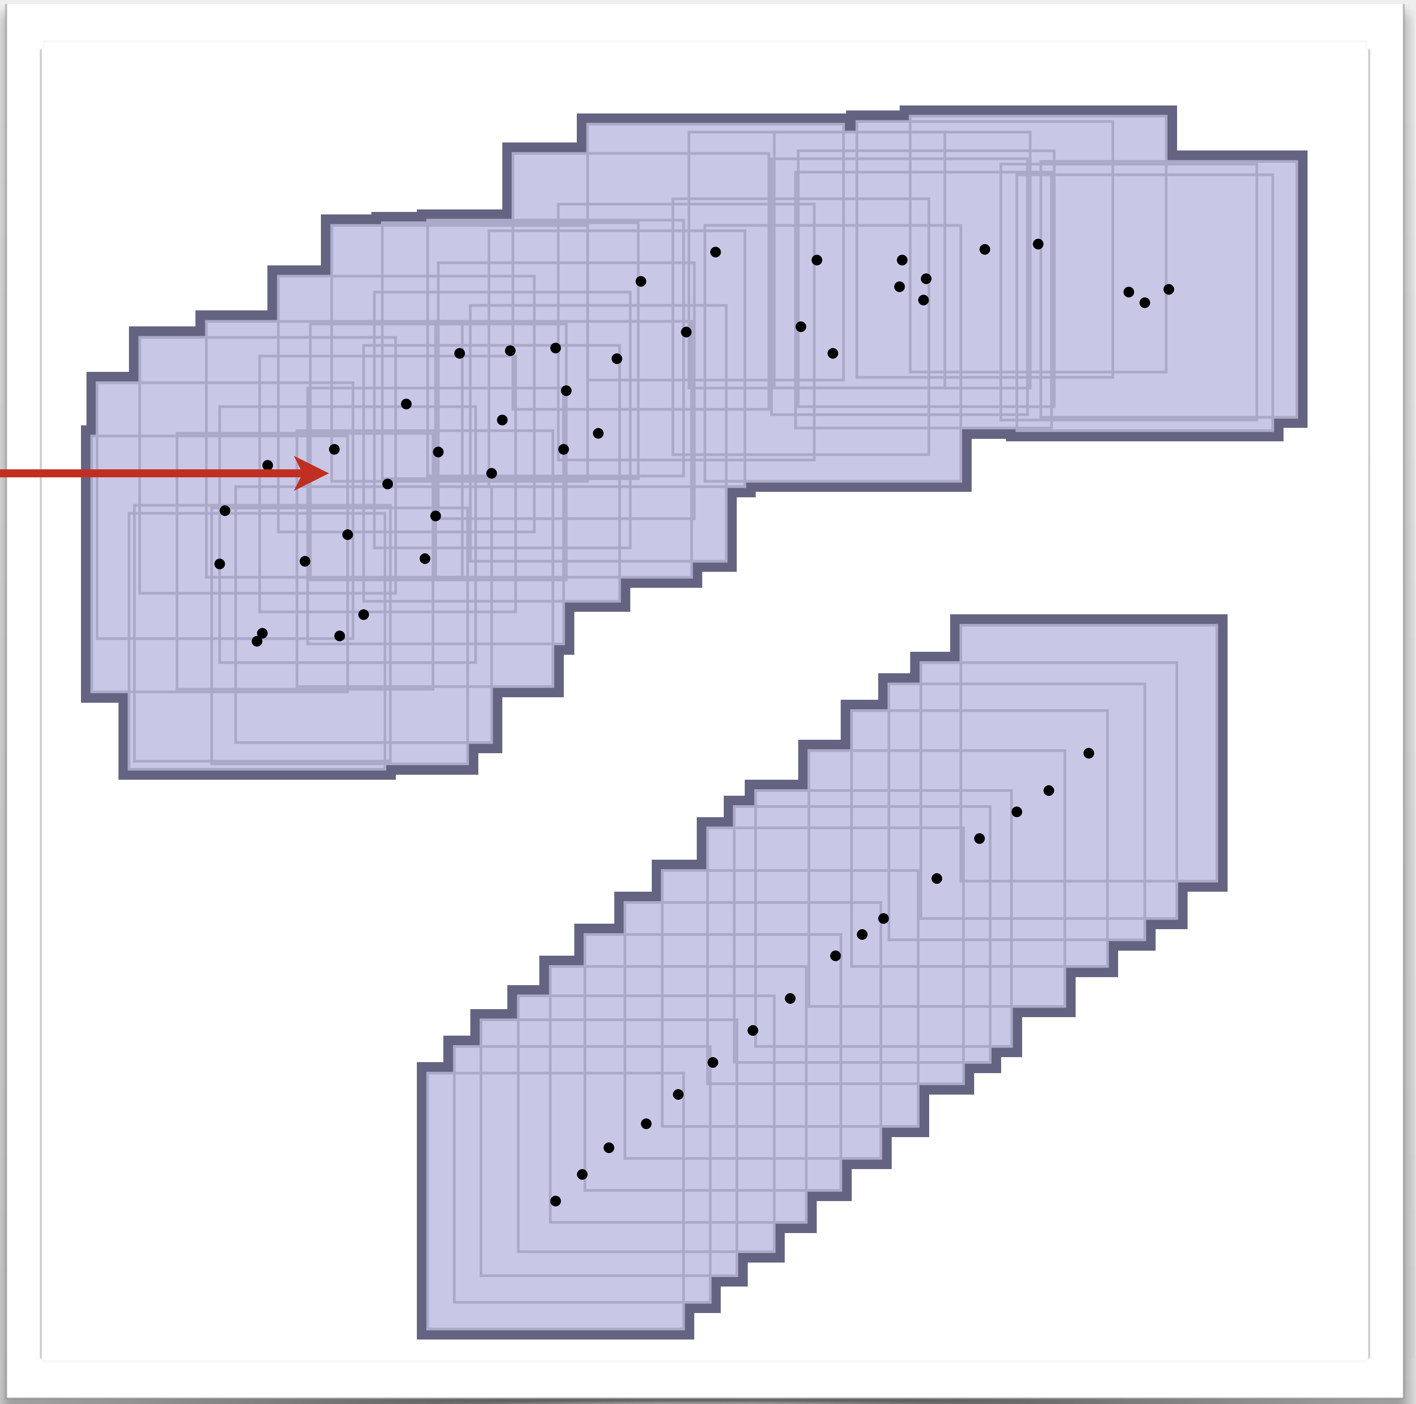

Intuition: Remove points that are covered by their neighbors.

No change
in topology

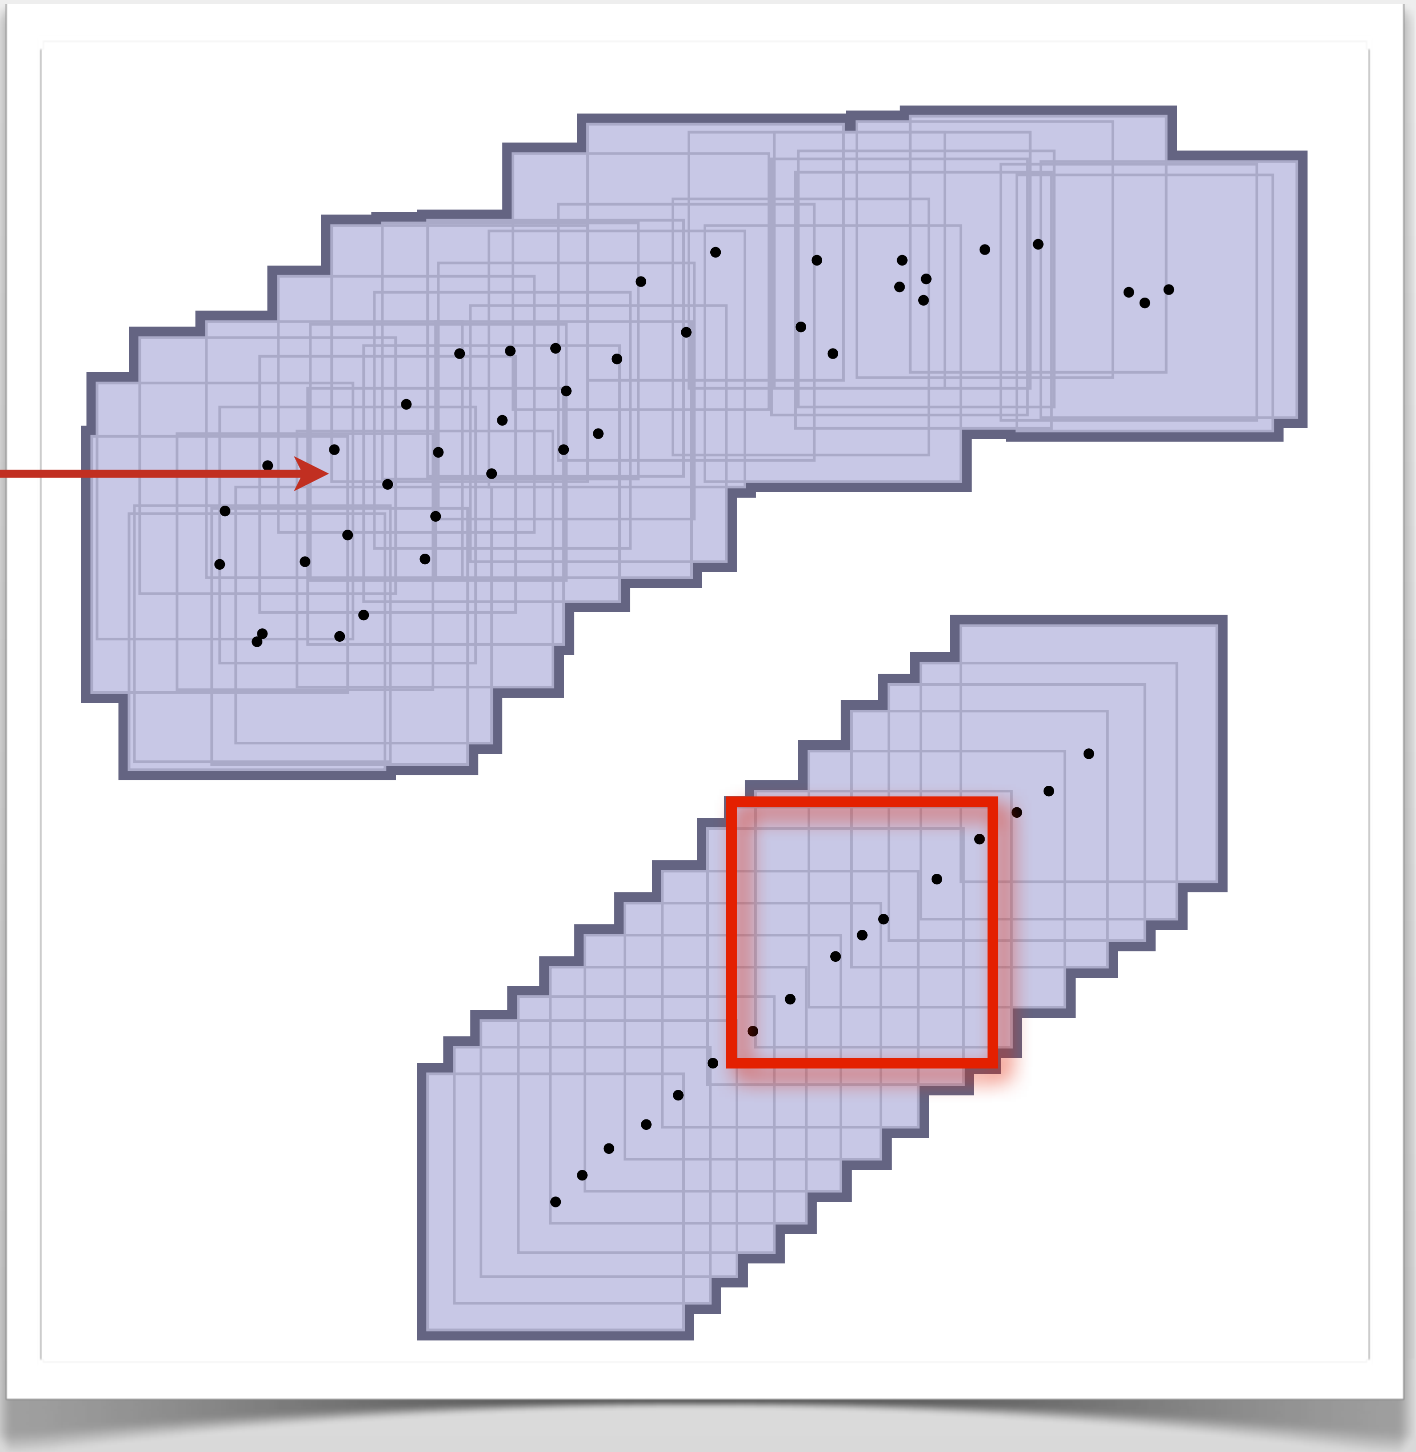

Intuition: Remove points that are covered by their neighbors.

No change in topology

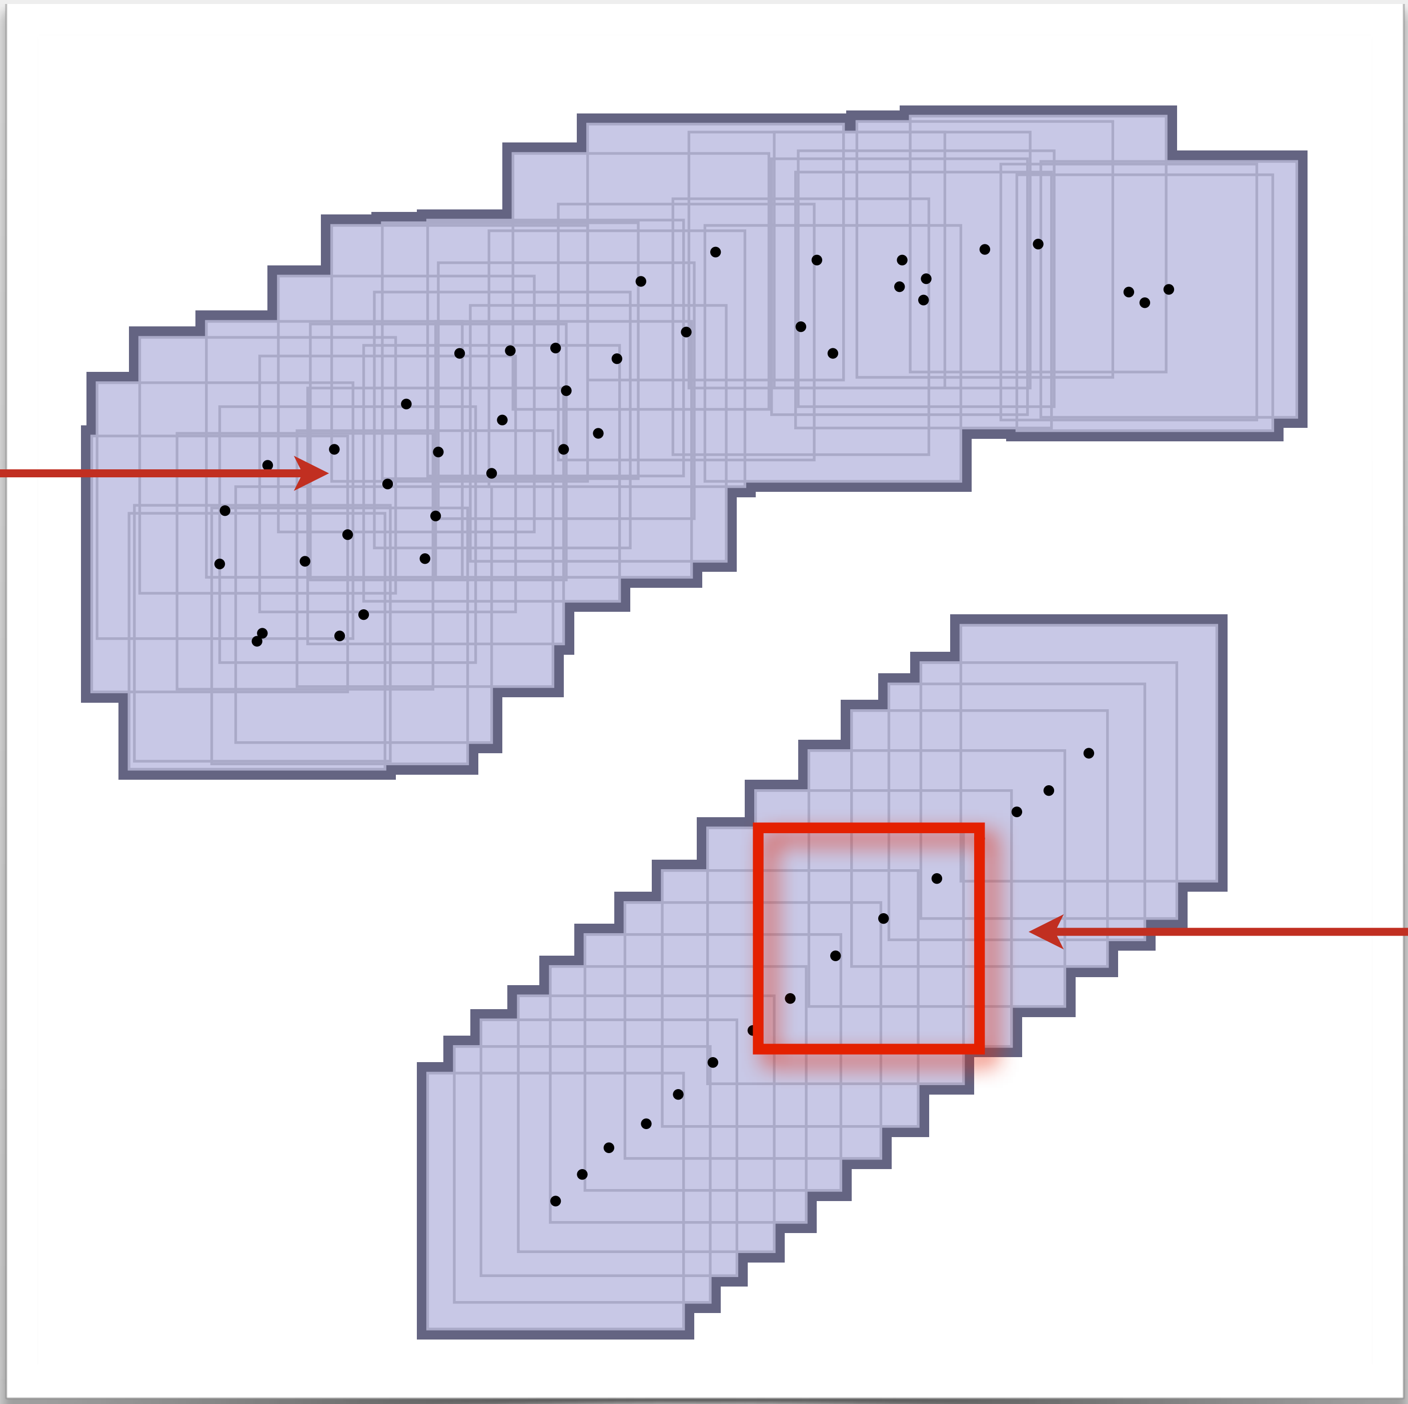

Perturb the metric

Two tricks:

Two tricks:

1. Embed the zigzag in a topologically equivalent filtration.

Two tricks:

1. Embed the zigzag in a topologically equivalent filtration.

$$
\begin{array}{lccccccccc}
\text{Standard filtration} & \hookrightarrow & \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\
 & & \big\uparrow & & \big\uparrow & & \big\uparrow & \\
\text{Zigzag filtration} & \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow
\end{array}
$$

Two tricks:

1 Embed the zigzag in a topologically equivalent filtration.

2 Perturb the metric so the persistence module does not zigzag.

$$
\begin{array}{lccccccc}
\text{Standard filtration} & \hookrightarrow \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\
 & \uparrow & & \uparrow & & \uparrow & \\
\text{Zigzag filtration} & \hookleftarrow Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow
\end{array}
$$

Two tricks:

1 Embed the zigzag in a topologically equivalent filtration.

2 Perturb the metric so the persistence module does not zigzag.

Standard filtration $\hookrightarrow \hat{R}_\alpha \hookrightarrow \hat{R}_\beta \hookrightarrow \hat{R}_\gamma \hookrightarrow$

$$\begin{array}{ccccccc} & \hat{R}_\alpha & & \hat{R}_\beta & & \hat{R}_\gamma & \\ & \uparrow & & \uparrow & & \uparrow & \\ \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow \end{array}$$

Zigzag filtration $\hookleftarrow Q_\alpha \hookrightarrow Q_\beta \hookleftarrow Q_\gamma \hookrightarrow$

$$\leftarrow \mathrm{H}(Q_\alpha) \rightarrow \mathrm{H}(Q_\beta) \leftarrow \mathrm{H}(Q_\gamma) \rightarrow$$

Two tricks:

1 Embed the zigzag in a topologically equivalent filtration.

2 Perturb the metric so the persistence module does not zigzag.

$$\text{Standard filtration} \quad \hookrightarrow \hat{R}_\alpha \hookrightarrow \hat{R}_\beta \hookrightarrow \hat{R}_\gamma \hookrightarrow$$

$$\begin{array}{ccc} \uparrow & \uparrow & \uparrow \end{array}$$

$$\text{Zigzag filtration} \quad \hookleftarrow Q_\alpha \hookrightarrow Q_\beta \hookleftarrow Q_\gamma \hookrightarrow$$

$$\xleftarrow{\cong} \mathrm{H}(Q_\alpha) \to \mathrm{H}(Q_\beta) \xleftarrow{\cong} \mathrm{H}(Q_\gamma) \to$$

Two tricks:

1 Embed the zigzag in a topologically equivalent filtration.

2 Perturb the metric so the persistence module does not zigzag.

$$\begin{array}{lccccccc} \text{Standard filtration} & \hookrightarrow & \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\ & & \uparrow & & \uparrow & & \uparrow & \\ \text{Zigzag filtration} & \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow \end{array}$$

$$\boxed{\xrightarrow{\cong}} \mathrm{H}(Q_\alpha) \to \mathrm{H}(Q_\beta) \boxed{\xrightarrow{\cong}} \mathrm{H}(Q_\gamma) \to$$

Two tricks:

1 Embed the zigzag in a topologically equivalent filtration.

2 Perturb the metric so the persistence module does not zigzag.

$$\begin{array}{ll}\text{Standard filtration} & \hookrightarrow \hat{R}_\alpha \hookrightarrow \hat{R}_\beta \hookrightarrow \hat{R}_\gamma \hookrightarrow \\ & \quad\;\; \uparrow \qquad\; \uparrow \qquad\; \uparrow \\ \text{Zigzag filtration} & \hookleftarrow Q_\alpha \hookrightarrow Q_\beta \hookleftarrow Q_\gamma \hookrightarrow\end{array}$$

$$\xrightarrow{\cong} \mathrm{H}(Q_\alpha) \to \mathrm{H}(Q_\beta) \xrightarrow{\cong} \mathrm{H}(Q_\gamma) \to$$

At the homology level, there is no zigzag.

The Result:

The Result:

Given an n point metric (P, d), there exists a zigzag filtration of size $O(n)$ whose persistence diagram $(1+\varepsilon)$-approximates that of the Rips filtration.

The Result:

Given an *n* point metric *(P, d)*, there exists a zigzag filtration of size *O(n)* whose persistence diagram (1+ ε)-approximates that of the Rips filtration.

(Big-O hides doubling dimension and approximation factor, ε)

The Result:

Given an n point metric (P, d), there exists a zigzag filtration of size $O(n)$ whose persistence diagram $(1+\varepsilon)$-approximates that of the Rips filtration.

(Big-O hides doubling dimension and approximation factor, ε)

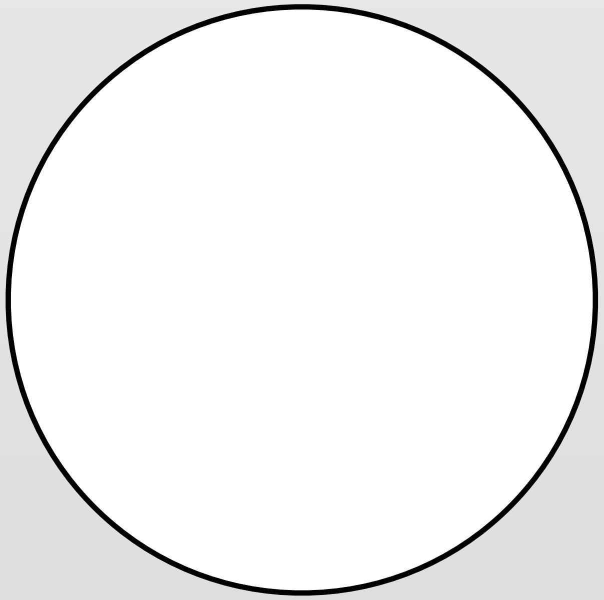

The Result:

Given an n point metric (P, d), there exists a zigzag filtration of size $O(n)$ whose persistence diagram $(1+\varepsilon)$-approximates that of the Rips filtration.

(Big-O hides doubling dimension and approximation factor, ε)

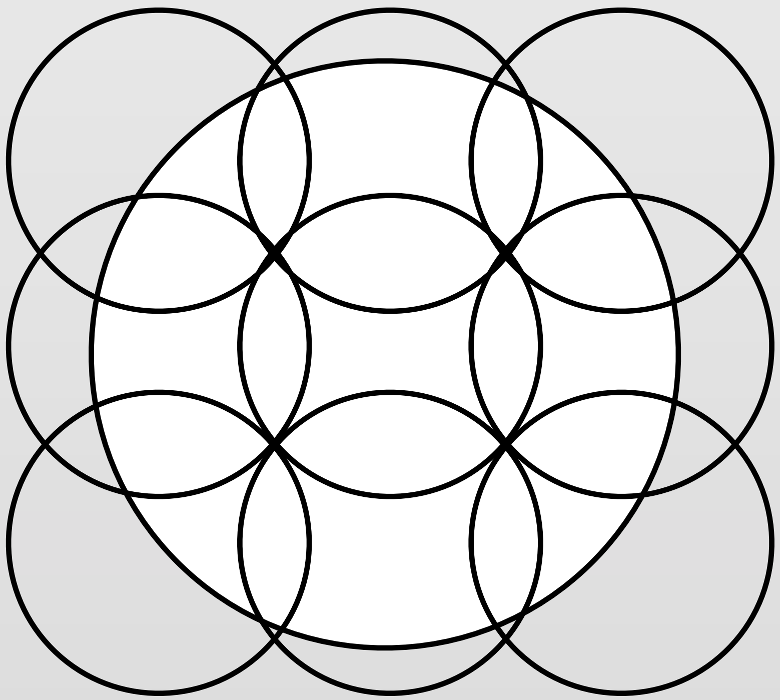

The Result:

Given an n point metric (P, d), there exists a zigzag filtration of size $O(n)$ whose persistence diagram $(1+\varepsilon)$-approximates that of the Rips filtration.

(Big-O hides doubling dimension and approximation factor, ε)

A metric with doubling dimension d is one for which every ball of radius $2r$ can be covered by 2^d balls of radius r for all r.

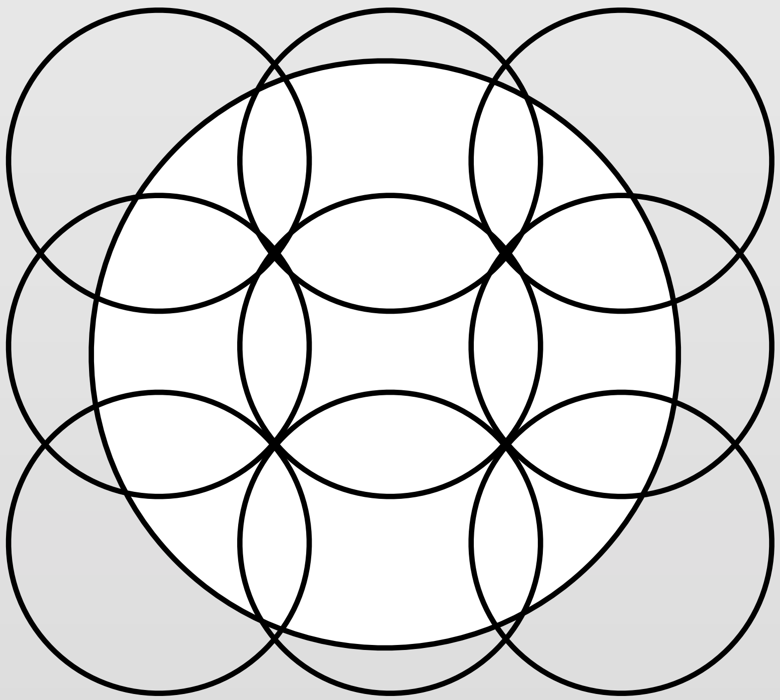

How to perturb the metric.

How to perturb the metric.

Let t_p be the time when point p is removed.

How to perturb the metric.

Let t_p be the time when point p is removed.

$$w_p(\alpha) = \begin{cases} 0 & \text{if } \alpha \leq (1-\varepsilon)t_p \\ \alpha - (1-\varepsilon)t_p & \text{if } (1-\varepsilon)t_p < \alpha < t_p \\ \varepsilon\alpha & \text{if } t_p \leq \alpha \end{cases}$$

How to perturb the metric.

Let t_p be the time when point p is removed.

$$w_p(\alpha) = \begin{cases} 0 & \text{if } \alpha \leq (1-\varepsilon)t_p \\ \alpha - (1-\varepsilon)t_p & \text{if } (1-\varepsilon)t_p < \alpha < t_p \\ \varepsilon\alpha & \text{if } t_p \leq \alpha \end{cases}$$

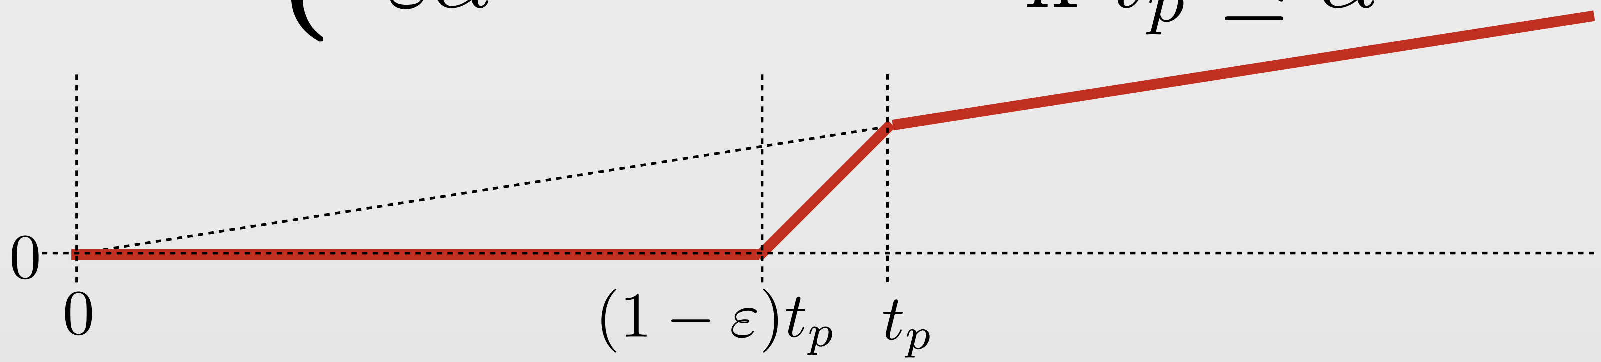

How to perturb the metric.

Let t_p be the time when point p is removed.

$$w_p(\alpha) = \begin{cases} 0 & \text{if } \alpha \leq (1-\varepsilon)t_p \\ \alpha - (1-\varepsilon)t_p & \text{if } (1-\varepsilon)t_p < \alpha < t_p \\ \varepsilon\alpha & \text{if } t_p \leq \alpha \end{cases}$$

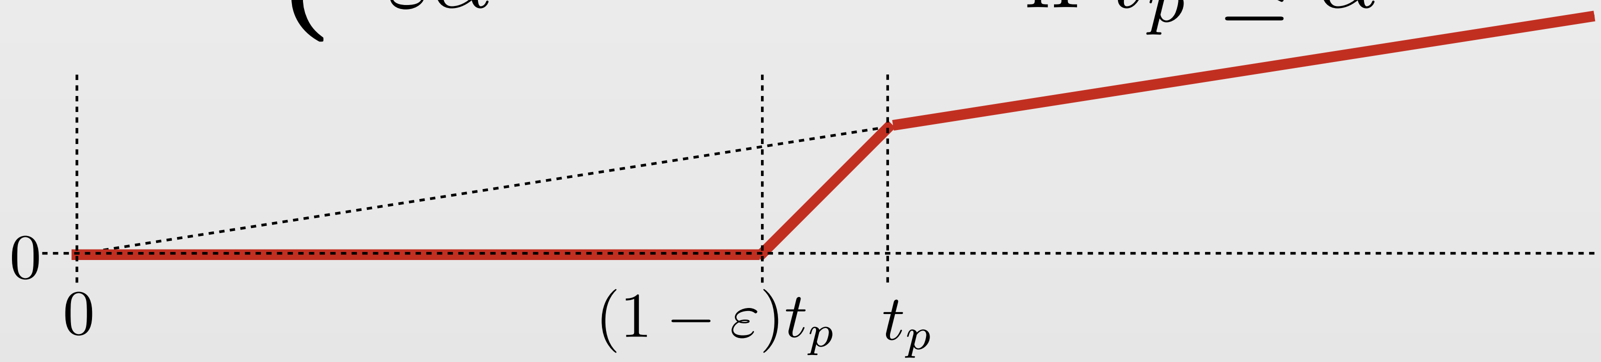

$$\hat{\mathbf{d}}_\alpha(p,q) = \mathbf{d}(p,q) + w_p(\alpha) + w_q(\alpha)$$

How to perturb the metric.

Let t_p be the time when point p is removed.

$$w_p(\alpha) = \begin{cases} 0 & \text{if } \alpha \leq (1-\varepsilon)t_p \\ \alpha - (1-\varepsilon)t_p & \text{if } (1-\varepsilon)t_p < \alpha < t_p \\ \varepsilon\alpha & \text{if } t_p \leq \alpha \end{cases}$$

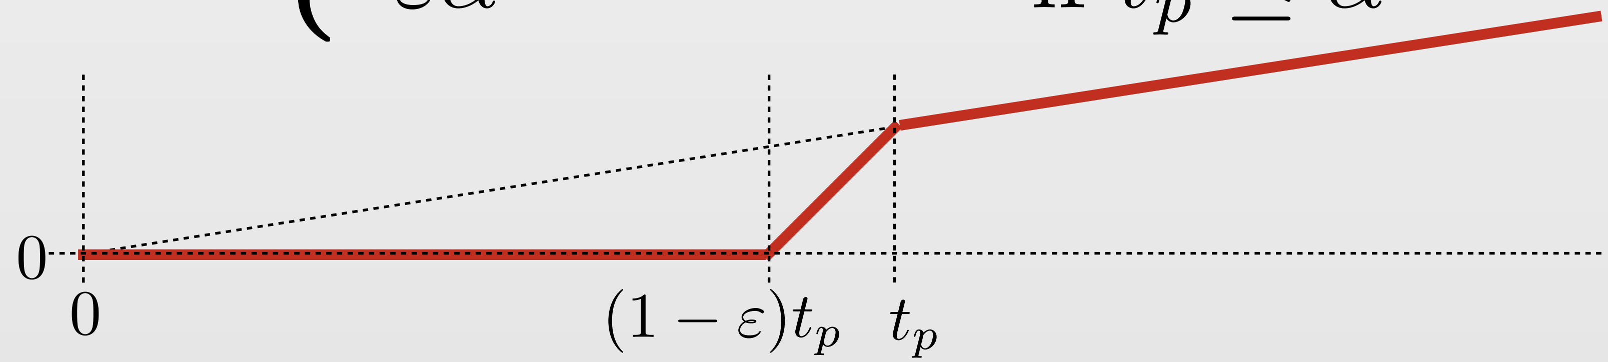

$$\hat{\mathbf{d}}_\alpha(p,q) = \mathbf{d}(p,q) + w_p(\alpha) + w_q(\alpha)$$

Rips Complex: $\sigma \in R_\alpha \Leftrightarrow \mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$

How to perturb the metric.

Let t_p be the time when point p is removed.

$$w_p(\alpha) = \begin{cases} 0 & \text{if } \alpha \leq (1-\varepsilon)t_p \\ \alpha - (1-\varepsilon)t_p & \text{if } (1-\varepsilon)t_p < \alpha < t_p \\ \varepsilon\alpha & \text{if } t_p \leq \alpha \end{cases}$$

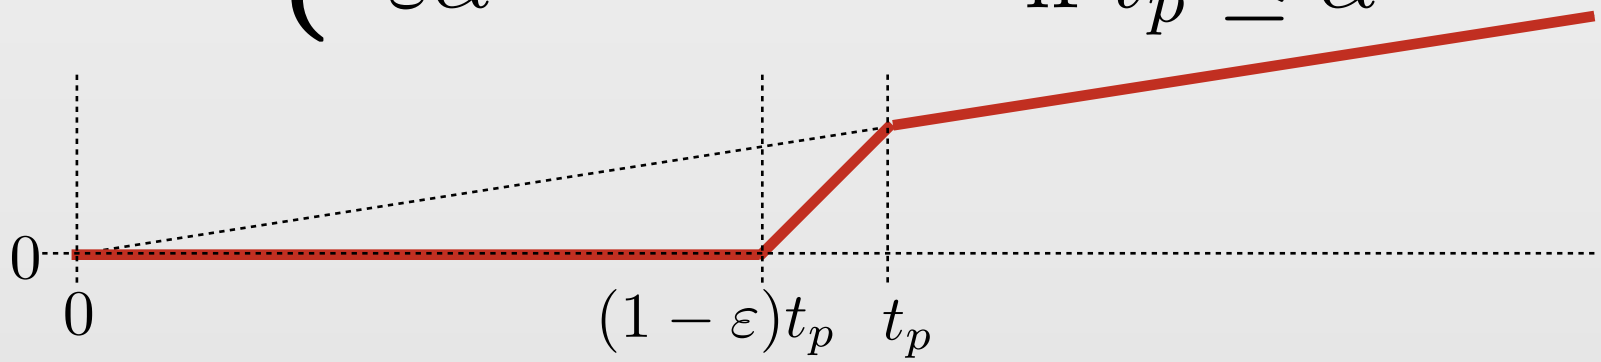

$$\hat{\mathbf{d}}_\alpha(p,q) = \mathbf{d}(p,q) + w_p(\alpha) + w_q(\alpha)$$

Rips Complex: $\sigma \in R_\alpha \Leftrightarrow \mathbf{d}(p,q) \leq 2\alpha$ for all $p,q \in \sigma$

Relaxed Rips: $\sigma \in \hat{R}_\alpha \Leftrightarrow \hat{\mathbf{d}}_\alpha(p,q) \leq 2\alpha$ for all $p,q \in \sigma$

How to perturb the metric.

Let t_p be the time when point p is removed.

$$w_p(\alpha) = \begin{cases} 0 & \text{if } \alpha \leq (1-\varepsilon)t_p \\ \alpha - (1-\varepsilon)t_p & \text{if } (1-\varepsilon)t_p < \alpha < t_p \\ \varepsilon\alpha & \text{if } t_p \leq \alpha \end{cases}$$

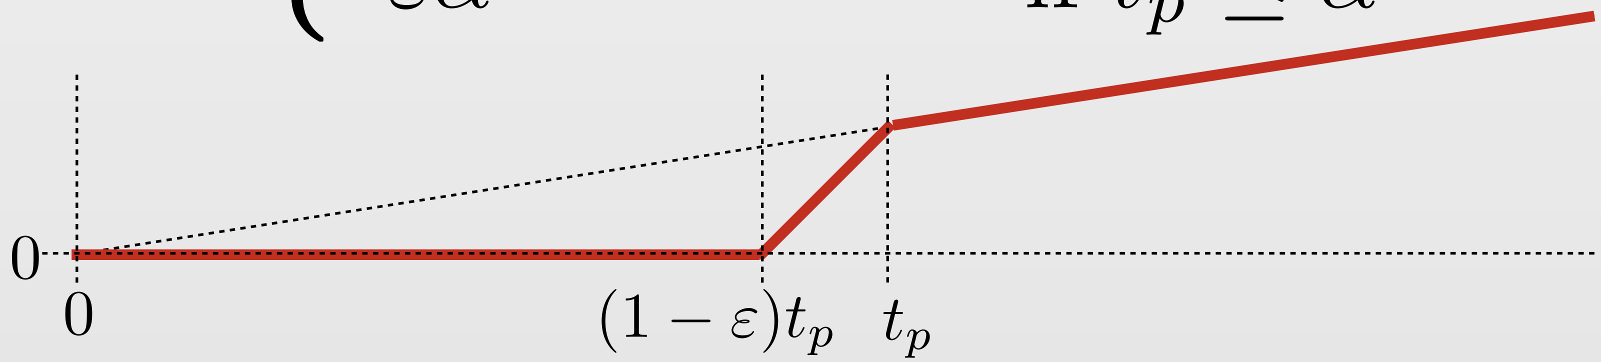

$$\hat{\mathbf{d}}_\alpha(p,q) = \mathbf{d}(p,q) + w_p(\alpha) + w_q(\alpha)$$

Rips Complex: $\sigma \in R_\alpha \Leftrightarrow \mathbf{d}(p,q) \leq 2\alpha$ for all $p, q \in \sigma$

Relaxed Rips: $\sigma \in \hat{R}_\alpha \Leftrightarrow \hat{\mathbf{d}}_\alpha(p,q) \leq 2\alpha$ for all $p, q \in \sigma$

$$R_{\frac{\alpha}{1+\varepsilon}} \subseteq \hat{R}_\alpha \subseteq R_\alpha$$

Net-trees

Net-trees

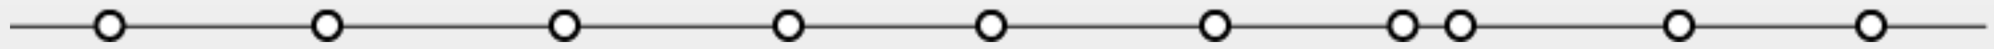

Net-trees

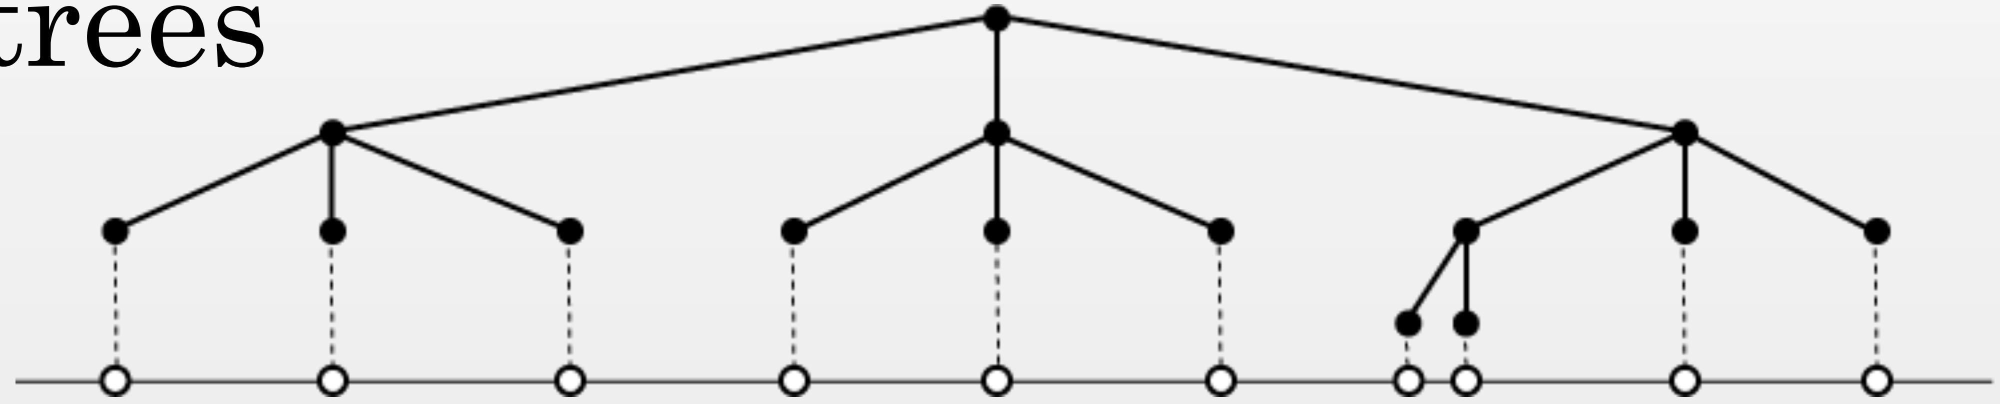

Net-trees

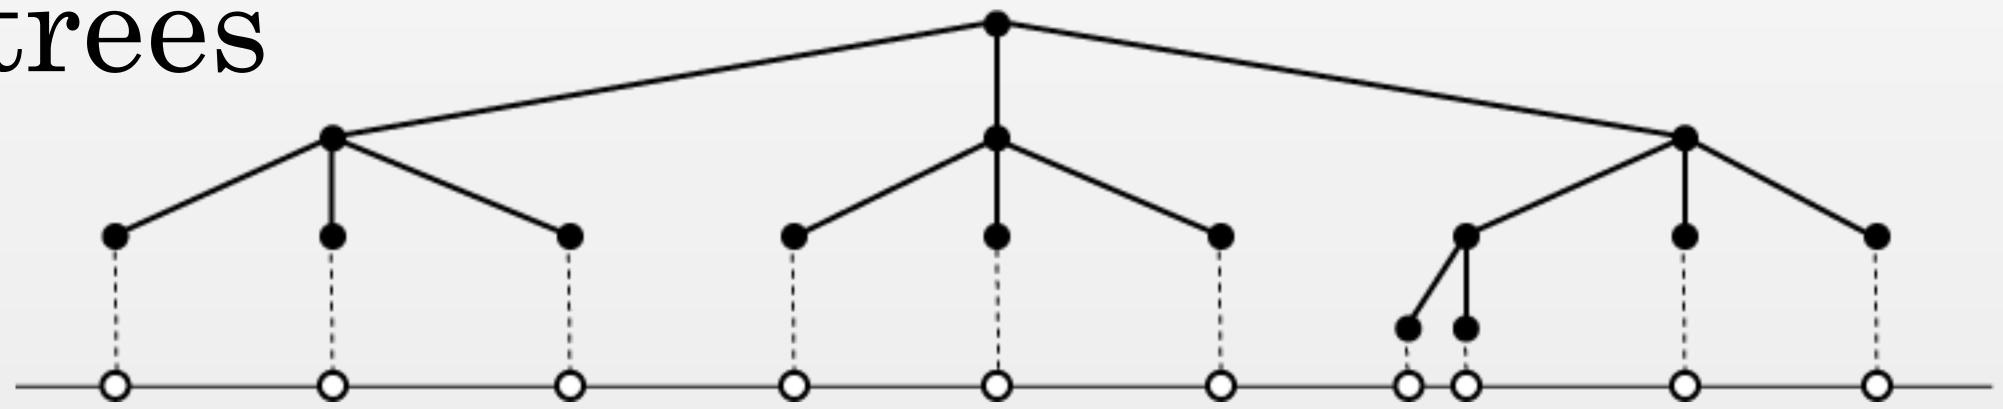

A generalization of quadtrees to metric spaces.

Net-trees

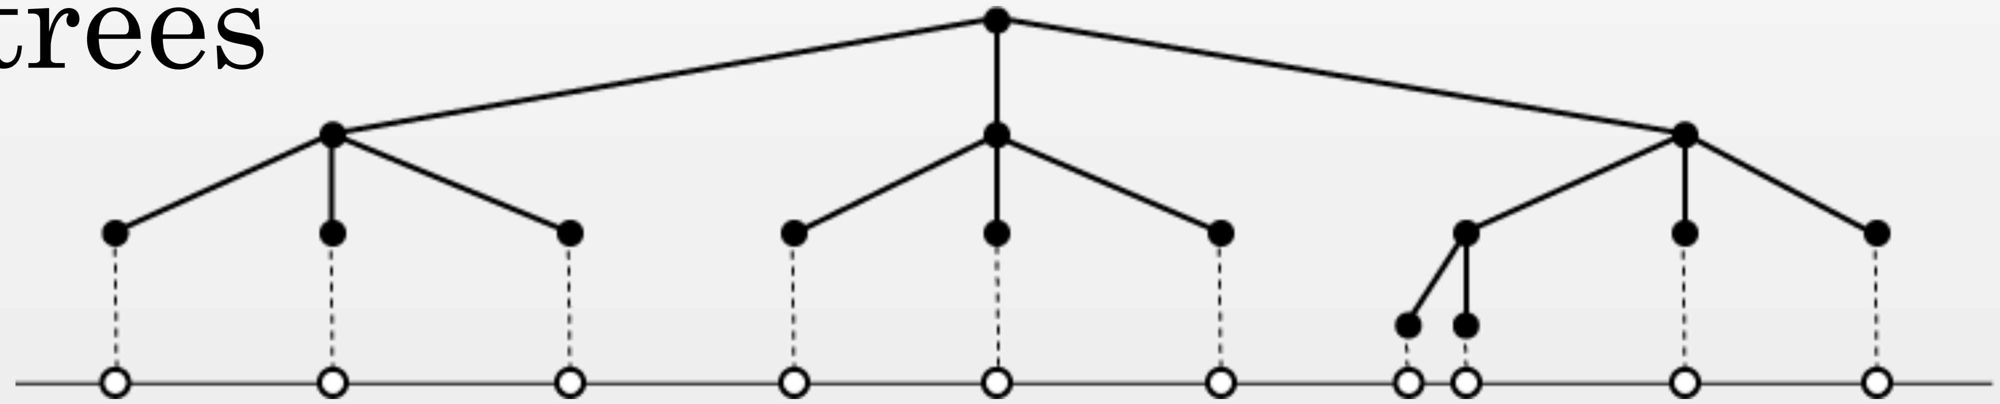

A generalization of quadtrees to metric spaces.
One leaf per point of P.

Net-trees

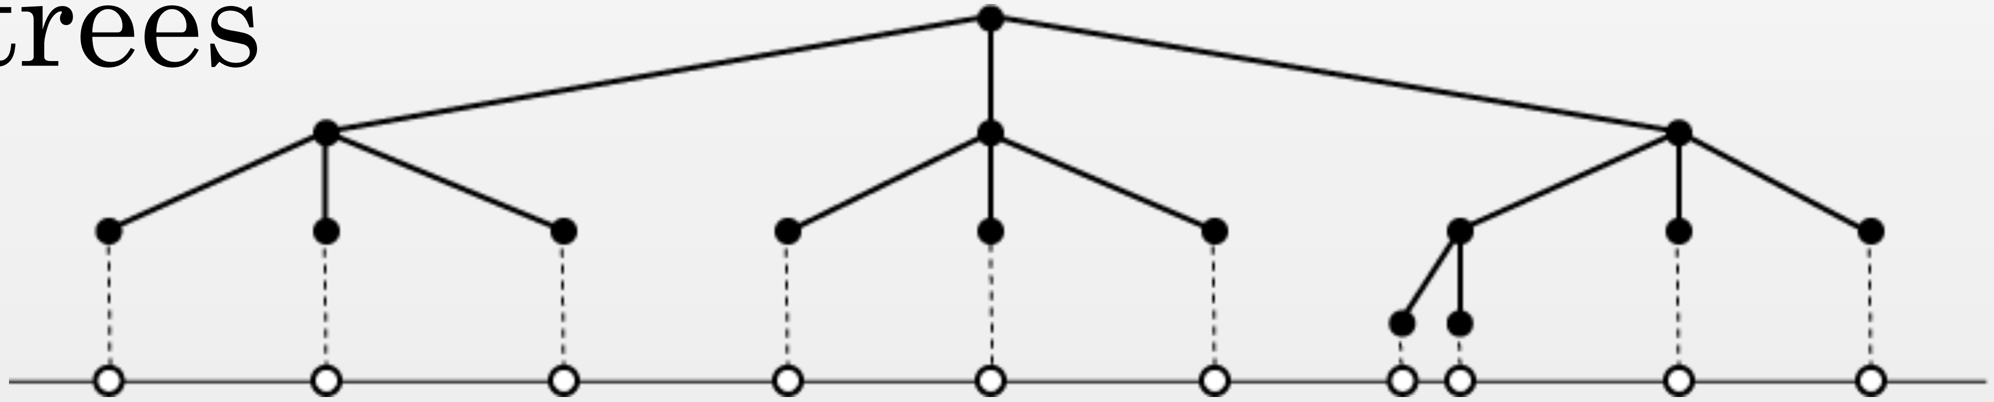

A generalization of quadtrees to metric spaces.
One leaf per point of P.
Each node u has a representative *rep(u)* in P and a radius *rad(p)*.

Net-trees

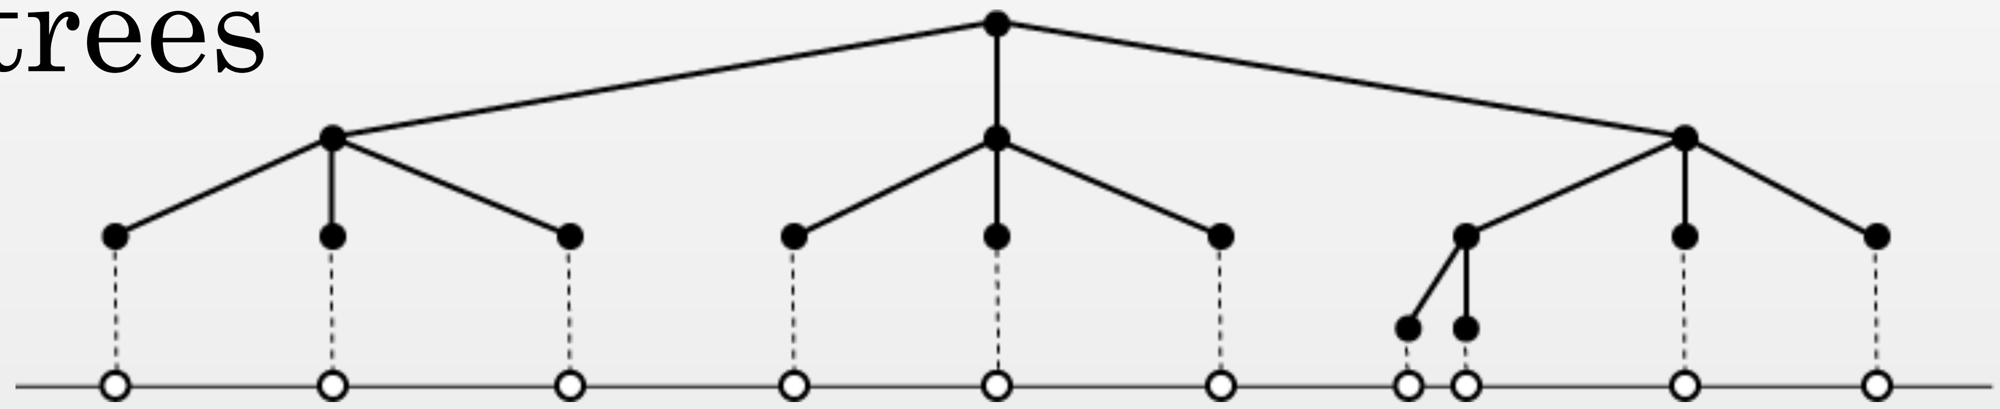

A generalization of quadtrees to metric spaces.
One leaf per point of P.
Each node u has a representative *rep(u)* in P and a radius *rad(p)*.
Three properties:

Net-trees

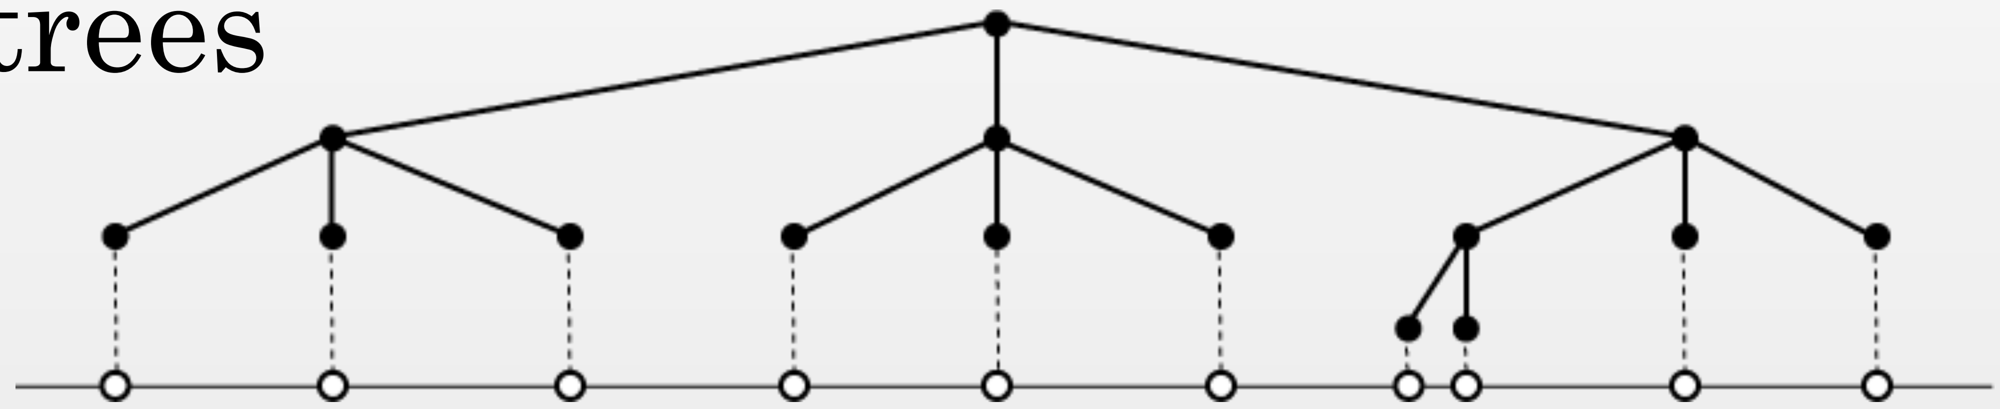

A generalization of quadtrees to metric spaces.
One leaf per point of P.
Each node u has a representative $rep(u)$ in P and a radius $rad(p)$.
Three properties:

1. Inheritance: Every nonleaf u has a child with the same rep.

Net-trees

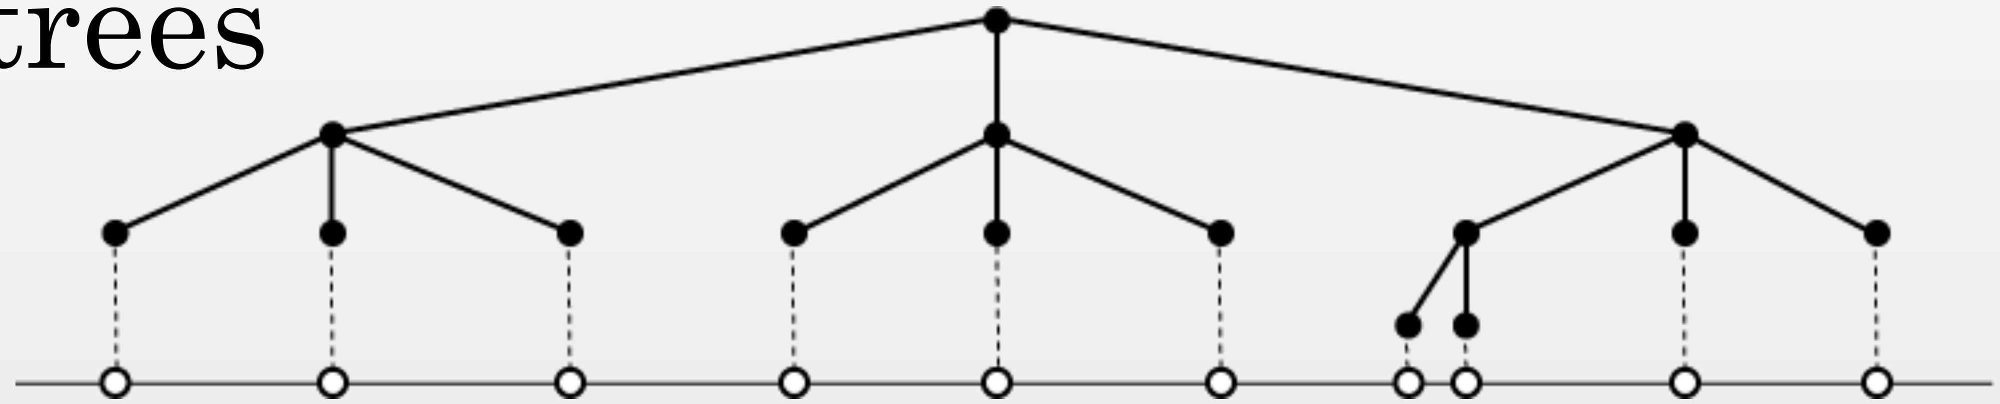

A generalization of quadtrees to metric spaces.
One leaf per point of P.
Each node u has a representative *rep(u)* in P and a radius *rad(p)*.
Three properties:

1. Inheritance: Every nonleaf u has a child with the same rep.
2. Covering: Every heir is within *ball(rep(u), rad(u))*

Net-trees

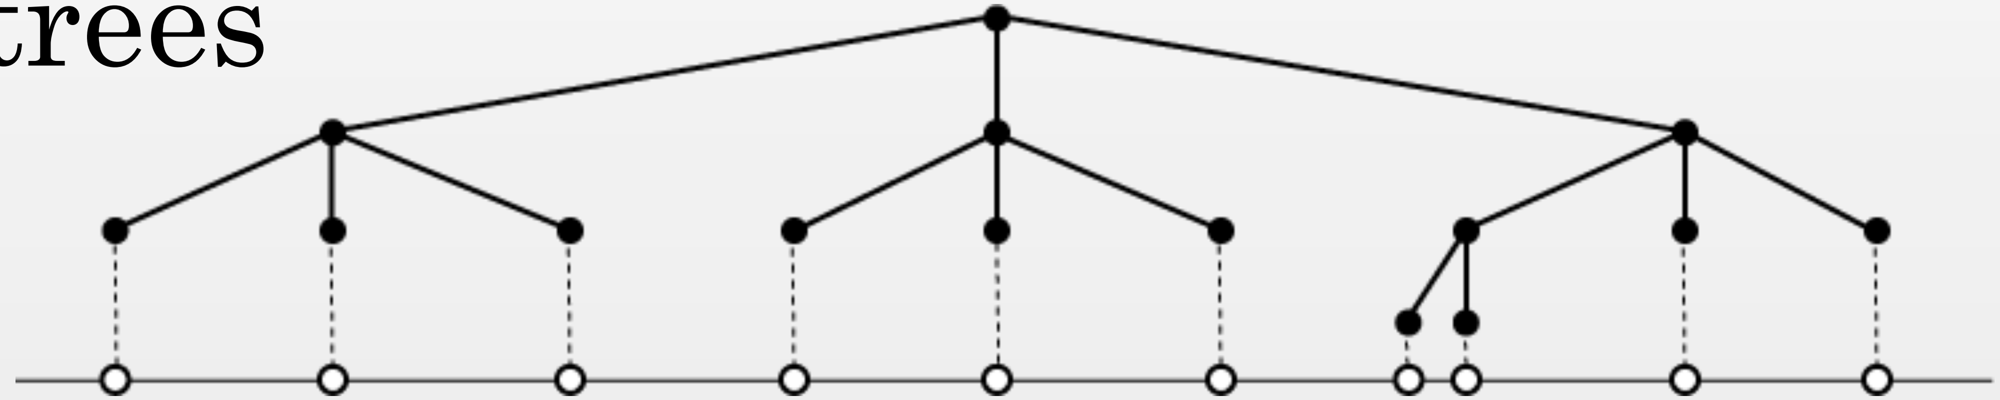

A generalization of quadtrees to metric spaces.
One leaf per point of P.
Each node u has a representative *rep(u)* in P and a radius *rad(p)*.
Three properties:

1. Inheritance: Every nonleaf u has a child with the same rep.
2. Covering: Every heir is within *ball(rep(u), rad(u))*
3. Packing: Any children v,w of u have $d(rep(v),rep(w)) > K\, rad(u)$

Net-trees

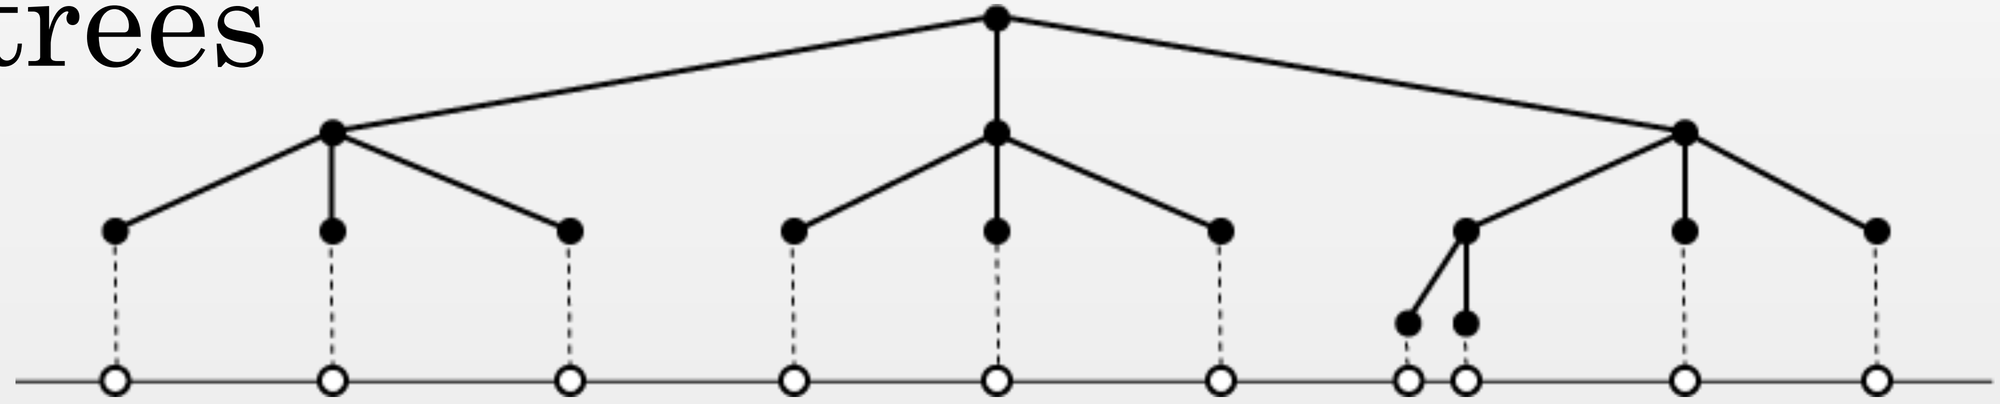

A generalization of quadtrees to metric spaces.
One leaf per point of P.
Each node u has a representative *rep(u)* in P and a radius *rad(p)*.
Three properties:

1. Inheritance: Every nonleaf u has a child with the same rep.
2. Covering: Every heir is within *ball(rep(u), rad(u))*
3. Packing: Any children v,w of u have *d(rep(v),rep(w))* > *K rad(u)*

Let u_p be the ancestor of all nodes represented by p.

Time to remove p: $t_p = \frac{1}{\varepsilon(1-\varepsilon)} rad(parent(u_p))$

Projection onto a net

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Q_α is the subcomplex of $\hat{R}_\alpha$ induced on N_α.

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Q_α is the subcomplex of $\hat{R}_\alpha$ induced on N_α.

$$\pi_\alpha(p) = \begin{cases} p & \text{if } p \in N_\alpha \\ \arg\min_{q \in N_\alpha} \hat{d}(p, q) & \text{otherwise} \end{cases}$$

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Q_α is the subcomplex of $\hat{R}_\alpha$ induced on N_α.

$$\pi_\alpha(p) = \begin{cases} p & \text{if } p \in N_\alpha \\ \arg\min_{q \in N_\alpha} \hat{d}(p, q) & \text{otherwise} \end{cases}$$

N_α is a Delone set.

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Q_α is the subcomplex of $\hat{R}_\alpha$ induced on N_α.

$$\pi_\alpha(p) = \begin{cases} p & \text{if } p \in N_\alpha \\ \arg\min_{q \in N_\alpha} \hat{d}(p, q) & \text{otherwise} \end{cases}$$

N_α is a Delone set.

1 **Covering:** For all $p \in P$, there is a $q \in N_\alpha$ such that $\mathbf{d}(p, q) \leq \varepsilon(1 - \varepsilon)\alpha$.

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Q_α is the subcomplex of $\hat{R}_\alpha$ induced on N_α.

$$\pi_\alpha(p) = \begin{cases} p & \text{if } p \in N_\alpha \\ \arg\min_{q \in N_\alpha} \hat{d}(p,q) & \text{otherwise} \end{cases}$$

N_α is a Delone set.

1. Covering: For all $p \in P$, there is a $q \in N_\alpha$ such that $\mathbf{d}(p,q) \leq \varepsilon(1-\varepsilon)\alpha$.

2. Packing: For all distinct $p, q \in N_\alpha$, $\mathbf{d}(p,q) \geq K_{pack}\varepsilon(1-\varepsilon)\alpha$.

Projection onto a net

$$N_\alpha = \{p \in P : t_p \geq \alpha\}$$

Q_α is the subcomplex of $\hat{R}_\alpha$ induced on N_α.

$$\pi_\alpha(p) = \begin{cases} p & \text{if } p \in N_\alpha \\ \arg\min_{q \in N_\alpha} \hat{d}(p,q) & \text{otherwise} \end{cases}$$

N_α is a Delone set.

1 Covering: For all $p \in P$, there is a $q \in N_\alpha$ such that $\mathbf{d}(p,q) \leq \varepsilon(1-\varepsilon)\alpha$.

2 Packing: For all distinct $p, q \in N_\alpha$, $\mathbf{d}(p,q) \geq K_{pack}\varepsilon(1-\varepsilon)\alpha$.

Key Fact: For all $p, q \in P$, $\hat{\mathbf{d}}(\pi_\alpha(p), q) \leq \hat{\mathbf{d}}(p,q)$.

Key Fact: For all $p, q \in P$, $\hat{\mathbf{d}}(\pi_\alpha(p), q) \leq \hat{\mathbf{d}}(p, q)$.

This can be used to show that projection onto a net is a homotopy equivalence.

$$\begin{array}{rccccccc}
\text{Relaxed Rips:} & \hookrightarrow & \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\
 & & \uparrow & & \uparrow & & \uparrow & \\
\text{Sparse Rips Zigzag:} & \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow
\end{array}$$

Key Fact: For all $p, q \in P$, $\hat{\mathbf{d}}(\pi_\alpha(p), q) \le \hat{\mathbf{d}}(p, q)$.

This can be used to show that projection onto a net is a homotopy equivalence.

$$
\begin{array}{rccccccc}
\text{Relaxed Rips:} & \hookrightarrow & \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\
 & & \uparrow & & \uparrow & & \uparrow & \\
\text{Sparse Rips Zigzag:} & \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow
\end{array}
$$

$$\leftarrow \mathrm{H}(Q_\alpha) \rightarrow \mathrm{H}(Q_\beta) \leftarrow \mathrm{H}(Q_\gamma) \rightarrow$$

Key Fact: For all $p, q \in P$, $\hat{\mathbf{d}}(\pi_\alpha(p), q) \leq \hat{\mathbf{d}}(p, q)$.

This can be used to show that projection onto a net is a homotopy equivalence.

$$
\begin{array}{lccccccc}
\text{Relaxed Rips:} & \hookrightarrow & \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\
 & & \uparrow & & \uparrow & & \uparrow & \\
\text{Sparse Rips Zigzag:} & \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow
\end{array}
$$

$$
\xleftarrow{\cong} \mathrm{H}(Q_\alpha) \to \mathrm{H}(Q_\beta) \xleftarrow{\cong} \mathrm{H}(Q_\gamma) \to
$$

Key Fact: For all $p, q \in P$, $\hat{\mathbf{d}}(\pi_\alpha(p), q) \leq \hat{\mathbf{d}}(p, q)$.

This can be used to show that projection onto a net is a homotopy equivalence.

$$\begin{array}{rccccccc}
\text{Relaxed Rips:} & \hookrightarrow & \hat{R}_\alpha & \hookrightarrow & \hat{R}_\beta & \hookrightarrow & \hat{R}_\gamma & \hookrightarrow \\
& & \uparrow & & \uparrow & & \uparrow & \\
\text{Sparse Rips Zigzag:} & \hookleftarrow & Q_\alpha & \hookrightarrow & Q_\beta & \hookleftarrow & Q_\gamma & \hookrightarrow
\end{array}$$

$$\boxed{\overset{\cong}{\to}} \mathrm{H}(Q_\alpha) \to \mathrm{H}(Q_\beta) \boxed{\overset{\cong}{\to}} \mathrm{H}(Q_\gamma) \to$$

Why is the filtration only linear size?

Why is the filtration only linear size?

Standard trick from Euclidean geometry:

Why is the filtration only linear size?

Standard trick from Euclidean geometry:

1. Charge each simplex to its vertex with the earliest deletion time.

Why is the filtration only linear size?

Standard trick from Euclidean geometry:

1. Charge each simplex to its vertex with the earliest deletion time.

2. Apply a packing argument to the larger neighbors.

Why is the filtration only linear size?

Standard trick from Euclidean geometry:

1. Charge each simplex to its vertex with the earliest deletion time.
2. Apply a packing argument to the larger neighbors.
3. Conclude average O(1) simplices per vertex.

Why is the filtration only linear size?

Standard trick from Euclidean geometry:

1 Charge each simplex to its vertex with the earliest deletion time.

2 Apply a packing argument to the larger neighbors.

3 Conclude average O(1) simplices per vertex.

$$\left(\frac{1}{\varepsilon}\right)^{O(d^2)}$$

Really getting rid of the zigzags.

Really getting rid of the zigzags.

$$X_\alpha = \bigcup_{\beta \leq \alpha} Q_\alpha$$

Really getting rid of the zigzags.

$$X_\alpha = \bigcup_{\beta \le \alpha} Q_\alpha$$

This is (almost) the clique complex of a hierarchical spanner!

Summary

Summary

Approximate the VR filtration with a Zigzag filtration.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Distance to a measure?

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Distance to a measure?
Other types of simplification.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Distance to a measure?
Other types of simplification.
Construct the net-tree in O(n log n) time.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Distance to a measure?
Other types of simplification.
Construct the net-tree in O(n log n) time.
An implementation.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Distance to a measure?
Other types of simplification.
Construct the net-tree in O(n log n) time.
An implementation.

Thank You.

Summary

Approximate the VR filtration with a Zigzag filtration.
Remove points using a hierarchical net-tree.
Perturb the metric to straighten out the zigzags.
Use packing arguments to get linear size.
Union trick eliminates zigzag.

The Future

Distance to a measure?
Other types of simplification.
Construct the net-tree in O(n log n) time.
An implementation.

Thank You.