Course 8: Geometry and Neural Networks

Eric Goubault & Sylvie Putot
MPRI

17th February 2026

Contents of the course

1 Motivation for this course

2 Persistent homology
- Motivation
- Barcodes
- Efficient calculation of unidimensional persistent homology

3 From data to simplicial sets/complexes
- Rips and Cech complexes
- Persistence modules and distances
- Interleaving distance and stability
- Sample calculation of the interleaving distance

4 Homotopy type for neural nets and classifiers
- Pavings and homology
- Nerves and the nerve lemma
- Robustness of a neural network

Neural network classification

The “shape of classes”

A neural network used as a classifier, trained over a set of labelled data, is going to be of quality if it “recognizes” the shapes of the different input classes.

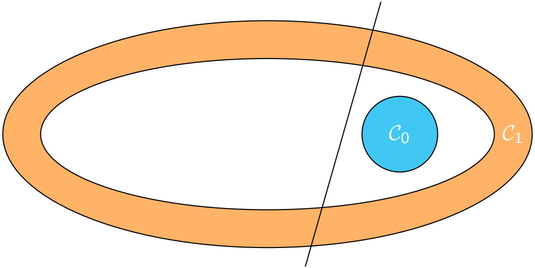

Here: two nested classes with a linear separator; about 80% of the samples will be classified correctly, but we clearly missed the point here.

More in depth: actual training phases!

Complexity of classification problems

(already a glimpse on that in the last example)

Example: complexity of data classification

When the classes labeled data has "complex" shapes, this has to reflect into the complexity of any classifier (e.g. neural network).

Suppose for now the data forms a "real" topological space (not just a discrete set of points). Then having m connected components means we need at least an algorithm with complexity $O(\log\ m)$ to classify them ("set membership problem" using algebraic decision trees)

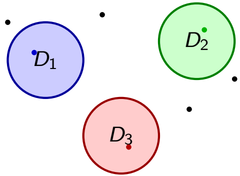

X together with points sampled in the ambient space

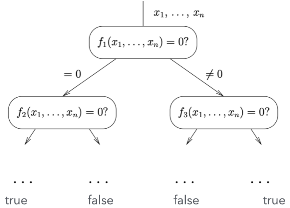

Corresponding decision tree

Topological adequacy/robustness: what we are going to do

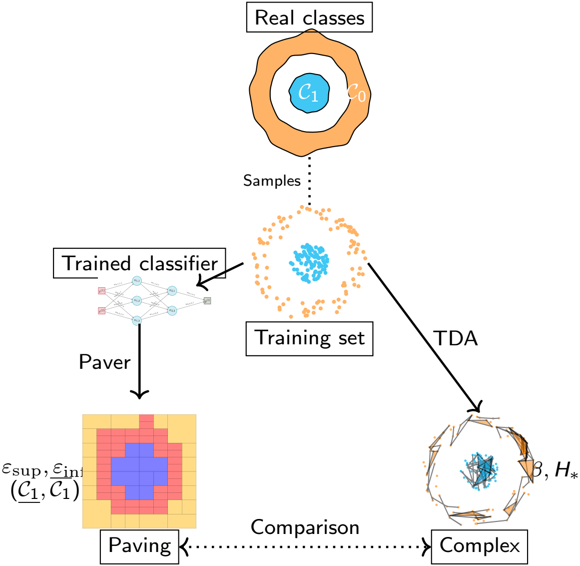

TDA/Persistent homology?

Given a finite set of points sampled from an unknown space:

- How many connected components are there?
- When do loops appear?
- Which features persist across scales?

Persistent homology answers these questions by tracking homology across a filtration.

Point clouds and scale

Example

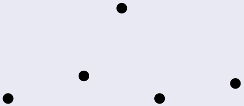

- At very small scale: each point is isolated.
- As the scale increases, points begin to connect.

E.g. Vietoris–Rips filtration

For a scale parameter $\varepsilon > 0$:

- Vertices: the data points
- Edges: connect points at distance $\leq \varepsilon$
- Triangles: fill in when all edges are present

This gives a nested sequence:

$$R_{\varepsilon_0} \subset R_{\varepsilon_1} \subset R_{\varepsilon_2} \subset \cdots$$

We compute homology at each stage.

Chain complex in dimension 1

We focus on:

$$C_1 \xrightarrow{\partial_1} C_0$$

where:

- C_0: free abelian group on vertices
- C_1: free abelian group on edges

Homology:

$$H_0 = \ker 0 / \operatorname{im} \partial_1$$

Tracks connected components.

Boundary matrix

Order vertices $v_1, \ldots, v_5$ and edges $e_1, \ldots, e_4$.

$$\partial_1 = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

This matrix evolves as ε increases.

Smith normal form

Over $\mathbb{Z}$, any integer matrix can be diagonalized:

$$UAV = \mathrm{diag}(d_1, \ldots, d_r, 0, \ldots, 0)$$

with $d_i \mid d_{i+1}$.

For boundary matrices:

- Zero columns $\leftrightarrow$ births
- Pivot entries $\leftrightarrow$ deaths

Birth and death of components

Long bars correspond to persistent features.

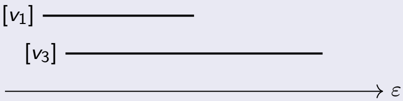

Persistent homology in degree 1

When loops appear (three edges forming a cycle):

- A class is born in H_1
- It dies when the triangle fills in

This is encoded in the boundary map:

$$C_2 \xrightarrow{\partial_2} C_1$$

and its Smith normal form.

Barcode summary

Persistent homology separates signal from noise

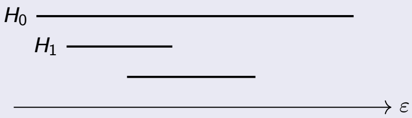

Barcode example

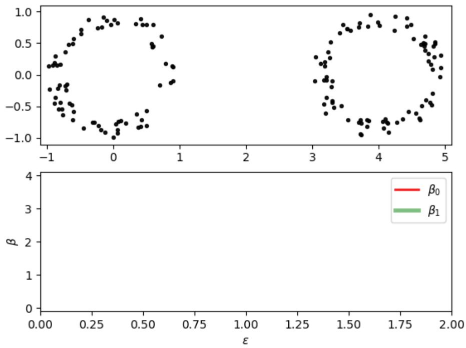

Barcode example

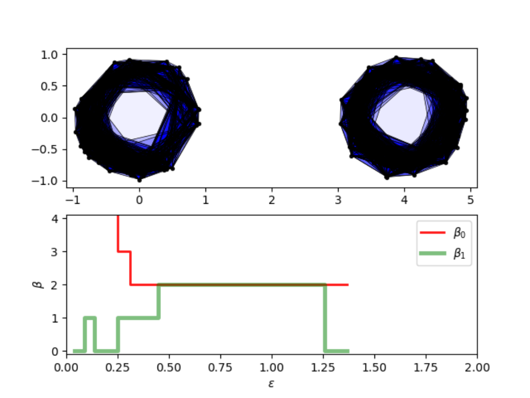

How to compute things efficiently?

The method we have seen (computing homology at each “time step” in the filtration) is not practical, let us step back for a minute...

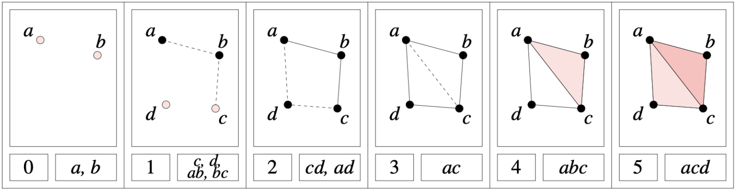

Filtered complex

- Time 0: 2 points, so 2 components
- Time 1: add 2 more points, but new edges merge components: still 2 components
- Time 2: add 2 more edge, merging components, but creating a hole etc.

Representation

How to represent a unidimensional filtration in one go?

Idea: z-transform! (algebraically, this is linked to quiver representation theory)

- Encode the boundary matrices in the ring $\mathbb{F}[t]$ of polynomials over a field $\mathbb{F}$,
- a, $t.a, \ldots, t^n.a$ are respectively simplices a at time 0, $1, \ldots, n$ plus the time of their appearance,
- this is some form of "time delay" operator.

When $\mathbb{F}$ is a field, the ring of univariate polynomials $\mathbb{F}[t]$ is a Euclidean ring (so in particular a PID - Principal Ideal Domain) in which we can compute the Smith Normal Form of matrices.

See e.g. Zomorodian, A., Carlsson, G. "Computing Persistent Homology", Discrete Comput Geom 33, (2005).

Computation of the persistence modules

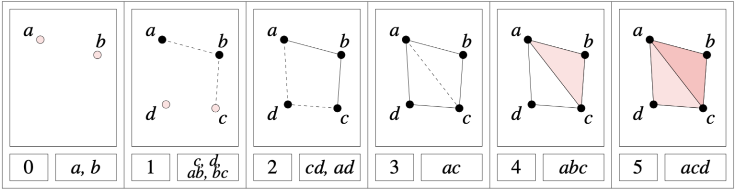

Matrix for ∂_1

$$M_1 = \left(\begin{array}{c|ccccc} & ab & bc & cd & ad & ac \\ \hline d & 0 & 0 & t & t & 0 \\ c & 0 & 1 & t & 0 & t^2 \\ b & t & t & 0 & 0 & 0 \\ a & t & 0 & 0 & t^2 & t^3 \end{array}\right)$$

Computation of the persistence modules

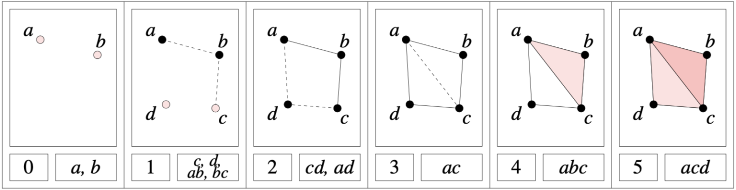

Matrix for ∂_2

$$M_2 = \left(\begin{array}{c|cc} & abc & acd \\ ac & 0 & t^2 \\ ad & 0 & t^3 \\ cd & 0 & t^3 \\ bc & t^3 & 0 \\ ab & t^3 & 0 \end{array} \right)$$

Row-echelon and column-echelon forms of matrices

Row-echelon

- All rows having only zero entries are at the bottom,
- The leading entry (that is, the leftmost non-zero entry) of every non-zero row, called the pivot, is to the right of the leading entry of every row above.

Obtained by using only the row operations of the SNF algorithm (enough for computing kernels).

Column-echelon

- Similarly, when the transpose is in row-echelon form.

Obtained by using only the column operations of the SNF algorithm (enough for computing images).

Matrices in row-echelon and column-echelon form

M_1 in column-echelon form

$$\tilde{M}_1 = \left(\begin{array}{c|ccccc} & cd & bc & ab & z_1 & z_2 \\ \hline d & t & 0 & 0 & 0 & 0 \\ c & t & 1 & 0 & 0 & 0 \\ b & 0 & t & t & 0 & 0 \\ a & 0 & 0 & t & 0 & 0 \end{array}\right)$$

with $z_1 = ad - cd - t.bc - t.ab$ and $z_2 = ac - t^2.bc - t^2.ab$.

M_2 in row-echelon form

$$\tilde{M}_2 = \left(\begin{array}{c|cc} & abc & acd \\ \hline z_2 & t & t^2 \\ z_1 & 0 & t^3 \end{array}\right)$$

Cech complexes

Cech complex

Given a *metric* space X, and $V \subseteq X$, and $\varepsilon > 0$, the Čech complex associated to V and ε is the nerve of the covering $\{B(v,\varepsilon) \mid v \in V\}$:

$$\check{C}(V,\varepsilon) = \mathcal{N}\left(\{B(v,\varepsilon) \mid v \in V\}\right)$$

Intuitively, according to the nerve lemma , if ε is small enough and the elements of V are spaced out enough in X to "capture" X, then $\check{C}(V,\varepsilon)$ has the same homotopy as X.

Vietoris-Rips complexes

A variant: Vietoris-Rips complex

The Vietoris-Rips complex associated to ε, noted $VR(X, \varepsilon)$, is the simplicial complex:

- with elements of X as vertices,
- and $\{x_1, \ldots, x_k\}$ is a k-face if all of its elements are distant from less than ε:

$$\forall i, j, \mathbf{d}(x_i, x_j) \leq \varepsilon$$

This method is somewhat less practical in practice because the complex is in general much larger.

Note that

$$\check{C}(X, \varepsilon) \subseteq VR(X, 2\varepsilon) \subseteq \check{C}(X, 2\varepsilon)$$

Formally: Persistence modules

A real persistence module $\mathbb{V}$ contains:

- objects V_t indexed by a real parameter t
- morphisms $v_b^a : V_a \to V_b$ whenever $a \leq b$
- for all t, v_t^t is an isomorphism, and $v_c^b \circ v_b^a = v_c^a$ for adapted a, b, c

To summarise, an $\mathbb{R}$-persistence module is a functor from $(\mathbb{R}, \leq)$ to the category of vector spaces.

More generally, $\mathcal{P}$-persistence modules are defined similarly for any poset $\mathcal{P}$, the most common case being **Z**-persistence modules.

Interleaving distance

Interleaving distance measures how far from being isomorphic the persistence modules are.

Interleavings

Let $\mathbb{V}$ and $\mathbb{W}$ be two $\mathbb{R}$-persistence modules. A δ-interleaving is a pair of family of morphisms $\{\phi_t : V_t \to W_{t+\delta} \ s.t. \ t \in \mathbb{R}\}$ and $\{\psi_t : W_t \to V_{t+\delta} \ s.t. \ t \in \mathbb{R}\}$ such that the following diagrams commute for all $a, b \in \mathbb{R}$:

$$\begin{array}{ccccc} V_a & & \xrightarrow{v_b^a} & & V_b \\ & \searrow^{\phi_a} & & & & \searrow^{\phi_b} \\ & & W_{a+\delta} & \xrightarrow{w_{b+\delta}^{a+\delta}} & & W_{b+\delta} \end{array} \qquad \begin{array}{ccccc} W_a & & \xrightarrow{w_b^a} & & W_b \\ & \searrow^{\psi_a} & & & & \searrow^{\psi_b} \\ & & V_{a+\delta} & \xrightarrow{v_{b+\delta}^{a+\delta}} & & V_{b+\delta} \end{array}$$

$$\begin{array}{ccc} W_{a-\delta} & \xrightarrow{w_{a+\delta}^{a-\delta}} & W_{a+\delta} \\ \searrow^{\psi_{a-\delta}} & & \nearrow_{\phi_a} \\ & V_a & \end{array} \qquad \begin{array}{ccc} V_{a-\delta} & \xrightarrow{v_{a+\delta}^{a-\delta}} & V_{a+\delta} \\ \searrow^{\phi_{a-\delta}} & & \nearrow_{\psi_a} \\ & W_a & \end{array}$$

Interleaving distance

Interleaving distance measures how far from being isomorphic the persistence modules are.

Interleaving distance

The interleaving distance $\mathbf{d}_I$ between $\mathbb{V}$ and $\mathbb{W}$ is an element of $\mathbb{R}_+ \cup \{+\infty\}$ defined by

$$\mathbf{d}_I(\mathbb{V}, \mathbb{W}) = \inf_{\delta > 0} \{\exists\ \delta\text{-interleaving between } \mathbb{V} \text{ and } \mathbb{W}\}$$

Hausdorff and Gromov-Hausdorff distances

Hausdorff distance

$\mathbf{d}_H$: Hausdorff distance between sets in a metric space

$$\begin{aligned}\mathbf{d}_H(A,B) &\stackrel{\text{def}}{=} \max\left\{\sup_{a\in A}\mathbf{d}(a,B), \sup_{b\in B}\mathbf{d}(A,b)\right\}\\ &= \max\left\{\sup_{a\in A}\inf_{b\in B}\mathbf{d}(a,b), \sup_{b\in B}\inf_{a\in A}\mathbf{d}(a,b)\right\}\\ &= \inf\left\{\varepsilon>0 \mid \begin{cases} A & \subseteq & B\oplus\mathbf{B}(0,\varepsilon)\\ B & \subseteq & A\oplus\mathbf{B}(0,\varepsilon)\end{cases}\right\}\end{aligned}$$

Gromov-Hausdorff distance

$\mathbf{d}_{GH}$: Gromov-Hausdorff distance between metric spaces:

$$\mathbf{d}_{GH}(A,B) \stackrel{\text{def}}{=} \inf_{f,g}\{\mathbf{d}_H(f(A),g(B))\}$$

for $f : A \to Z$ and $g : B \to Z$ isometric embedding in the same metric space Z.

Comparison theorems

Correspondence

A *correspondence* $C : X \rightrightarrows Y$ is a surjective multi-valued map from X to Y, in other words:

- a subset C of $X \times Y$ such that $\pi_{X_{|C}} : C \to X$ and $\pi_{Y_{|C}} : C \to Y$ are both surjective,
- with $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ the canonical projections.

Distortion of a correspondance C

Its distorsion is $\mathrm{dist}(C) \stackrel{\text{def}}{=} \sup_{(x,y),(x',y') \in C} |\mathbf{d}_X(x,x') - \mathbf{d}_Y(y,y')|$.

Also, note the following:

Partial monotony of the Gromov-Hausdorff distance

For all $X \subseteq \mathbb{R}^n$, the function $\mathbf{d}_{GH}(X, \cdot) : \left| \begin{array}{ccc} (\{A \mid X \subseteq A\}, \subseteq) & \longrightarrow & (\mathbb{R}, \leq) \\ Y & \longmapsto & \mathbf{d}_{GH}(X, Y) \end{array} \right.$ is increasing.

Comparison theorems

Gromov-Hausdorff distance expressed in terms of distorsion

$$\mathbf{d}_{GH}(X,Y) = \frac{1}{2} \inf_{C:X \rightrightarrows Y} \operatorname{dist}(C)$$

Proof

With $A \subseteq B$, let's show $\mathbf{d}_{GH}(X,A) \leq \mathbf{d}_{GH}(X,B)$, in other words

$$\inf_{C:X \rightrightarrows A} \operatorname{dist}(C) \leq \inf_{C:X \rightrightarrows B} \operatorname{dist}(C)$$

To each correspondence from X to B corresponds its restriction from X to A that is also a correspondence. So for any fixed $C^* : X \rightrightarrows B$:

$$\sup_{(x,y),(x',y') \in C^*_{|A}} \left| \mathbf{d}_X(x,x') - \mathbf{d}_B(y,y') \right|$$
$$\leq \sup_{(x,y),(x',y') \in C^*} \left| \mathbf{d}_X(x,x') - \mathbf{d}_B(y,y') \right|$$

This inequality is precisely equivalent to

$$\operatorname{dist}(C^*_{|A}) \leq \operatorname{dist}(C^*)$$

Stability theorem

The most important part for our application!

Stability theorem

Let $(X, \mathbf{d}_X)$ and $(Y, \mathbf{d}_Y)$ be two metric spaces, and $C(X)$ (resp. $C(Y)$) the persistence module associated to the Cech or Rips complex of X (resp. Y). Then

$$\mathbf{d}_I(C(X), C(Y)) \leq 2\mathbf{d}_{GH}(X, Y)$$

Means that the "geometric signature" can serve a (hopefully more tractable) way to find lower-bounds for Gromov-Hausdorff distances, which are very difficult to calculate.

Comparison of the persistence homologies of the Cech simplices of two triangles

Triangles

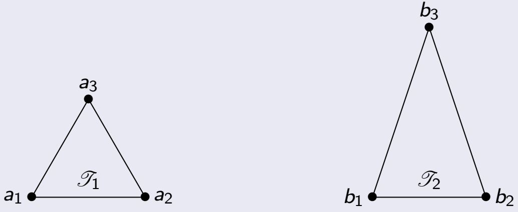

If we connect the dots, the first case is an equilateral triangle, with sizes 1. The second is an isosceles triangle, with sizes 1, 2 and 2.

Persistence diagrams

For the equilateral triangle

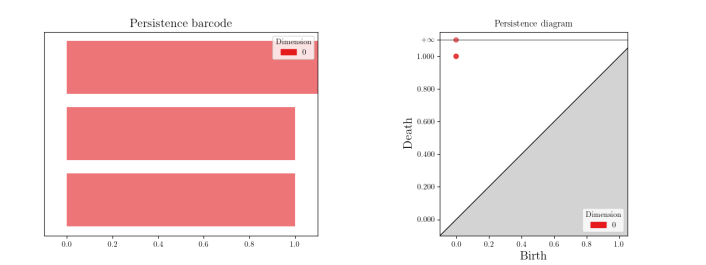

Persistence diagrams

For the isosceles triangle

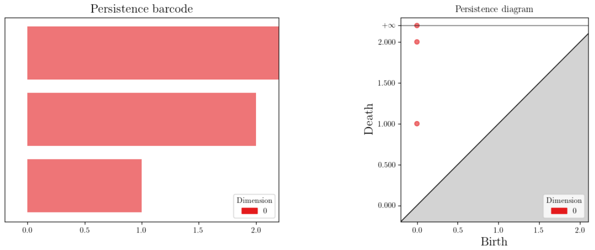

Interleaving distance

Principle

- To compute the interleaving distance, we must look closely at the morphisms of a persistence module,
- The idea is to map each generator to what it becomes in the future,
- The considered bases for matrix representations are (a_1, a_2, a_3) for $\mathscr{T}_1$ and (b_1, b_2, b_3) for $\mathscr{T}_2$.

In the case of $\mathscr{T}_1$, all the generators a_1, a_2 and a_3 merge at time $t = \frac{1}{2}$.

For $\mathscr{T}_2$, b_1 and b_2 merge at time $t = \frac{1}{2}$, and then into b_3 at time $t = 1$. Between these times, nothing changes.

Interleaving distance

Morphisms of $\mathscr{T}_1$

$$\begin{array}{ll} 0 \leq t < \frac{1}{2} & u_t^t = \mathrm{id}_{\mathbb{R}^3} \\ \frac{1}{2} \leq t < \infty & u_t^t = \mathrm{id}_{\mathbb{R}^1} \end{array}$$

and for

$$\begin{array}{lcl} 0 \leq t_1 \leq t_2 < \frac{1}{2} & u_{t_2}^{t_1} = & \mathrm{id}_{\mathbb{R}^3} \\ 0 \leq t_1 < \frac{1}{2} \leq t_2 & u_{t_2}^{t_1} = & \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} \\ \frac{1}{2} < t_1 \leq t_2 & u_{t_2}^{t_1} = & \mathrm{id}_{\mathbb{R}^1} \end{array}$$

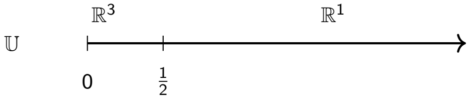

Interleaving distance

Morphisms of $\mathscr{T}_2$

$$\begin{array}{cc} 0 \leq t < \frac{1}{2} & v_t^t = \mathrm{id}_{\mathbb{R}^3} \\ \frac{1}{2} \leq t < a & v_t^t = \mathrm{id}_{\mathbb{R}^2} \\ a \leq t < \infty & v_t^t = \mathrm{id}_{\mathbb{R}^1} \end{array}$$

and for $0 \leq t_1 < \dfrac{1}{2} \leq t_2 < 1 \leq t_3$

$$\begin{array}{ccc} v_{t_2}^{t_1} & = & \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ v_{t_3}^{t_2} & = & \begin{bmatrix} 1 & 1 \end{bmatrix} \\ v_{t_3}^{t_1} & = & \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} \end{array}$$

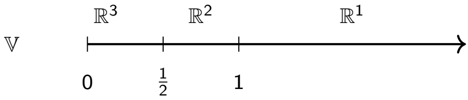

Interleaving distance

$d_I \leq \frac{1}{2}$

$\mathscr{T}_1$ and $\mathscr{T}_2$ are $\frac{1}{2}-$interleaved. We can set

- $\phi_t : \mathbb{U}_t \to \mathbb{V}_{t+\frac{1}{2}}, x \mapsto \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} x$, for $t \in [0, \frac{1}{2}[$, and then $\phi_t : \mathbb{U}_t \to \mathbb{V}_{t+\frac{1}{2}}, x \mapsto \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} x$, for $\frac{1}{2} \leq t < 1$ and $\phi_t : \mathbb{U}_t \to \mathbb{V}_{t+\frac{1}{2}}, x \mapsto x$, for $1 \leq t$;
- $\psi_t : \mathbb{V}_t \to \mathbb{U}_{t+\frac{1}{2}}, x \mapsto \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} x$, for $t \in [0, \frac{1}{2}[$, $\psi_t : \mathbb{V}_t \to \mathbb{U}_{t+\frac{1}{2}}, x \mapsto \begin{bmatrix} 1 & 1 \end{bmatrix} x$, for $t \in [\frac{1}{2}, 1[$ and then $\psi_t : \mathbb{V}_t \to \mathbb{U}_{t+\frac{1}{2}}, x \mapsto x$, for $t \geq 1$.

Interleaving distance

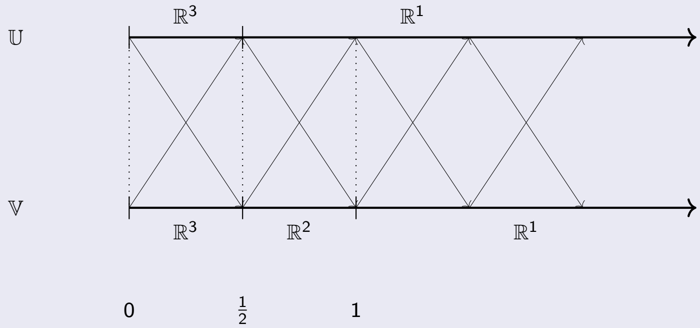

Interleaving distance

$d_I = \frac{1}{2}$

$\mathscr{T}_1$ and $\mathscr{T}_2$ are not $(\frac{1}{2} - \varepsilon)-$interleaved for $\varepsilon > 0$. If they were, then

- $\psi_{\frac{3\varepsilon}{2}} : \mathbb{V}_{\frac{3\varepsilon}{2}} \to \mathbb{U}_{\frac{1}{2}+\frac{\varepsilon}{2}}$
- $\phi_{\frac{1}{2}+\frac{\varepsilon}{2}} : \mathbb{U}_{\frac{1}{2}+\frac{\varepsilon}{2}} \to \mathbb{V}_{1-\frac{\varepsilon}{2}}$

must satisfy $\phi_{\frac{1}{2}+\frac{\varepsilon}{2}} \circ \psi_{\frac{3\varepsilon}{2}} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$, but that is impossible since $\psi_{\frac{3\varepsilon}{2}}$ has rank 1.

Interleaving distance

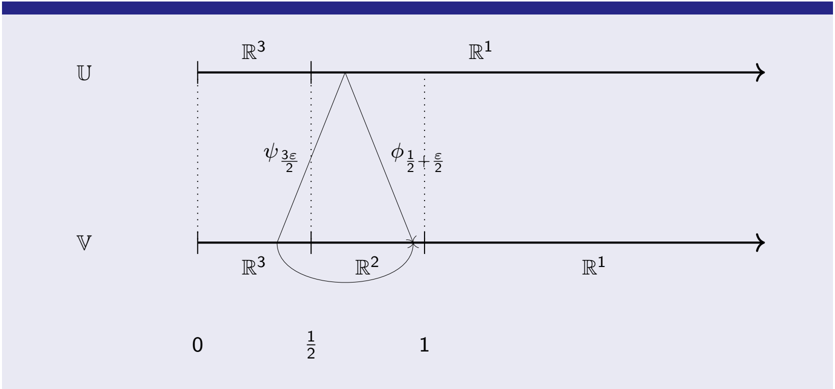

So

$$\mathbf{d}_I(\mathscr{T}_1, \mathscr{T}_2) = \frac{1}{2}$$

Pavings

Rectangles

A *rectangle* of $\mathbb{R}^n$ is a product of n intervals $I_1 \times \cdots \times I_n \subseteq \mathbb{R}^n$.

Pavings

A *paving* of $\mathbb{R}^n$ is a set of axis-aligned rectangles.

Pav is the set of pavings, and $(\mathrm{Pav}, \subseteq)$ is a poset, where we note $P_1 \subseteq P_2$ if $\bigcup_{r \in P_1} r \subseteq \bigcup_{r' \in P_2} r'$.

A pair of pavings approximate a real set.

Approximation of sets by pavings

Concretisation function

A pair of pavings is an abstraction for a set of point clouds, hence the following concretisation function γ:

$$\gamma : \left| \begin{array}{ccc} (\mathrm{Pav}^2, \subseteq) & \longrightarrow & (\mathbf{P}(\mathbf{P}(\mathbb{R}^n)), \subseteq) \\ (\underline{P}, \overline{P}) & \longmapsto & \{ Y \subseteq \mathbb{R}^n \mid \underline{P} \subseteq Y \subseteq \overline{P} \} \end{array} \right.$$

A cloud point S is *compatible* with $(\underline{P}, \overline{P})$ iff $S \in \gamma(\underline{P}, \overline{P})$.

Approximation of classes classified by a NN

Pavings for NN^{-1}

Require: Classifier NN : $\mathbb{R}^n \to \{0,1\}$, hyper-rectangle of finite volume $H \subset \mathbb{R}^n$, real number $\eta > 0$

Ensure: $\mathbf{P} = (P_0, P_1, P_\perp)$ such that $P_0 \cup P_1 \cup P_\perp$ is a partition of H in hyperrectangles such $\forall h \in P_i, \text{NN}(P_i) = i$, and $\forall h \in P_\perp, \texttt{image}(\text{NN}, h) = \perp \wedge \mu(h) \leq \eta$

$P_0, P_1, P_\perp \leftarrow \varnothing$
$Q \leftarrow \{H\}$ ▷ Q is the priority queue of the uncertified zones
while Q is not empty and the measure of its biggest element is $> 2\eta$ **do**
  $h \leftarrow \text{dequeue}(Q)$ ▷ h has the biggest volume among Q
  if $\texttt{image}(\text{NN}, h) = \perp$ **then**
    $(h_1, h_2) \leftarrow \texttt{cut}(h)$
    enqueue h_1 in Q with priority $\mu(h_1)$
    enqueue h_2 in Q with priority $\mu(h_2)$
  else if $\texttt{image}(\text{NN}, h) = 0$ **then**
    $P_0 \leftarrow P_0 \cup \{h\}$
  else if $\texttt{image}(\text{NN}, h) = 1$ **then**
    $P_1 \leftarrow P_1 \cup \{h\}$
  end if
end while
$P_\perp \leftarrow Q$ ▷ $P_\perp$ gets all the elements of Q
$\mathbf{P} \leftarrow (P_0, P_1, P_\perp)$

Simple experiment

Dataset

A dataset consisting of 500 points is generated from the following classes:

$$\mathcal{C}_0 = \{x \in [-1,1]^2 \mid 3/4 \leq \|x\|_2\} \quad \text{and} \quad \mathcal{C}_1 = \{x \in [-1,1]^2 \mid \|x\|_\infty \leq 1/2\}$$

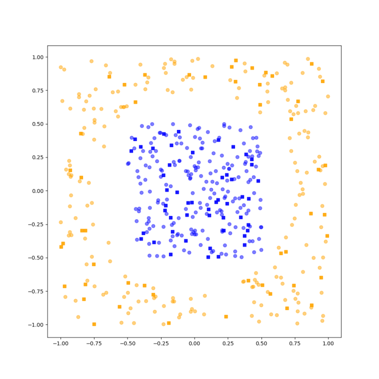

Example

Examples of pavings

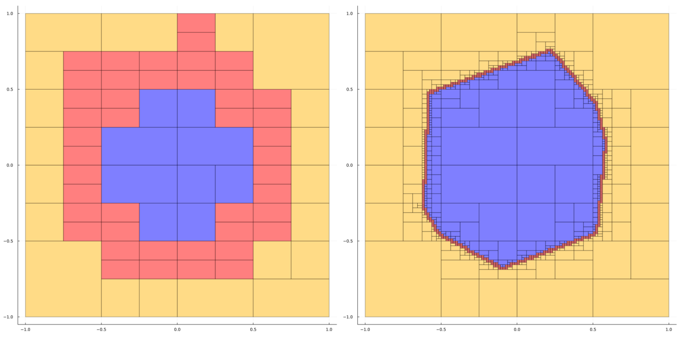

The Nerve of a Cover

How do we get the homology of such a space, decomposed in smaller chunks? Could use cubical complexes as for Conley, but more generally: nerve!!

Definition

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of a space X. The nerve $N(\mathcal{U})$ is the simplicial complex with

- a vertex v_i for each U_i,
- an k-simplex $[v_{i_0}, \ldots, v_{i_k}]$ whenever $U_{i_0} \cap \cdots \cap U_{i_k} \neq \emptyset$.

Good Covers

Definition

A cover $\mathcal{U}$ is a *good cover* if every finite nonempty intersection $U_{i_0} \cap \cdots \cap U_{i_k}$ is contractible (or empty).

Nerve Lemma:

If $\mathcal{U}$ is a good open cover of X (and X is covered by $\mathcal{U}$), then

$$X \simeq |N(\mathcal{U})|,$$

i.e. X is homotopy equivalent to the geometric realization of the nerve. Consequently

$$H_*(X) \cong H_*(N(\mathcal{U})).$$

Calculation of first homology group of S^1

Cover of S^1 by three open arcs

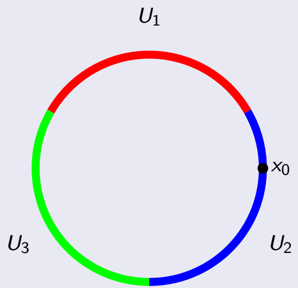

More explicitely

With the open cover of S^1 made up of 3 open arcs

The nerve $N(\{U_1, U_2, U_3\})$ is the 1-dimensional simplicial complex:

- with vertices v_1, v_2, v_3,
- and edges v_1v_2, v_2v_3, v_3v_1 (no triangle).

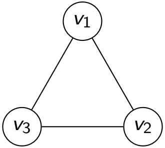

Simplicial Homology of the 3-cycle

Recap

Simplicial complex: vertices v_1, v_2, v_3, edges e_{12}, e_{23}, e_{31}, no 2-simplices.

Chain complex

Chain groups (with $\mathbb{Z}$ coefficients):

$$C_1 = \mathbb{Z}\langle e_{12}, e_{23}, e_{31}\rangle \cong \mathbb{Z}^3, \qquad C_0 = \mathbb{Z}\langle v_1, v_2, v_3\rangle \cong \mathbb{Z}^3.$$

Boundary map $\partial_1 : C_1 \to C_0$ (orientation chosen):

$$\begin{aligned} \partial_1(e_{12}) &= v_2 - v_1, \\ \partial_1(e_{23}) &= v_3 - v_2, \\ \partial_1(e_{31}) &= v_1 - v_3. \end{aligned}$$

Calculation of the first homology group

There are no 2-simplices, so $\operatorname{im} \partial_2 = 0$. Thus

$$H_1 = \ker \partial_1, \qquad H_0 = co\ker \partial_1.$$

Compute H_1 and H_0 for the 3-cycle

Calculation of the first homology group (continued)

Solve $\partial_1(ae_{12} + be_{23} + ce_{31}) = 0$:

$$a(v_2 - v_1) + b(v_3 - v_2) + c(v_1 - v_3) = 0.$$

Collect coefficients on v_1, v_2, v_3:

$$(-a + c)v_1 + (a - b)v_2 + (b - c)v_3 = 0.$$

This yields the linear system:

$$-a + c = 0, \quad a - b = 0, \quad b - c = 0.$$

So $a = b = c$ and elements of $\ker \partial_1$ are integer multiples of $e_{12} + e_{23} + e_{31}$. Therefore

$$H_1 \cong \mathbb{Z}.$$

Calculation of the 0th homology group (continued)

For H_0: since the graph is connected, $H_0 \cong \mathbb{Z}$. Higher homology vanishes.

Thus $H_*(S^1) \cong H_*(N(\mathcal{U}))$, as expected.

Nerve of a paving

Nerve of a paving P

Let $P = \left\{ r^{(j)} \;\middle|\; j \in J \right\}$:

- define $P_\epsilon \stackrel{\text{def}}{=} \left\{ r_\epsilon^{(j)} \;\middle|\; j \in J \right\}$.
- define its *nerve* $\mathcal{N}(P_\epsilon)$ as usual: it is the simplicial complex whose vertices are the elements of J, edges the $(j_0, j_1) \in P_\epsilon^2$ such that $r_\epsilon^{(j_0)} \cap r_\epsilon^{(j_1)} \neq \varnothing$,
- and more generally

$$\Sigma_k = \left\{ S \subseteq J, |J| = k, \bigcap_{s \in S} r_\epsilon^{(s)} \neq \varnothing \right\}$$

Lower bound of the robustness

Comparison theorem

Let X_1 be the subset of the training set for neural net NN, which should be classified as class 1. If $(\underline{P}, \overline{P})$ is an abstraction for $\mathrm{NN}^{-1}(1)$ with $\mathbf{d}_H(\underline{P}, \overline{P}) \leq \varepsilon$, if $Y_1 \in \gamma(\underline{P}, \overline{P})$:

$$2\mathbf{d}_{GH}(X_1, Y_1) \geq \mathbf{d}_I(X_1, \underline{P}) - 2\varepsilon$$

Proof

From stability and triangular inequality, we get

$$2\mathbf{d}_{GH}(X_1, Y) \geq \mathbf{d}_I(X_1, Y) \geq \mathbf{d}_I(X_1, \underline{P}) - \mathbf{d}_I(Y, \underline{P})$$
$$2\mathbf{d}_{GH}(Y, \underline{P}) \geq \mathbf{d}_I(Y, \underline{P})$$

By combining:

$$2\mathbf{d}_{GH}(X_1, Y) \geq \mathbf{d}_I(X, \underline{P}) - 2\mathbf{d}_{GH}(Y, \underline{P})$$

From the partial monotony of d_{GH}, we get $\mathbf{d}_{GH}(Y, \underline{P}) \leq \mathbf{d}_{GH}(\overline{P}, \underline{P}) = \varepsilon$, and so

$$2\mathbf{d}_{GH}(X_1, Y) \geq \mathbf{d}_I(X, \underline{P}) - 2\varepsilon$$

Lower and upper bounds of the classification difference of two neural nets

Comparison theorem

Let NN_1 and NN_2 two classifiers for the same classification problem, $(\underline{P}_1, \overline{P}_1)$ and $(\underline{P}_2, \overline{P}_2)$ the abstractions for $NN_1^{-1}(c)$ and for $NN_2^{-1}(c)$ respectively, for some class c. Then for all $Y_1 \in \gamma(\underline{P}_1, \overline{P}_1)$ and for all $Y_2 \in \gamma(\underline{P}_2, \overline{P}_2)$, we have:

$$2d_{GH}(Y_1, Y_2) \geq d_I(\underline{P}_1, \underline{P}_2) - 2\epsilon_1 - 2\epsilon_2$$

where $d_H(\underline{P}_1, \overline{P}_1) \leq \epsilon_1$ and $d_H(\underline{P}_2, \overline{P}_2) \leq \epsilon_2$.

Lower and upper bounds of the classification difference of two neural nets

Comparison theorem

Then for all $Y_1 \in \gamma(\underline{P}_1, \overline{P}_1)$ and for all $Y_2 \in \gamma(\underline{P}_2, \overline{P}_2)$, we have:
$2d_{GH}(Y_1, Y_2) \geq d_I(\underline{P}_1, \underline{P}_2) - 2\epsilon_1 - 2\epsilon_2$ where $d_H(\underline{P}_1, \overline{P}_1) \leq \epsilon_1$ and $d_H(\underline{P}_2, \overline{P}_2) \leq \epsilon_2$.

Proof

$d_{GH}(\underline{P}_1, \underline{P}_2) \geq d_{GH}(\underline{P}_1, Y_1) + d_{GH}(Y_1, Y_2) + d_{GH}(Y_2, \underline{P}_2)$ by the triangular inequality.
Now,

$$\begin{aligned} d_{GH}(\underline{P}_1, Y_1) &\leq d_{GH}(\underline{P}_1, \overline{P}_2) && (1)\\ &\leq d_H(\underline{P}_1, \overline{P}_2) && (2)\\ &\leq \epsilon_1 && (3) \end{aligned}$$

by partial monotony of d_{GH} and similarly, $d_{GH}(Y_2, \underline{P}_2) \leq \epsilon_2$. Thus,

$$d_{GH}(\underline{P}_1, \underline{P}_2) \geq d_{GH}(\underline{P}_1, \underline{P}_2) + \epsilon_1 + \epsilon_2 \quad (4)$$

Finally, by stability, we conclude:

$$\begin{aligned} 2d_{GH}(Y_1, Y_2) &\geq 2d_{GH}(\underline{P}_1, \underline{P}_2) - 2\epsilon_1 - 2\epsilon_2 && (5)\\ &\geq d_I(\underline{P}_1, \underline{P}_2) - 2\epsilon_1 - 2\epsilon_2 && (6) \end{aligned}$$

Remark: upper bound of the robustness

Is simpler to find, when needed...

- By choosing any matching from the set of points for class 1, with the inner approx.
- (Or any matching using the bottleneck distance, if we want to keep things at the level of persistence modules)

Thank you!

Numerous research paths from there...

{goubault,putot}@lix.polytechnique.fr