Adapting persistent homology to invariance groups

Patrizio Frosini[1,2]

[1] Department of Mathematics, University of Bologna, Italy
[2] ARCES - Vision Mathematics Group, University of Bologna, Italy
patrizio.frosini@unibo.it

Applied and Computational Algebraic Topology
Bremen, July 15-19, 2013

Outline

1. **The limitations of classical Persistent Homology**
2. ***G*-invariant persistent homology via quotient spaces**
3. ***G*-invariant persistent homology via *G*-operators**

The point of this talk

- It is well known that classical persistent homology is invariant under the action of the group $\mathrm{Homeo}(X)$ of all self-homeomorphisms of a topological space X. As a consequence, this theory is not able to distinguish two filtering functions $\varphi, \psi : X \to \mathbb{R}$ if a homeomorphism $h : X \to X$ exists, such that $\psi = \varphi \circ h$.
- However, in several applications the existence of a homeomorphism $h : X \to X$ such that $\psi = \varphi \circ h$ is not sufficient to consider φ and ψ equivalent to each other.
- How can we adapt the concept of persistence in order to get invariance just under the action of a proper subgroup of $\mathrm{Homeo}(X)$ rather than under the action of the whole group $\mathrm{Homeo}(X)$?

Example

These data are equivalent for classical Persistent Homology

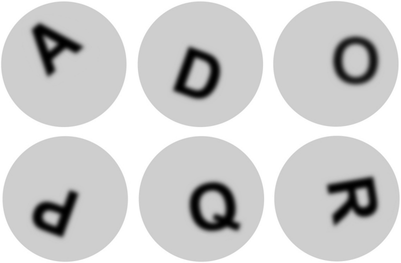

Figure: Examples of letters A, D, O, P, Q, R represented by functions $\varphi_A, \varphi_D, \varphi_O, \varphi_P, \varphi_Q, \varphi_R$ from the unit disk $D \subset \mathbb{R}^2$ to the real numbers. Each function $\varphi_Y : D \to \mathbb{R}$ describes the grey level at each point of the topological space D, with reference to the considered instance of the letter Y. Black and white correspond to the values 0 and 1, respectively (so that light grey corresponds to a value close to 1).

Example (continuation)

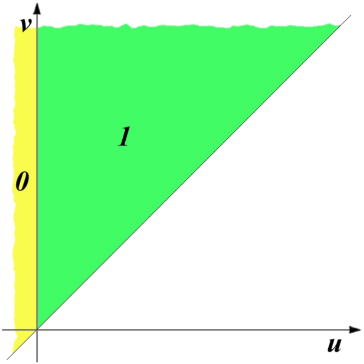

Figure: The persistent Betti number function (i.e. the rank invariant) in degree 0 for all images in the previous figure (“letters *A*, *D*, *O*, *P*, *Q*, *R*”).

Example (continuation)

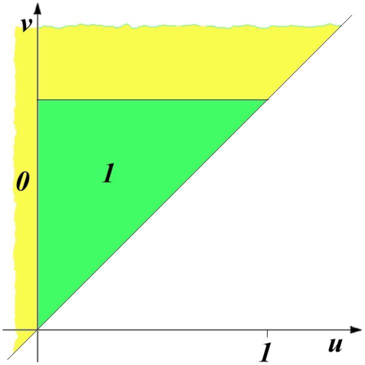

Figure: The persistent Betti number function (i.e. the rank invariant) in degree 1 for all images in the previous figure ("letters *A*, *D*, *O*, *P*, *Q*, *R*").

Example (continuation)

In our example classical persistent homology fails in distinguishing the letters because it is invariant under the action of homeomorphisms, and our six images are equivalent up to homeomorphisms.

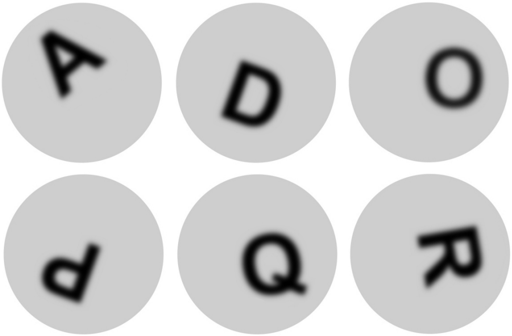

The main point

Classical persistent homology is not tailored to study invariance with respect to a group G different from the group of all self-homeomorphisms of a topological space.

In this talk we will show two ways to adapt classical persistent homology to the group G, in order to use it for shape comparison.

Observation

One could think of solving the problem described in the previous example by using other filtering functions, possibly defined on different topological spaces. For example, we could extract the boundaries of our letters and consider the distance from the center of mass of each boundary as a new filtering function. This approach presents some problems:

1. It usually requires an extra computational cost (e.g., to extract the boundaries of the letters in our previous example).
2. It can produce a different topological space for each new filtering function (e.g., the letters of the alphabet can have non-homeomorphic boundaries). Working with several topological spaces instead of just one can be a disadvantage.
3. It is not clear how we can translate the invariance that we need into the choice of new filtering functions defined on new topological spaces.

Before proceeding we need a *"ground truth"*.

In this talk, our ground truth will be the *natural pseudo-distance*.

Definition (Natural pseudo-distance)

Let X be a topological space. Let G be a subgroup of the group $\mathrm{Homeo}(X)$ of all self-homeomorphisms of X. Let S be a subset of the set $C^0(X, \mathbb{R})$ of all continuous functions from X to $\mathbb{R}$. The pseudo-distance $d_G : S \times S \to \mathbb{R}$ defined by setting

$$d_G(\varphi_1, \varphi_2) = \inf_{g \in G} \|\varphi_1 - \varphi_2 \circ g\|_\infty$$

is called the natural pseudo-distance associated with the group G.

$*$ P. Donatini and P. Frosini, Natural pseudodistances between closed surfaces, *Journal of the European Mathematical Society*, vol. 9 (2007), n. 2, 331-353

$*$ F. Cagliari, B. Di Fabio and C. Landi, The natural pseudo-distance as a quotient pseudo-metric, and applications, *Forum Mathematicum* (in press)

The rationale of using the natural pseudo-distance d_G

Important property

The natural pseudo-distance d_G is G-invariant. This means that $d_G(\varphi_1, \varphi_2 \circ g) = d_G(\varphi_1, \varphi_2)$ for every $g \in G$ and every $\varphi_1, \varphi_2 \in C^0(X, \mathbb{R})$.

The rationale of using the natural pseudo-distance d_G consists in considering two shapes σ_1 and σ_2 equivalent to each other if a transformation exists in the group G, taking the measurements on σ_1 to the measurements on σ_2.

BASIC ASSUMPTION

The observer has the right to change the invariance group G according to his/her judgement. Therefore we look at G as a variable in our problem.

The rationale of using the natural pseudo-distance d_G

Example: Two gray-level pictures can be considered equivalent if a gray-level-preserving rigid motion exists, transforming one picture into the other.

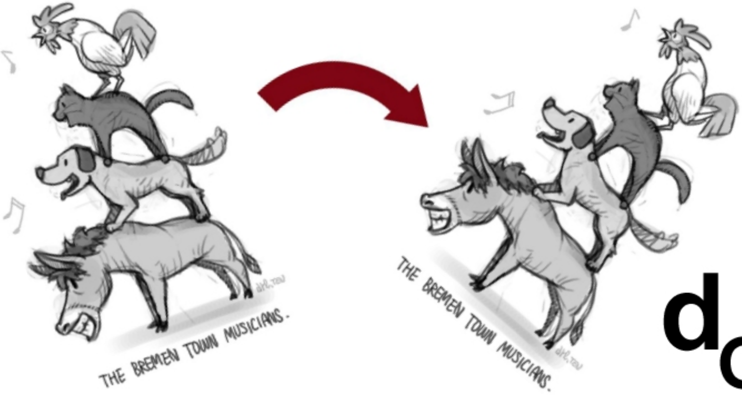

$$d_G = 0$$

Remark: the case G equal to the trivial group

Assume that $\mathbf{I} = \{id\}$ is the trivial group, containing only the identical homeomorphism. We observe that

$$d_G(\varphi_1, \varphi_2) \leq d_{\mathbf{I}}(\varphi_1, \varphi_2) = \|\varphi_1 - \varphi_2\|_\infty$$

for every continuous function $\varphi_1, \varphi_2 : X \to \mathbb{R}$.

Another reason to use the natural pseudo-distance d_G

The natural pseudo-distance d_G allows to obtain a stability result for persistent diagrams that is better than the classical one, involving d_∞:

$$d_{match}(\rho_{\varphi_1}, \rho_{\varphi_2}) \leq d_G(\varphi_1, \varphi_2) \leq \|\varphi_1 - \varphi_2\|_\infty$$

for every continuous function $\varphi_1, \varphi_2 : X \to \mathbb{R}$.

EXAMPLE: here $d_{match}(\rho_{\varphi_1}, \rho_{\varphi_2}) = 0 = d_G(\varphi_1, \varphi_2) < \|\varphi_1 - \varphi_2\|_\infty = 1$

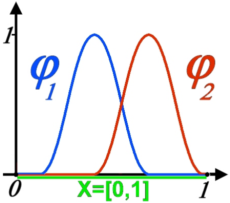

Figure: These two functions have the same persistent homology ($d_{match}(\rho_{\varphi_1}, \rho_{\varphi_2}) = 0$, but $\|\varphi_1 - \varphi_2\|_\infty = 1$).They are equivalent w.r.t. $G = \text{Homeo}([0, 1])$.

Remark

We need to apply persistent homology in a way that is invariant under the action of a given subgroup G of the group Homeo(X).

We could think of using the well known concept of Equivariant Homology. In other words, in the case that G acts freely on X, one could think of considering the topological quotient space X/G, endowed with the filtering functions $\hat{\varphi}, \hat{\psi}$ that take each orbit ω of the group G to the maximum of φ and ψ on ω, respectively.

We observe that this approach would not be of help in the case that the action of the group G is transitive (such as in the "letters example"), since the quotient of X/G is just a singleton. As a consequence, if we considered two filtering functions $\varphi, \psi : X \to \mathbb{R}$ with $\max \varphi = \max \psi$, the persistent homology of the induced functions $\hat{\varphi}, \hat{\psi} : X/G \to \mathbb{R}$ would be the same.

Therefore, we need to use a different procedure.

Our approach

Our first approach is based on the choice of a suitable group H associated with the group G.

- We choose a subgroup H of Homeo(X) such that
  1. H is finite (i.e. $H = \{h_1, \ldots, h_r\}$);
  2. $g \circ h \circ g^{-1} \in H$ for every $g \in G$ and every $h \in H$. (This implies that the restriction to H of the conjugacy action of each $g \in G$ is a permutation of H.)
- We compute the persistent homology group of the quotient space $\frac{X}{H}$, with respect to the filtering function $\hat{\varphi}$ that takes each orbit ω of the group H to the maximum of φ on ω.

We shall use the symbol $PH_n^{\hat{\varphi}}(u, v)$ to denote the persistent homology group in degree n of $\frac{X}{H}$ with respect to the filtering function $\hat{\varphi} : \frac{X}{H} \to \mathbb{R}$, computed at the point (u, v).

Remark

If G is Abelian, a simple way of getting a subgroup H of Homeo(X) verifying properties 1 and 2 consists in setting H equal to a finite subgroup of G. However, we have to observe that in most of the applications, the group G is not Abelian.

If G is finite, a trivial way of getting a subgroup H of Homeo(X) verifying properties 1 and 2 consists in setting $H = G$. This choice leads to consider the quotient space X/G. However, we have to observe that in most of the applications, the group G is not finite.

The trivial choice $H = \mathbf{I} = \{id\}$ can be always made, but it leads to compute the classical persistent homology of the topological space X.

Two key properties of $PH_n^{\hat{\varphi}}$ are expressed by the following results:

Theorem (Invariance with respect to the group G)

If $g \in G$ and $u, v \in \mathbb{R}$ with $u < v$, the groups $PH_n^{\hat{\varphi}\circ g}(u, v)$ and $PH_n^{\hat{\varphi}}(u, v)$ are isomorphic.

Under suitable assumptions about the topological space X the next statement holds:

Theorem (Stability)

For every $n \in \mathbb{Z}$, let us set $\rho_n^{\hat{\varphi}}(u, v) := \mathrm{rank}\left(PH_n^{\hat{\varphi}}(u, v)\right)$ and $\rho_n^{\hat{\psi}}(u, v) := \mathrm{rank}\left(PH_n^{\hat{\psi}}(u, v)\right)$. Then

$$d_{match}(\rho_n^{\hat{\varphi}}, \rho_n^{\hat{\psi}}) \leq d_G(\varphi, \psi) \leq d_{id}(\varphi, \psi) = \|\varphi - \psi\|_\infty.$$

We recall that d_{match} is the classical matching distance between the persistent diagrams associated with $PH_n^{\hat{\varphi}}$ and $PH_n^{\hat{\psi}}$.

Some problems with our method:

The previous method presents some drawbacks:

- It requires to find a suitable nontrivial group H, associated with G. This group could be difficult to find or not exist at all.
- The computation of the persistent homology group in degree n with respect to the filtering function $\hat{\varphi} : \frac{X}{H} \to \mathbb{R}$ requires a fine enough triangulation of X that is invariant under the action of H. This triangulation could be difficult to find or not exist at all.

An alternative approach based on G-operators

Fortunately, an alternative approach is available.

We will describe it in the second part of this talk.

This part is based on an ongoing joint research project with

Grzegorz Jabłoński
Jagiellonian University - Kraków

An alternative approach based on *G*-operators

Informal description of our idea

Instead of changing the topological space X, we can get invariance with respect to the group G by changing the "glasses" that we use "to observe" the filtering functions. In our approach, these "glasses" are G-operators F_i, which act on the filtering functions.

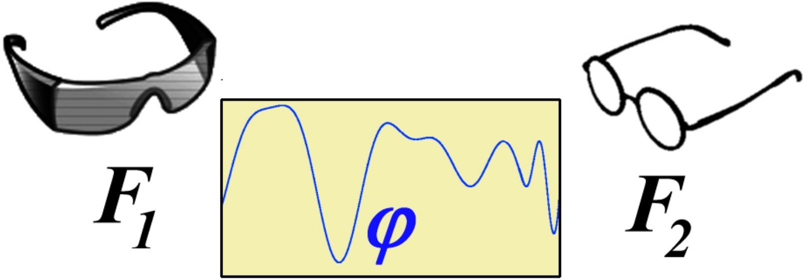

Let us consider the following objects:

- A triangulable space X with nontrivial homology in degree k.
- A set $\mathcal{C}$ of continuous functions from X to $\mathbb{R}$, that contains the set of all constant functions.
- A topological subgroup G of Homeo(X) that acts on $\mathcal{C}$ by composition on the right.
- The natural pseudo-distance d_G on $\mathcal{C}$ with respect to G, defined by setting $d_G(\varphi_1, \varphi_2) := \inf_{g \in G} \|\varphi_1 - \varphi_2 \circ g\|_\infty$ for every $\varphi_1, \varphi_2 \in \mathcal{C}$.
- The distance d_∞ on $\mathcal{C}$, defined by setting $d_\infty(\varphi_1, \varphi_2) := \|\varphi_1 - \varphi_2\|_\infty$. This is just the natural pseudo-distance d_G in the case that G is the trivial group $\mathbf{I} = \{id\}$, containing only the identical homeomorphism.
- A subset $\mathcal{F}$ of the set $\mathcal{F}^{\text{all}}(\mathcal{C}, G)$ of all non-expansive G-operators from $\mathcal{C}$ to $\mathcal{C}$.

The operator space $F \in \mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$

In plain words, $F \in \mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$ means that

1. $F : \mathcal{C} \to \mathcal{C}$
2. $F(\varphi \circ g) = F(\varphi) \circ g$. *(F is a G-operator)*
3. $\|F(\varphi_1) - F(\varphi_2)\|_\infty \leq \|\varphi_1 - \varphi_2\|_\infty$. *(F is non-expansive)*

The operator F is not required to be linear.

Some simple examples of F, taking $\mathcal{C}$ equal to the set of all continuous functions $\varphi : S^1 \to \mathbb{R}$ and G equal to the group of all rotations of S^1:

- $F(\varphi) :=$ the constant function $\psi : S^1 \to \mathbb{R}$ taking everywhere the value $\max \varphi$;
- $F(\varphi) :=$ the function $\psi : S^1 \to \mathbb{R}$ defined by setting $\psi(x) = \max\left\{\varphi\left(x - \frac{\pi}{8}\right), \varphi\left(x + \frac{\pi}{8}\right)\right\}$;
- $F(\varphi) :=$ the function $\psi : S^1 \to \mathbb{R}$ defined by setting $\psi(x) = \frac{1}{2}\left(\varphi\left(x - \frac{\pi}{8}\right) + \varphi\left(x + \frac{\pi}{8}\right)\right)$.

The pseudo-metric $D^{\mathcal{F}}_{\text{match}}$

For every $\varphi_1, \varphi_2 \in \mathcal{C}$ we set

$$D^{\mathcal{F}}_{\text{match}}(\varphi_1, \varphi_2) := \sup_{F \in \mathcal{F}} d_{match}(\rho_k(F(\varphi_1)), \rho_k(F(\varphi_2)))$$

where $\rho_k(\psi)$ denotes the persistent Betti number function (i.e. the rank invariant) of ψ in degree k.

Proposition

$D^{\mathcal{F}}_{match}$ is a G-invariant and stable pseudo-metric on $\mathcal{C}$.

The G-invariance of $D^{\mathcal{F}}_{match}$ means that
$D^{\mathcal{F}}_{\text{match}}(\varphi_1, \varphi_2 \circ g) = D^{\mathcal{F}}_{\text{match}}(\varphi_1, \varphi_2)$ for every $\varphi_1, \varphi_2 \in \mathcal{C}$ and every $g \in G$.

Theoretical results

THEORETICAL RESULTS

OUR MAIN RESULT:

We observe that the pseudo-distance $D^{\mathcal{F}}_{\mathrm{match}}$ and the natural pseudo-distance d_G are defined in quite different ways.

In particular, the definition of $D^{\mathcal{F}}_{\mathrm{match}}$ is based on persistent homology, while the natural pseudo-distance d_G is based on the group of homeomorphisms G.

In spite of this, the following statement holds:

Theorem

If $\mathcal{F} = \mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$, then the pseudo-distance $D^{\mathcal{F}}_{\mathrm{match}}$ coincides with the natural pseudo-distance d_G on $\mathcal{C}$.

Our idea

The previous theorem suggests the following approach.

Let us choose a finite subset $\mathcal{F}^*$ of $\mathcal{F}$, and consider the pseudo-metric

$$D_{\mathrm{match}}^{\mathcal{F}^*}(\varphi_1, \varphi_2) := \max_{F \in \mathcal{F}^*} d_{match}(\rho_k(F(\varphi_1)), \rho_k(F(\varphi_2)))$$

for every $\varphi_1, \varphi_2 \in \mathcal{C}$.

Obviously, $D_{\mathrm{match}}^{\mathcal{F}^*} \leq D_{\mathrm{match}}^{\mathcal{F}}$.

Furthermore, if $\mathcal{F}^*$ is dense enough in $\mathcal{F}$, then the new pseudo-distance $D_{\mathrm{match}}^{\mathcal{F}^*}$ is close to $D_{\mathrm{match}}^{\mathcal{F}}$.

In order to make this point clear, we need the next theoretical result.

Compactness of $\mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$

The following result holds:

Theorem

If $(\mathcal{C}, d_\infty)$ is a compact metric space and G is a compact topological group, then $\mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$ is a compact metric space with respect to the distance d defined by setting

$$d(F_1, F_2) := \max_{\varphi \in \mathcal{C}} \|F_1(\varphi) - F_2(\varphi)\|_\infty$$

for every $F_1, F_2 \in \mathcal{F}$.

Approximation of $\mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$

This statement follows:

Corollary

Assume that both the metric space $(\mathcal{C}, d_\infty)$ and the topological group G are compact. Let $\mathcal{F}$ be a subset of $\mathcal{F}^{\mathrm{all}}(\mathcal{C}, G)$. For every $\epsilon > 0$, a finite subset $\mathcal{F}^*$ of $\mathcal{F}$ exists, such that

$$\left| D_{match}^{\mathcal{F}^*}(\varphi_1, \varphi_2) - D_{match}^{\mathcal{F}}(\varphi_1, \varphi_2) \right| \leq \epsilon$$

for every $\varphi_1, \varphi_2 \in \mathcal{C}$.

This corollary implies that the pseudo-distance $D_{match}^{\mathcal{F}}$ can be approximated computationally, at least in the compact case.

Let us check what happens in practice

A RETRIEVAL EXPERIMENT ON A DATASET OF CURVES

Let us check what happens in practice

We have considered

1. a dataset of 10000 functions from S^1 to $\mathbb{R}$, depending on five random parameters $(*)$;
2. these three invariance groups:
   - the group Homeo(S^1) of all self-homeomorphisms of S^1
   - the group $R(S^1)$ of all rotations of S^1
   - the trivial group $\mathbf{I}(S^1) = \{id\}$, containing just the identity of S^1.

Obviously,

$$\text{Homeo}(S^1) \supset R(S^1) \supset \mathbf{I}(S^1).$$

$(*)$ For $1 \leq i \leq 10000$ we have set $\bar{\varphi}_i(x) = r_1 \sin(3x) + r_2 \cos(3x) + r_3 \sin(4x) + r_4 \cos(4x)$, with $r_1, .., r_4$ randomly chosen in the interval $[-2, 2]$; the i-th function in our dataset is the function $\varphi_i := \bar{\varphi}_i \circ \gamma_i$, where $\gamma_i(x) := 2\pi(\frac{x}{2\pi})^{r_5}$ and r_5 is randomly chosen in the interval $[\frac{1}{2}, 2]$.

Let us check what happens in practice

The choice of Homeo(S^1) as an invariance group implies that the following two functions are considered equivalent. Their graphs are obtained from each other by applying a horizontal stretching. Also shifts are accepted as legitimate transformations.

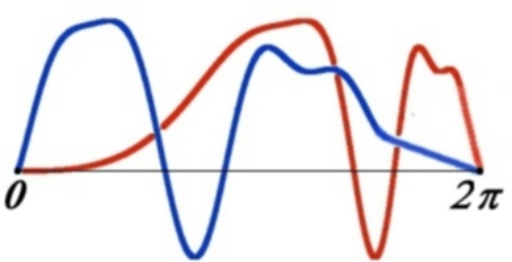

Let us check what happens in practice

The choice of $R(S^1)$ as an invariance group implies that the following two functions are considered equivalent. Their graphs are obtained from each other by applying a rotation of S^1. Stretching is not accepted as a legitimate transformation.

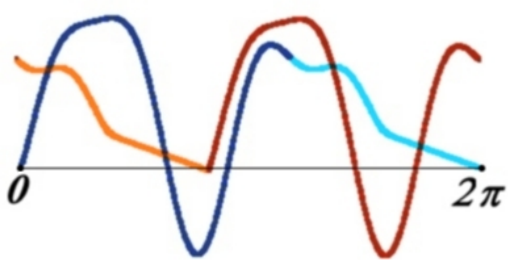

Finally, the choice of $\mathbf{I}(S^1) = \{id\}$ as an invariance group means that two functions are considered equivalent if and only if they coincide everywhere.

Theoretical results

What happens if we decide to assume

that the invariance group is the group Homeo(S^1)

of all self-homeomorphisms of S^1?

The results of an experiment: the group Homeo(S^1)

If we choose $G = \mathrm{Homeo}(S^1)$, to proceed we need to choose a finite set of non-expansive Homeo(S^1)-operators. In our experiment we have considered these three non-expansive Homeo(S^1)-operators:

- $F_0 = id$ (i.e., $F_0(\varphi) = \varphi$);
- $F_1 = -id$ (i.e., $F_0(\varphi) = -\varphi$);
- $F_2 = \frac{1}{5} \cdot \sup\{-\varphi(x_1) + \varphi(x_2) - \frac{1}{2}\varphi(x_3) + \frac{1}{2}\varphi(x_4) - \varphi(x_5) + \varphi(x_6)\}$, $(x_1, \ldots, x_6)$ varying among all the counterclockwise 6-tuples on S^1. $(*)$

This choice produces the Homeo(S^1)-invariant pseudo-distance

$$D_{match}^{\mathcal{F}^*}(\varphi_1, \varphi_2) := \max_{0 \le i \le 2} d_{match}(\rho_k(F_i(\varphi_1)), \rho_k(F_i(\varphi_2))).$$

Here $\mathcal{F}^* = \{F_0 := id, F_1, F_2, F_3\}$.

$(*)$ P. Frosini and C. Landi, Reparametrization invariant norms, *Transactions of the American Mathematical Society*, vol. 361 (2009), 407-452.

An important remark

It is important to use several operators. The use of just one operator still produces a pseudo-distance $D_{match}^{\mathcal{F}^*}$ that is invariant under the action of the group G, but this choice is far from guaranteeing a good approximation of the natural pseudo-distance d_G.

As an example in the case $G = \mathrm{Homeo}(S^1)$, if we use just the identity operator (i.e., we just apply classical persistent homology), we cannot distinguish these two functions $\varphi_1, \varphi_2 : S^1 \to \mathbb{R}$, despite the fact that they are different for d_G:

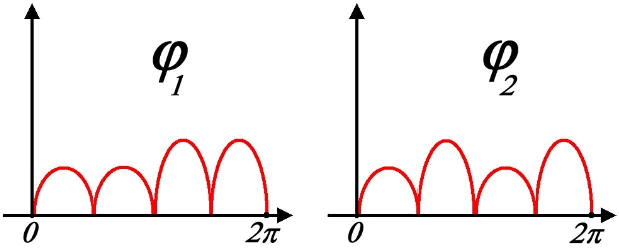

The results of an experiment: the group $\mathrm{Homeo}(S^1)$

Here is a query (in **blue**), and the first four retrieved functions (in **black**):

$$D_{\text{match}}^{\mathcal{F}^*}(\varphi_1, \varphi_2)$$

Mean	Max
1.648	4.831

Standard deviation
0.934

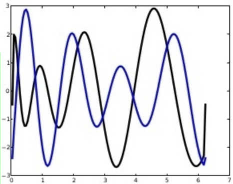

(a) φ_{516}, dist: 0.0465393

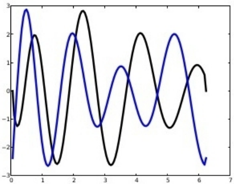

(b) φ_{381}, dist: 0.0541687

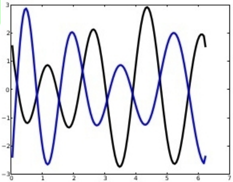

(c) φ_{7776}, dist: 0.0984192

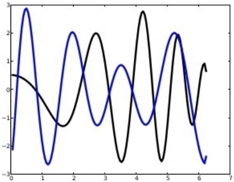

(d) φ_{6214}, dist: 0.10376

The results of an experiment: the group $\mathrm{Homeo}(S^1)$

Let's have a closer look at the query and at the first retrieved function:
Here is the query:

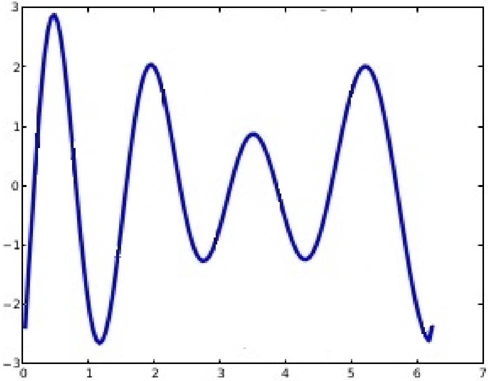

The results of an experiment: the group $\mathrm{Homeo}(S^1)$

Here is the first retrieved function with respect to $D_{match}^{\mathcal{F}^*}$:

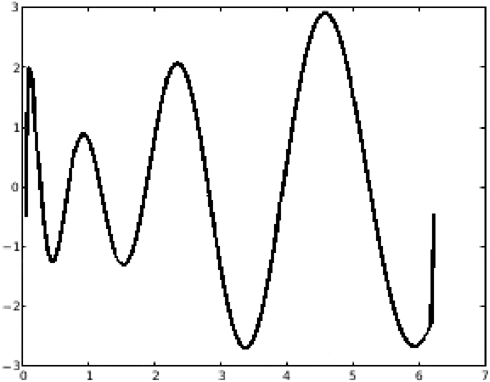

The results of an experiment: the group $\mathrm{Homeo}(S^1)$

Here is the query function after aligning it to the first retrieved function by means of a shift (in **red**).
The first retrieved function is represented in **black**.
The figure shows that the retrieved function is approximately equivalent to the query function, by applying a shift and a stretching.

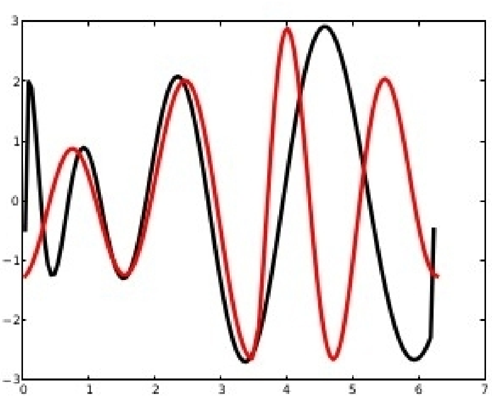

The results of an experiment: the group $\mathrm{Homeo}(S^1)$

Here is the query function after aligning it to the first four retrieved functions by means of a shift (in **red**).
The first four retrieved functions are represented in **black**.

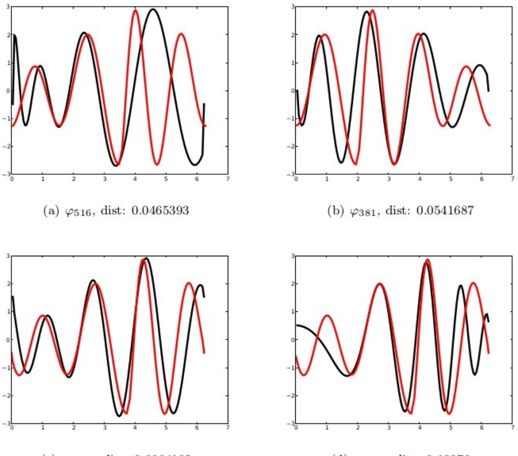

(a) φ_{516}, dist: 0.0465393

(b) φ_{381}, dist: 0.0541687

(c) φ_{7776}, dist: 0.0984192

(d) φ_{6214}, dist: 0.10376

Theoretical results

What happens if we decide to assume

that the invariance group is the group $R(S^1)$

of all rotations of S^1?

The results of an experiment: the group $R(S^1)$

If we choose $G = R(S^1)$, in order to proceed we need to choose a finite set of non-expansive $R(S^1)$-operators. Obviously, since F_0, F_1 and F_2 are Homeo(S^1)-invariant, they are also $R(S^1)$-invariant. In our experiment we have added these five non-expansive $R(S^1)$-operators (which are not Homeo(S^1)-invariant) to F_0, F_1 and F_2:

- $F_3(\varphi) := \max\{\varphi(x), \varphi(x+\pi)\}$
- $F_4(\varphi) := \frac{1}{2} \cdot (\varphi(x) + \varphi(x + \frac{\pi}{4}))$
- $F_5(\varphi) := \max\{\varphi(x), \varphi(x+\pi/10), \varphi(x + \frac{2\pi}{10}), \varphi(x + \frac{3\pi}{10})\}$
- $F_6(\varphi) := \frac{1}{3} \cdot (\varphi(x) + \varphi(x + \frac{\pi}{3}) + \varphi(x + \frac{\pi}{4}))$
- $F_7(\varphi) := \frac{1}{3} \cdot (\varphi(x) + \varphi(x + \frac{\pi}{3}) + \varphi(x + \frac{2\pi}{3}))$

This choice produces the $R(S^1)$-invariant pseudo-distance

$$D_{match}^{\mathcal{F}^*}(\varphi_1, \varphi_2) := \max_{0 \le i \le 7} d_{match}(\rho_k(F_i(\varphi_1)), \rho_k(F_i(\varphi_2))).$$

Here $\mathcal{F}^* = \{F_0 := id, F_1, F_2, F_3, F_4, F_5, F_6, F_7\}$.

The results of an experiment: the group $R(S^1)$

Here is a query (in **blue**), and the first four retrieved functions (in **black**):

$D^{\mathcal{F}^*}_{\mathrm{match}}(\varphi_1, \varphi_2)$

Mean	Max
1.938	4.831

Standard deviation
0.874

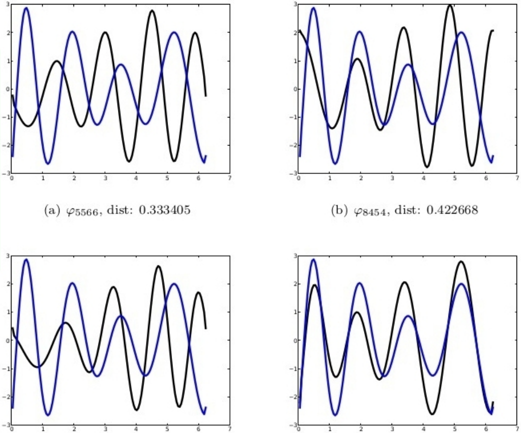

(a) φ_{5566}, dist: 0.333405

(b) φ_{8454}, dist: 0.422668

(c) φ_{8909}, dist: 0.453949

(d) φ_{4426}, dist: 0.46463

The results of an experiment: the group $R(S^1)$

Let's have a closer look at the query and at the first retrieved function:
Here is the query:

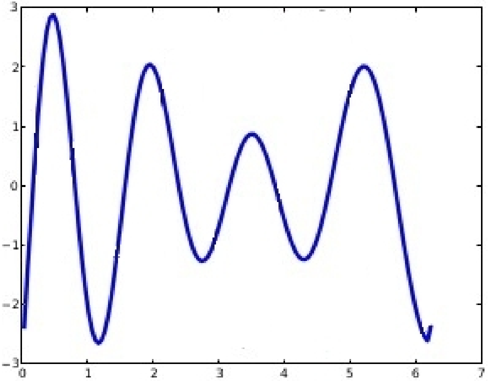

The results of an experiment: the group $R(S^1)$

Here is the first retrieved function with respect to $D_{match}^{\mathcal{F}^*}$:

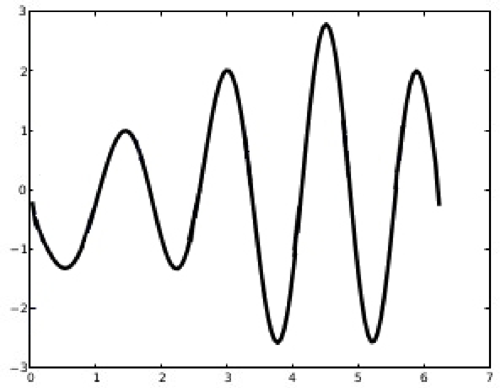

The results of an experiment: the group $R(S^1)$

Here is the query function after aligning it to the first retrieved function by means of a shift (in **red**).
The first retrieved function is represented in **black**.
The figure shows that the retrieved function is approximately equivalent to the query function, via a shift.

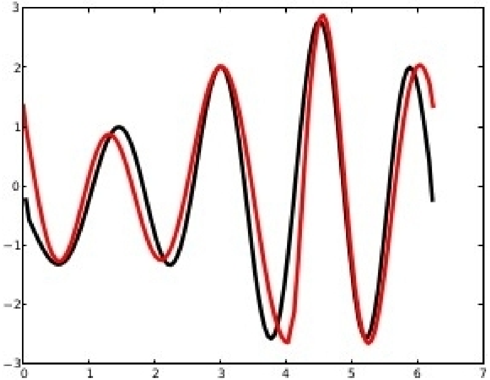

The results of an experiment: the group $R(S^1)$

Here is the query function after aligning it to the first four retrieved functions by means of a shift (in **red**).
The first four retrieved functions are represented in **black**.

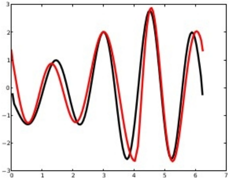

(a) φ_{5566}, dist: 0.333405

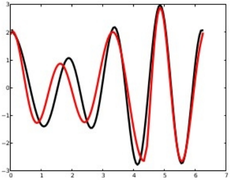

(b) φ_{8454}, dist: 0.422668

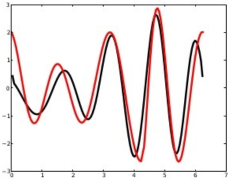

(c) φ_{8909}, dist: 0.453949

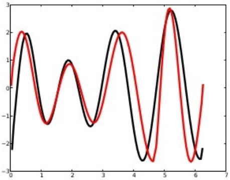

(d) φ_{4426}, dist: 0.46463

Theoretical results

Finally, what happens if we decide to assume that the invariance group is the group $\mathbf{I}(S^1) = \{id\}$ containing only the identity of S^1?

This means that the "perfect" retrieved function should coincide with our query.

The results of an experiment: the group $\mathbf{I}(S^1) = \{id\}$

If we choose $G = \mathbf{I}(S^1)$, in order to proceed we need to choose a finite set of non-expansive operators (obviously, every operator is an $\mathbf{I}(S^1)$-operator).
In our experiment we have considered these three non-expansive operators (which are not $R(S^1)$-operator):

- $F_8(\varphi) := \sin(x)\varphi(x)$
- $F_9(\varphi) := \frac{\sqrt{2}}{2}\sin(x)\varphi(x) + \frac{\sqrt{2}}{2}\cos(x)\varphi(x + \frac{\pi}{2})$
- $F_{10}(\varphi) := \sin(2x)\varphi(x)$

We have added F_8, F_9, F_{10} to $F_1, \ldots, F_7$.

This choice produces the pseudo-distance

$$D^{\mathcal{F}^*}_{match}(\varphi_1, \varphi_2) := \max_{0 \le i \le 10} d_{match}(\rho_k(F_i(\varphi_1)), \rho_k(F_i(\varphi_2))).$$

Here $\mathcal{F}^* = \{F_0 := id, F_1, F_2, F_3, F_4, F_5, F_6, F_7, F_8, F_9, F_{10}\}$.

The results of an experiment: the group $\mathbf{I}(S^1) = \{id\}$

Here is a query (in **blue**), and the first four retrieved functions (in **black**):

$$D_{\text{match}}^{\mathcal{F}^*}(\varphi_1, \varphi_2)$$

Mean	Max
2.022	4.831

Standard deviation
0.828

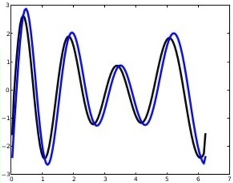

(a) φ_{7133}, dist: 0.415802

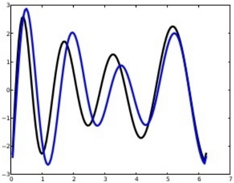

(b) φ_{7001}, dist: 0.598145

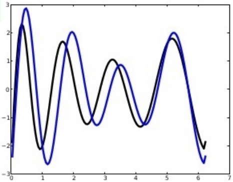

(c) φ_{389}, dist: 0.617218

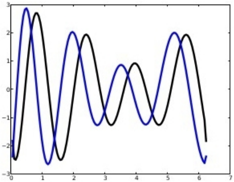

(d) φ_{5723}, dist: 0.617981

The results of an experiment: the group $\mathbf{I}(S^1) = \{Id\}$

Let's have a closer look at the query and at the first retrieved function:
Here is the query:

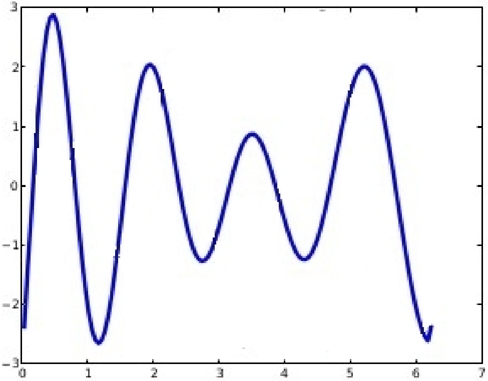

The results of an experiment: the group $\mathbf{I}(S^1) = \{Id\}$

Here is the first retrieved function with respect to $D^{\mathcal{F}}_{match}$:

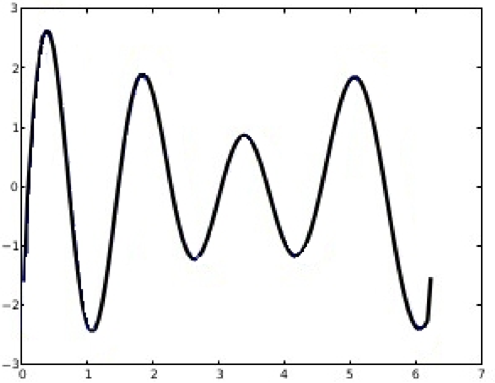

The results of an experiment: the group $\mathbf{I}(S^1) = \{Id\}$

The first retrieved function is represented in **black.**
As expected, no aligning shift is necessary here.
The figure shows that the retrieved function is approximately equal to the query function.

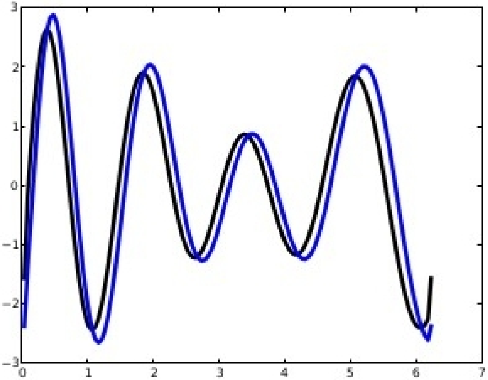

The results of an experiment: the group $\mathbf{I}(S^1) = \{Id\}$

Here we show again the query function and the first four retrieved functions (in **black**).
The figure shows that the retrieved functions are approximately coinciding with the query function.

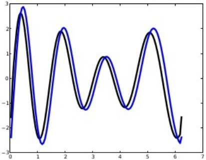

(a) φ_{7133}, dist: 0.415802

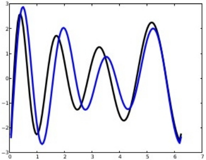

(b) φ_{7001}, dist: 0.598145

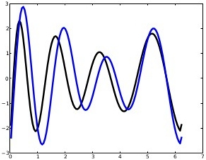

(c) φ_{389}, dist: 0.617218

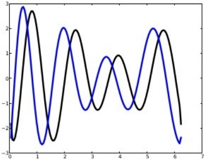

(d) φ_{5723}, dist: 0.617981

An open problem

We have proven that if $\mathcal{C}$ and G are compact, then $D^{\mathcal{F}}_{match}$ can be approximated computationally.

However, this result does not say which set of operators allows for both a good approximation of $D^{\mathcal{F}}_{match}$ and a fast computation.

Further research is needed in this direction.

Conclusions

In this talk we have shown that

- Persistent homology can be adapted to proper subgroups of the group of all self-homeomorphisms of a triangulable space, in two different ways. Both of these methods are stable with respect to noise.
- In particular, the approach based on non-expansive G-operators can be used for any subgroup G of $\mathrm{Homeo}(S^1)$. An experiment concerning this method has been illustrated, showing the possible use of this approach for data retrieval.

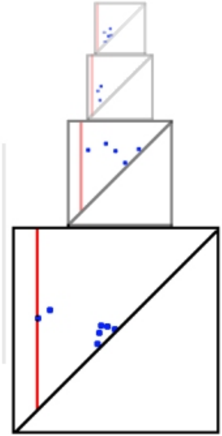

THANKS FOR YOUR ATTENTION