An approach to topological data analysis via persistent topology and invariant operators

Patrizio Frosini

Department of Mathematics and ARCES, University of Bologna
patrizio.frosini@unibo.it

Homotopy Probability Theory - Saarbrücken, 4-8 September 2017

Outline

Our basic questions

Assumptions in our model

Mathematical setting and theoretical results

An experiment concerning functions from $\mathbb{R}$ to $\mathbb{R}$

A first step towards the application of our model: GIPHOD

Some work in progress

Our basic questions

We are interested in these questions:

- Is there a general metric model to compare data in TDA?
- What should be the role of the observer in such a model?
- How could we approximate the metric used in that model?

Our talk will be devoted to illustrate these questions and to propose some answers by means of a mathematical approach based on persistent homology and group invariant non-expansive operators.

Assumptions in our model

Truth often depends on the observer's perspective:

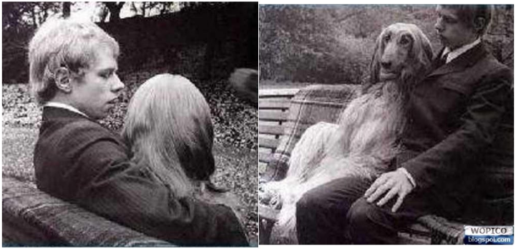

Multiple perspectives are usually unavoidable! In the past this observation was mostly confined to the philosophical debate, but nowadays it starts to be quite relevant also in several scientific applications involving Information Technology.

Assumptions in our model

We will make these assumptions:

1. No object can be studied in a direct and absolute way. Any object is only knowable through acts of measurement made by an observer.
2. Any act of measurement can be represented as a function defined on a topological space.
3. The observer usually acquires measurement data by applying operators to the functions describing these data. These operators are frequently endowed with some invariances that are relevant for the observer.
4. Only the observer is entitled to decide about data similarity.

Assumptions in our model

In some sense, we could summarize our assumptions by saying that

DATA ANALYSIS IS ALWAYS THE ANALYSIS OF AN OBSERVER

Our goal is not to approximate objects but to approximate observers.

Assumptions in our model

The observer usually acquires measurement data by applying operators to the functions describing the data.

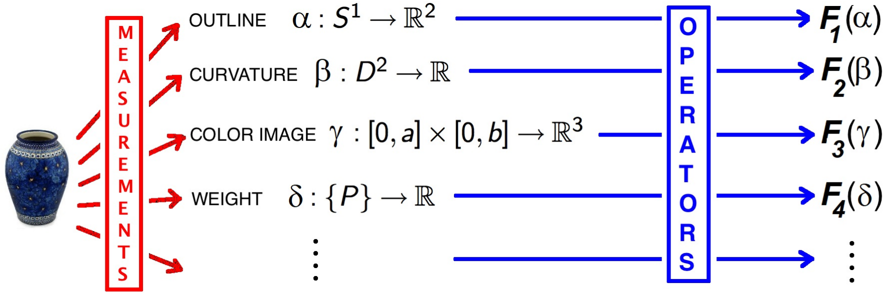

Let us give some examples of measurements and operators.

An example of measurement

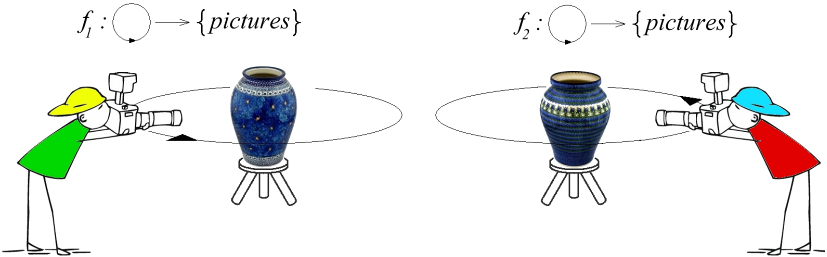

$\Phi =$ set of continuous functions from S^1 to $C^0(\mathbb{R}^2, \mathbb{R})$
(in the case of grey level images).

Another example of measurement

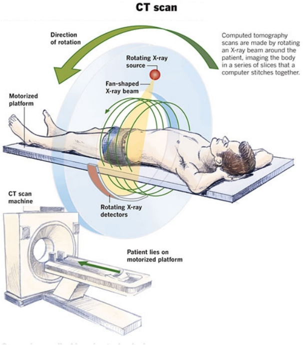

$\Phi =$ set of continuous functions from $\mathbb{R}$ to $\mathbb{R}$.

Another example of measurement

Φ = set of functions from a singleton to $\mathbb{R}$.

An example of operator

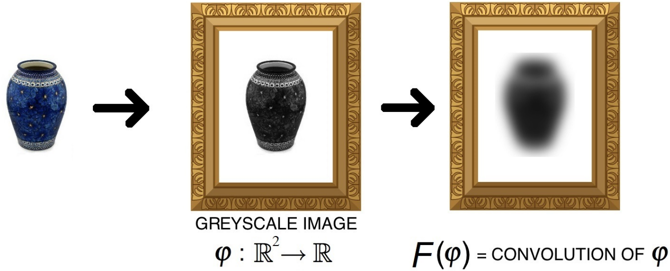

$$F(\varphi)(x) := \frac{1}{2\pi\sigma^2} \int_{\mathbb{R}^2} \varphi(y) e^{-\frac{\|x-y\|^2}{2\sigma^2}} \, dy.$$

F is G-invariant for G equal to the group of isometries of $\mathbb{R}^2$.

F is also non-expansive with respect to the sup-norm.

Another example of operator

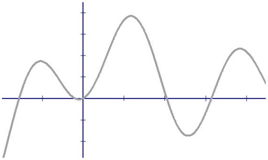

$\varphi : \mathbb{R} \to [0,1]$

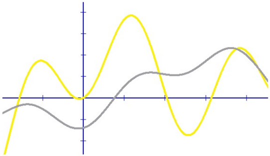

$F(\varphi)(x) := \frac{1}{2}(\varphi(x-1) + \varphi(x+1))$

F is G-invariant for G equal to the group of translations of $\mathbb{R}$.
F is also non-expansive with respect to the sup-norm.

Another example of operator

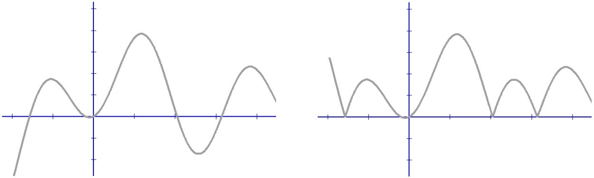

$\varphi : \mathbb{R} \to [0,1]$ $\qquad$ $F(\varphi) = |\varphi|$

F is G-invariant for G equal to the group of homeomorphisms of $\mathbb{R}$.
F is also non-expansive with respect to the sup-norm.

Choice of the operators

- The observer cannot usually choose the functions representing the measurement data, but he/she can often choose the operators that will be applied to those functions.
- The choice of the operators reflects the invariances that are relevant for the observer.
- In some sense we could state that the observer can be represented as a collection of (suitable) operators, endowed with the invariance he/she has chosen.

In the first part of this talk we will mainly examine the case of operators that act on a space Φ of continuous functions and take Φ to itself. We will also assume that these operators preserve the self-homeomorphisms of X.

From comparing sets in $\mathbb{R}^n$ to comparing functions

Instead of directly focusing on the objects we are interested in, we focus on the filtering functions describing the measurements we make on them, and on the "glasses" that we use "to observe" the functions. In our approach, these "glasses" are G-operators which act on the filtering functions.
These operators represent the observer's perspective.

In some sense, the family of operators defines the observer.

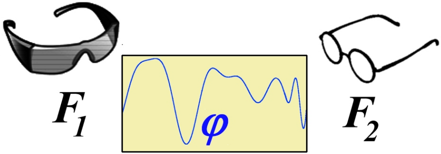

Natural pseudo-distance associated with a group G

First of all we need a definition allowing us to formalize the comparison of data in our model.

Definition

Let X be a compact space. Let G be a subgroup of the group $\mathrm{Homeo}(X)$ of all homeomorphisms $f : X \to X$. The pseudo-distance $d_G : C^0(X,\mathbb{R}) \times C^0(X,\mathbb{R}) \to \mathbb{R}$ defined by setting

$$d_G(\varphi,\psi) = \inf_{g\in G} \max_{x\in X} |\varphi(x) - \psi(g(x))|$$

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

In plain words, the definition of d_G is based on the attempt of finding the best correspondence between the functions φ, ψ by means of homeomorphisms in G.

A possible objection

A possible objection: "The use of the group of homeomorphisms makes the natural pseudo-distance d_G difficult to apply. For example, in shape comparison two objects are usually not homeomorphic, hence this pseudo-metric cannot be applied to real problems."

This objection can be faced by recalling that **the homeomorphisms do not concern the "objects" but the space where the measurements are made**. For example, if we take a grey level image, our measurement space can be modelled as the real plane and each image can be represented as a function from $\mathbb{R}^2$ to $\mathbb{R}$. Therefore, the space X is not given by the (possibly non-homeomorphic) objects displayed in the picture, but by the topological space $\mathbb{R}^2$.
Analogously, each subset of the 3D space can be associated with a probability density describing the probability that each point $p \in \mathbb{R}^3$ belongs to the considered object. In this case the space X is $\mathbb{R}^3$.

G-invariant non-expansive operators

The natural pseudo-distance d_G represents our ground truth.

Unfortunately, d_G is difficult to compute. This is also a consequence of the fact that we can easily find subgroups G of $\mathrm{Homeo}(X)$ that cannot be approximated with arbitrary precision by smaller finite subgroups of G (i.e. $G =$ group of rigid motions of $X = \mathbb{R}^3$).

Nevertheless, in this talk we will show that d_G can be approximated with arbitrary precision by means of a **DUAL** approach based on persistent homology and G-invariant non-expansive operators.

This research is based on an ongoing joint research project with Grzegorz Jabłoński (IST - Austria)

G-invariant non-expansive operators

Let us consider the following objects:

- A triangulable space X with nontrivial homology in degree k (e.g., for $k = 0$, $X =$ the real plane).
- A set Φ of continuous functions from X to $\mathbb{R}$, containing at least the set of the constant functions taking every finite value c with $|c| \leq \sup_{\varphi \in \Phi} \|\varphi\|_\infty$ (e.g., $\Phi =$ set of grey level pictures, i.e. functions from the real plane to $[0,1]$).
- A topological subgroup G of $\mathrm{Homeo}(X)$ that acts on Φ by composition on the right (e.g., $G =$ group of rigid motions of $\mathbb{R}^2$).
- A subset $\mathscr{F}$ of the set $\mathscr{F}^{\mathrm{all}}(\Phi, G)$ of all non-expansive G-operators from Φ to Φ.

The operator space $\mathscr{F}^{\mathrm{all}}(\Phi, G)$

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

1. $F : \Phi \to \Phi$
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.

The operators verifying properties 1, 2, 3 are called G-invariant non-expansive operators (**GINO**s) for (Φ, G).

The operator space $\mathscr{F}^{\mathrm{all}}(\Phi, G)$

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

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

What is persistent homology?

If $\varphi : X \to \mathbb{R}$ is a continuous functions, we can consider the sublevel sets $X_t := \{x \in X : \varphi(x) \le t\}$. When t varies we see the birth and death of k-dimensional holes.

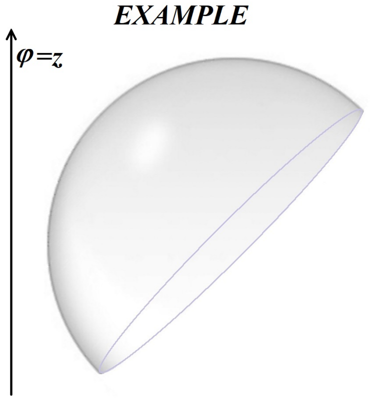

What is persistent homology?

If $\varphi : X \to \mathbb{R}$ is a continuous functions, we can consider the sublevel sets $X_t := \{x \in X : \varphi(x) \leq t\}$. When t varies we see the birth and death of k-dimensional holes.

EXAMPLE

No 1-dimensional hole — *Birth of a 1-dimensional hole* — *Death of the 1-dimensional hole*

What is persistent homology?

In plain words, the persistence diagram in degree k of φ is the collection of the pairs (b_i, d_i) where b_i and d_i are the times of birth and death of the i-th hole of dimension k.

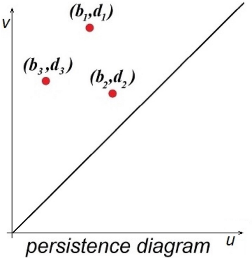

The points of the persistence diagram are endowed with multiplicity. Each point of the diagonal $u = v$ is assumed to be a point of the persistence diagram, endowed with infinite multiplicity.

What are persistent Betti number functions?

Persistence diagrams are not quite suitable for statistical purposes, because no good definition of average of persistence diagrams exists.

Persistent Betti number functions are more suitable for statistics.

Definition

The k-th persistent Betti number $\beta_k(u, v)$ is the number of holes of dimension k whose time of birth is greater than u and whose time of death is smaller than v.

The average of persistent Betti number functions can be trivially defined as the usual average of real-valued functions.

What are persistent Betti number functions?

The use of averages of persistent Betti number functions in degree 0 firstly appeared in the papers

- Pietro Donatini, Patrizio Frosini, Alberto Lovato, *Size functions for signature recognition*, Proceedings of SPIE, Vision Geometry VII, vol. 3454 (1998), 178183.
- Massimo Ferri, Patrizio Frosini, Alberto Lovato, Chiara Zambelli, *Point selection: A new comparison scheme for size functions (With an application to monogram recognition)*, Proceedings Third Asian Conference on Computer Vision, Lecture Notes in Computer Science 1351, vol. I, R. Chin, T. Pong (editors) Springer-Verlag, Berlin Heidelberg (1998), 329337.

What are persistent Betti number functions?

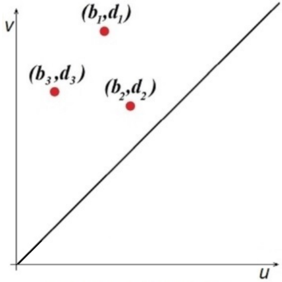

persistence diagram

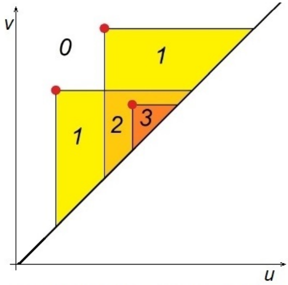

persistent Betti number function

If we use Čech homology, persistence diagrams are equivalent to persistent Betti number functions.

Comparison of persistent Betti number functions

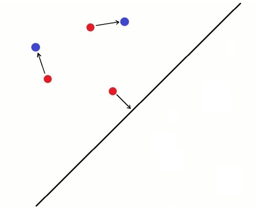

A way to change D_1 into D_2.

Persistence diagrams (and hence persistent Betti number functions) can be compared by means of the bottleneck distance. The bottleneck distance between two persistence diagrams D_1, D_2 is the minimum cost of changing the points of D_1 into the points of D_2, where the cost of moving each point is given by the max-norm distance in $\mathbb{R}^2$. Moving a point to the diagonal is equivalent to delete it.

The pseudo-metric $D^{\mathscr{F}}_{\mathrm{match}}$

Choose a set $\mathscr{F} \subseteq \mathscr{F}^{\mathrm{all}}(\Phi, G)$. For every $\varphi_1, \varphi_2 \in \Phi$ we set

$$D^{\mathscr{F}}_{\mathrm{match}}(\varphi_1, \varphi_2) := \sup_{F \in \mathscr{F}} d_{match}(\beta_k(F(\varphi_1)), \beta_k(F(\varphi_2)))$$

where $\beta_k(\psi)$ denotes the persistent Betti number function (i.e. the rank invariant) of ψ in degree k, while d_{match} denotes the usual bottleneck distance that is used to compare the persistence diagrams associated with $\beta_k(F(\varphi_1))$ and $\beta_k(F(\varphi_2))$.

Proposition

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

The G-invariance of $D^{\mathscr{F}}_{match}$ means that for every $\varphi_1, \varphi_2 \in \Phi$ and every $g \in G$ the equality $D^{\mathscr{F}}_{\mathrm{match}}(\varphi_1, \varphi_2 \circ g) = D^{\mathscr{F}}_{\mathrm{match}}(\varphi_1, \varphi_2)$ holds.

An equivalence result

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

In particular, the definition of $D^{\mathscr{F}}_{\mathrm{match}}$ 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 $\mathscr{F} = \mathscr{F}^{\mathrm{all}}(\Phi, G)$, then the pseudo-distance $D^{\mathscr{F}}_{\mathrm{match}}$ coincides with the natural pseudo-distance d_G on Φ.

Our main idea

The previous theorem suggests to study $D_{\mathrm{match}}^{\mathscr{F}}$ instead of d_G.

To this end, let us choose a finite subset $\mathscr{F}^*$ of $\mathscr{F}$, and consider the pseudo-metric

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

for every $\varphi_1, \varphi_2 \in \Phi$.

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

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

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

Compactness of $\mathscr{F}^{\mathrm{all}}(\Phi, G)$

The following result holds:

Theorem

If Φ is a compact metric space with respect to the sup-norm, then $\mathscr{F}^{\mathrm{all}}(\Phi, G)$ is a compact metric space with respect to the distance d defined by setting

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

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

Approximation of $\mathscr{F}^{\text{all}}(\Phi, G)$

This statement follows:

Corollary

Assume that the metric space Φ is compact with respect to the sup-norm. Let $\mathscr{F}$ be a subset of $\mathscr{F}^{\text{all}}(\Phi, G)$. For every $\varepsilon > 0$, a finite subset $\mathscr{F}^*$ of $\mathscr{F}$ exists, such that

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

for every $\varphi_1, \varphi_2 \in \Phi$.

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

Two references

- Patrizio Frosini, Grzegorz Jabłoński, *Combining persistent homology and invariance groups for shape comparison*, Discrete & Computational Geometry, vol. 55 (2016), n. 2, pages 373-409.
- Patrizio Frosini, *Towards an observer-oriented theory of shape comparison*, Proceedings of the 8th Eurographics Workshop on 3D Object Retrieval, Lisbon, Portugal, May 7-8, 2016, A. Ferreira, A. Giachetti, and D. Giorgi (Editors), 5-8.

Recap

We have seen that d_G (our ground truth about measurements) can be approximated by means of the pseudo-distance $D_{match}^{\mathscr{F}^*}$, which is based on persistent homology and the concept of group invariant non-expansive operator. This result is important for applications in shape comparison and topological data analysis.

This fact naturally leads us to the need of studying the topological space of group invariant non-expansive operators.

An experiment concerning functions from $\mathbb{R}$ to $\mathbb{R}$

We conclude this talk by illustrating another experiment, concerning functions from $\mathbb{R}$ to $\mathbb{R}$.
We have considered

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

Obviously,

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

(#) 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 $\mathrm{Homeo}(\mathbf{S}^1)$ 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.

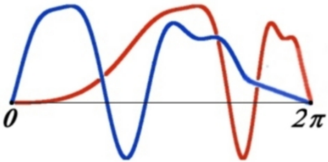

Let us check what happens in practice

The choice of $R(\mathbf{S}^1)$ 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 $\mathbf{S}^1$. Stretching is not accepted as a legitimate transformation.

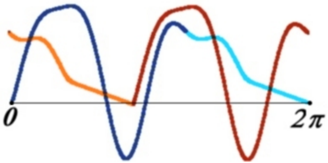

Finally, the choice of $\mathbf{I}(\mathbf{S}^1) = \{id\}$ as an invariance group means that two functions are considered equivalent if and only if they coincide everywhere.

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

What happens if we decide to assume

that the invariance group is the group $\mathrm{Homeo}(\mathbf{S}^1)$

of all self-homeomorphisms of $\mathbf{S}^1$?

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

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

- $F_0 := id$ (i.e., $F_0(\varphi) := \varphi$);
- $F_1 := -id$ (i.e., $F_0(\varphi) := -\varphi$);
- $F_2(\varphi) :=$ the constant function $\psi : \mathbf{S}^1 \to \mathbb{R}$ taking the value $\frac{1}{5} \cdot \sup\{-\varphi(x_1) + \varphi(x_2) - \frac{1}{2}\varphi(x_3) + \frac{1}{2}\varphi(x_4) - \varphi(x_5) + \varphi(x_6)\}$, $(x_1, \ldots, x_6)$ varying among all the counterclockwise 6-tuples on $\mathbf{S}^1$.

This choice produces the $\mathrm{Homeo}(\mathbf{S}^1)$-invariant pseudo-distance

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

An important remark

It is important to use several operators. The use of just one operator still produces a pseudo-distance $D_{match}^{\mathscr{F}^*}$ 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 $G = \mathrm{Homeo}(\mathbf{S}^1)$, if we use just the identity operator (i.e., we just apply classical persistent homology), we cannot distinguish these two functions $\varphi_1, \varphi_2 : \mathbf{S}^1 \to \mathbb{R}$, despite the fact that they are different for d_G:

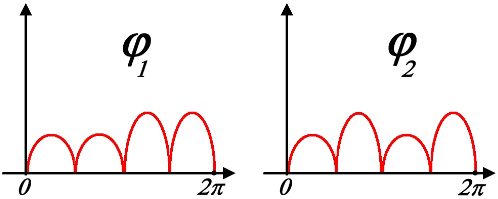

The results of an experiment: the group $\mathrm{Homeo}(\mathbf{S}^1)$

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
1.648	4.831

Standard deviation
0.934

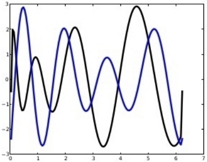

(a) φ_{516}, dist: 0.0465393

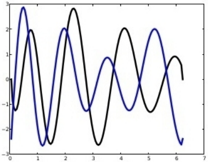

(b) φ_{381}, dist: 0.0541687

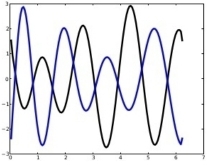

(c) φ_{7776}, dist: 0.0984192

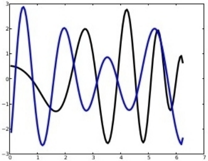

(d) φ_{6214}, dist: 0.10376

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

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

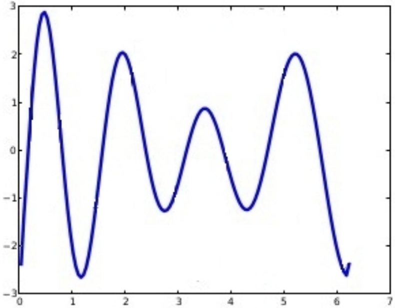

The results of an experiment: the group $\mathrm{Homeo}(\mathbf{S}^1)$

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

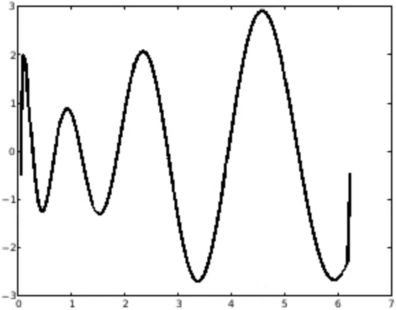

The results of an experiment: the group $\mathrm{Homeo}(\mathbf{S}^1)$

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.

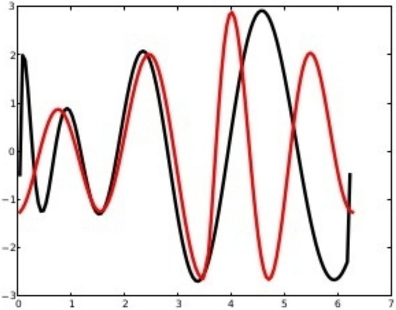

The results of an experiment: the group Homeo($\mathbf{S}^1$)

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**.

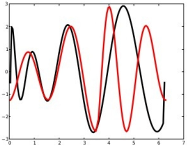

(a) φ_{516}, dist: 0.0465393

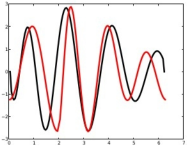

(b) φ_{381}, dist: 0.0541687

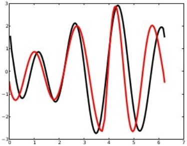

(c) φ_{7776}, dist: 0.0984192

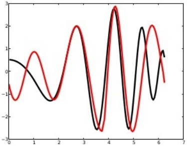

(d) φ_{6214}, dist: 0.10376

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

What happens if we decide to assume

that the invariance group is the group $R(\mathbf{S}^1)$

of all rotations of $\mathbf{S}^1$?

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

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

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

This choice produces the $R(\mathbf{S}^1)$-invariant pseudo-distance

$$D^{\mathscr{F}^*}_{match}(\varphi_1, \varphi_2) := \max_{0 \le i \le 7} d_{match}(\beta_k(F_i(\varphi_1)), \beta_k(F_i(\varphi_2))).$$

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

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

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

Mean	Max
1.938	4.831

Standard deviation
0.874

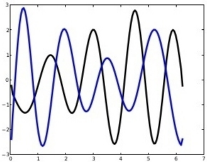

(a) φ_{5566}, dist: 0.333405

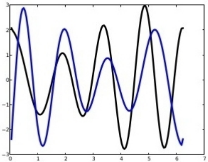

(b) φ_{8454}, dist: 0.422668

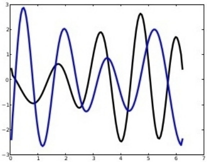

(c) φ_{8909}, dist: 0.453949

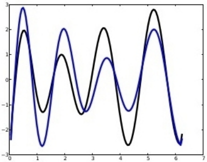

(d) φ_{4426}, dist: 0.46463

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

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

Here is the query:

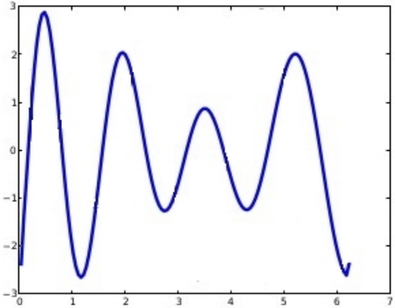

The results of an experiment: the group $R(\mathbf{S}^1)$

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

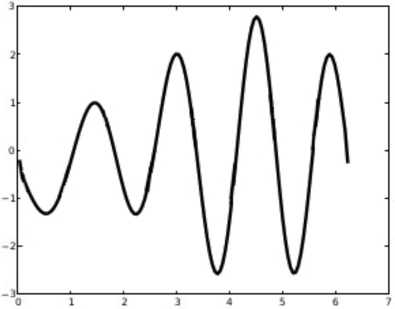

The results of an experiment: the group $R(\mathbf{S}^1)$

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.

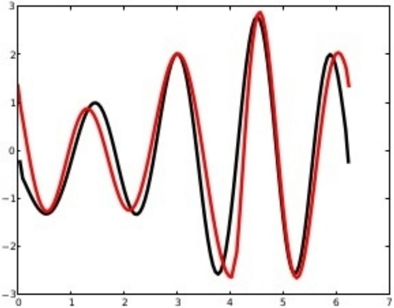

The results of an experiment: the group $R(\mathbf{S}^1)$

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**.

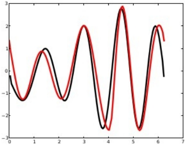

(a) φ_{5566}, dist: 0.333405

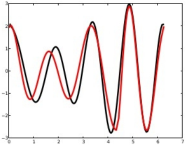

(b) φ_{8454}, dist: 0.422668

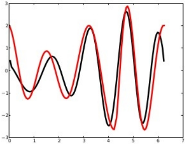

(c) φ_{8909}, dist: 0.453949

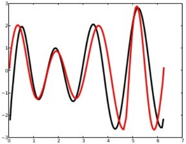

(d) φ_{4426}, dist: 0.46463

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

Finally, what happens if we decide to assume

that the invariance group is the group $\mathbf{I}(\mathbf{S}^1) = \{id\}$

containing only the identity of $\mathbf{S}^1$?

This means that the "perfect" retrieved function

should coincide with our query.

Remark: This is exactly the case where **we should not** use our dual approach! (Just compute $d_{\mathbf{I}(\mathbf{S}^1)}(\varphi_1, \varphi_2) = \|\varphi_1 - \varphi_2\|_\infty$ directly!)

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

If we choose $G = \mathbf{I}(\mathbf{S}^1) = \{id\}$, in order to proceed we need to choose a finite set of non-expansive operators (obviously, every operator is an $\mathbf{I}(\mathbf{S}^1)$-operator).

In our experiment we have considered these three non-expansive operators (which are <u>not</u> $R(\mathbf{S}^1)$-operators):

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

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

This choice produces the pseudo-distance

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

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

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

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

Mean	Max
2.022	4.831

Standard deviation
0.828

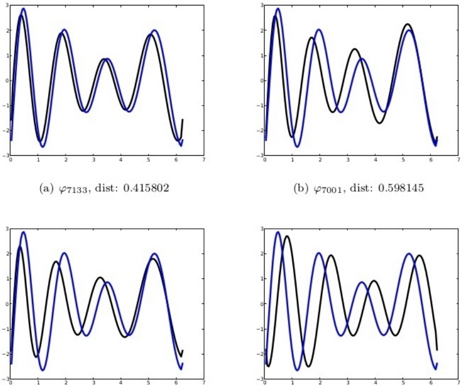

(a) φ_{7133}, dist: 0.415802

(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 $\mathbf{I}(\mathbf{S}^1)$

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

Here is the query:

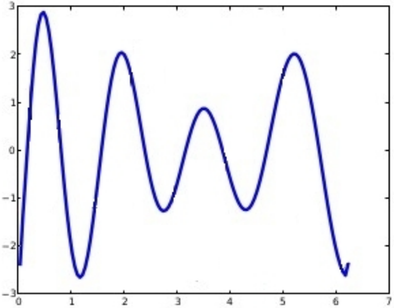

The results of an experiment: the group $\mathbf{I}(\mathbf{S}^1)$

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

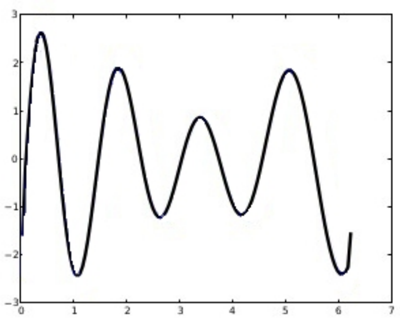

The results of an experiment: the group $\mathbf{I}(\mathbf{S}^1)$

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.

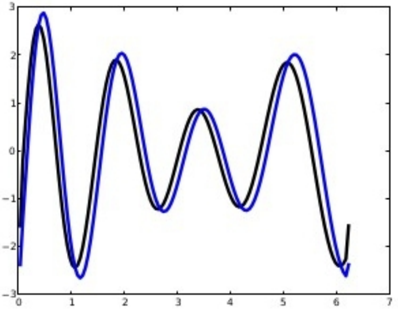

The results of an experiment: the group $\mathbf{I}(\mathbf{S}^1)$

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.

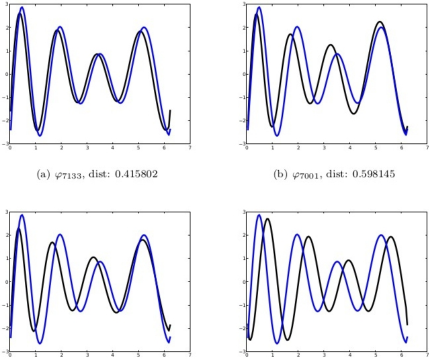

(a) φ_{7133}, dist: 0.415802

(b) φ_{7001}, dist: 0.598145

GIPHOD

Joint project with Grzegorz Jabłoński (IST - Austria) and Marc Ethier (Université de Saint-Boniface - Canada)

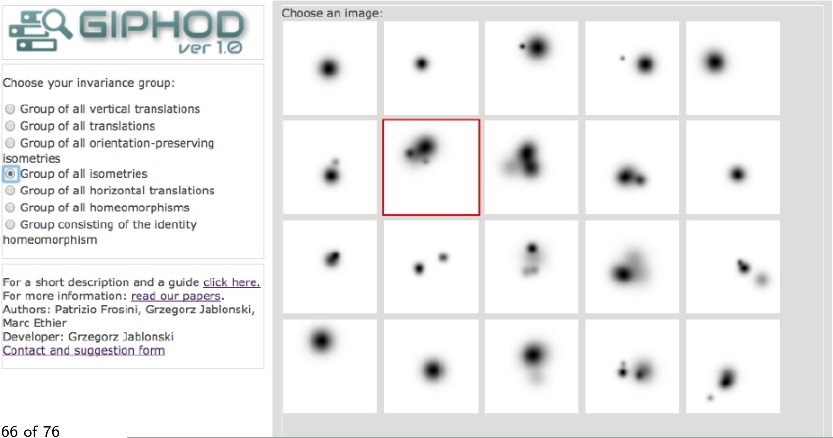

GIPHOD (Group Invariant Persistent Homology On-line Demonstrator)

GIPHOD is an on-line demonstrator, allowing the user to choose an image and an invariance group. GIPHOD searches for the most similar images in the dataset, with respect to the chosen invariance group.
Purpose: to show the use of our theoretical approach for image comparison.
Dataset: 10.000 quite simple grey level synthetic images obtained by adding randomly chosen bell-shaped functions. The images are coded as functions from $\mathbb{R}^2$ to $[0,1]$.

GIPHOD can be tested at `http://giphod.ii.uj.edu.pl/index2`.
All suggestions are greatly appreciated and welcomed (please send them to grzegorz.jablonski@uj.edu.pl)

GIPHOD (Group Invariant Persistent Homology On-line Demonstrator)

We will now show some results obtained by GIPHOD when the invariance group G is the group of isometries:
Some data about the pseudo-metric $D^{\mathscr{F}}_{\mathrm{match}}$ in this case:

- The images are coded as functions from $\mathbb{R}^2 \to [0, 1]$;
- Mean distance between images: 0.2984;
- Standard deviation of distance between images: 0.1377;
- Number of GINOs that have been used: 5.

GIPHOD (Group Invariant Persistent Homology On-line Demonstrator)

List of GINOs that have been used in the following image retrievals, where the invariance group G is the group of isometries:

- $F(\varphi) = \varphi$.
- $F(\varphi) :=$ constant function taking each point to the value $\int_{\mathbb{R}^2} \varphi(\boldsymbol{x})\ d\boldsymbol{x}$.
- $F(\varphi)$ defined by setting

$$F(\varphi)(\boldsymbol{x}) := \int_{\mathbb{R}^2} \varphi(\boldsymbol{x} - \boldsymbol{y}) \cdot \beta\left(\|\boldsymbol{y}\|_2\right)\ d\boldsymbol{y}$$

where $\beta : \mathbb{R} \to \mathbb{R}$ is an integrable function with $\int_{\mathbb{R}^2} \left|\beta\left(\|\boldsymbol{y}\|_2\right)\right|\ d\boldsymbol{y} \leq 1$. Three GINOs of this kind have been used.

GIPHOD: Examples for the group of isometries

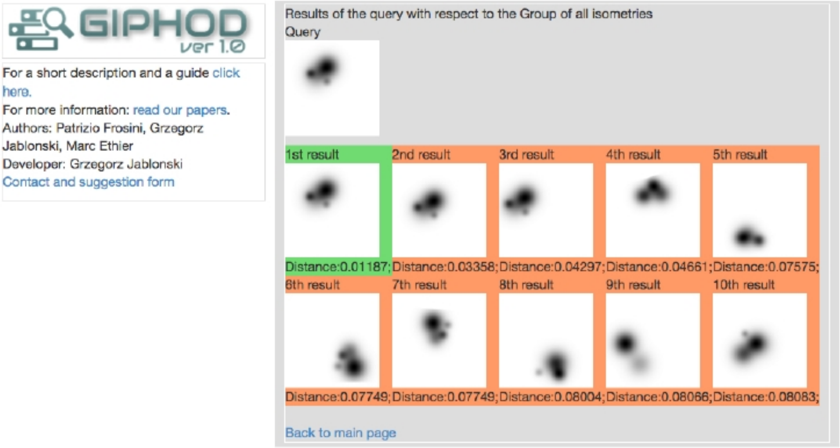

Mean distance: 0.2984; Standard deviation of distance: 0.1377; Number of GINOs: 5.

GIPHOD: Examples for the group of isometries

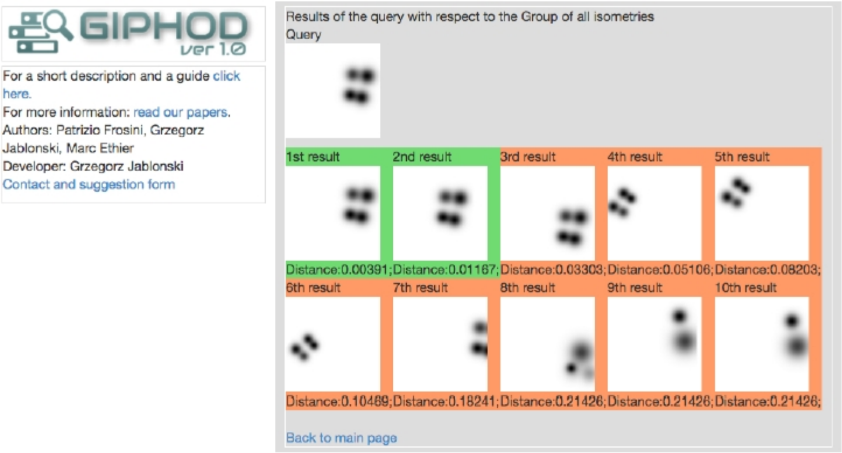

Mean distance: 0.2984; Standard deviation of distance: 0.1377; Number of GINOs: 5.

GIPHOD: Examples for the group of isometries

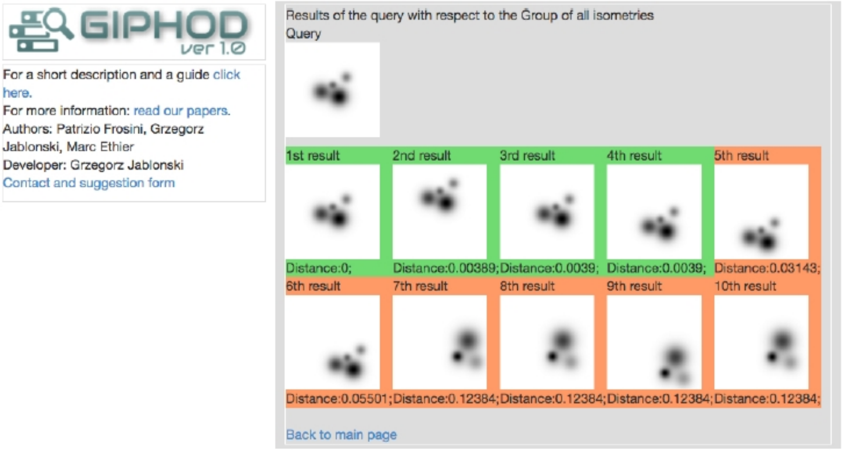

Mean distance: 0.2984; Standard deviation of distance: 0.1377; Number of GINOs: 5.

Our basic questions

Assumptions in our model

Mathematical setting and theoretical results

An experiment concerning functions from $\mathbb{R}$ to $\mathbb{R}$

A first step towards the application of our model: GIPHOD

Some work in progress

Some work in progress

Some work is in progress, concerning these three lines of research:

- Change of the topologies used on X and G.
- Extension of our approach to operators taking the space Φ (where a group G acts) to a different space of functions Ψ (where another group H acts). These operators also act on the group G, changing each $g \in G$ into an $h \in H$.
- Study of the metric space of GINOs both in the case $(\Phi, G) = (\Psi, H)$ and in the case $(\Phi, G) \neq (\Psi, H)$.

(Joint work with Nicola Quercioli - University of Bologna)

I will speak about this subject in my next talk.

Conclusions

In this talk we have supported these statements:

- Data comparison is based on acts of measurement made by an observer. The acts of measurement can be represented as a function defined on a topological space X. The observer can be seen as a collection of G-invariant non-expansive operators, applied to the functions describing the data.
- These functions can be compared by means of the natural pseudo-distance associated with any subgroup G of $\mathrm{Homeo}(X)$.
- Persistent homology can be used to approximate the natural pseudo-metric d_G. This can be done by means of a method that is based on G-invariant non-expansive operators. This method is stable with respect to noise.


Thanks for your attention!