Generalized persistent homologies: G-invariant and multi-dimensional persistence

2 – G-invariant persistent homology

Massimo Ferri[1,3], Patrizio Frosini[1,3], Claudia Landi[2,3]

[1]Dipartimento di Matematica, Università di Bologna

[2]DISMI, Università di Modena e Reggio Emilia

[3]ARCES, Università di Bologna

massimo.ferri@unibo.it, patrizio.frosini@unibo.it, claudia.landi@unimore.it

Homology: Theoretical and Computational Aspects
Genoa, 9-13 February 2015

Outline

Natural pseudo-distance associated with a group G

G-invariant persistent homology

G-invariant non-expansive operators

GIPHOD

Natural pseudo-distance associated with a group G

G-invariant persistent homology

G-invariant non-expansive operators

GIPHOD

An example in shape comparison

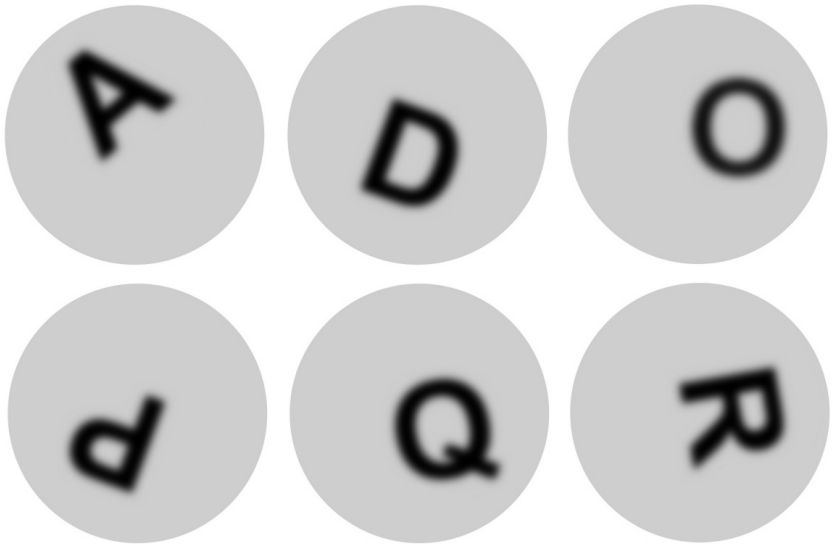

Figure: Examples of letters $\mathtt{A}, \mathtt{D}, \mathtt{O}, \mathtt{P}, \mathtt{Q}, \mathtt{R}$ represented by functions $\varphi_{\mathtt{A}}, \varphi_{\mathtt{D}}, \varphi_{\mathtt{O}}, \varphi_{\mathtt{P}}, \varphi_{\mathtt{Q}}, \varphi_{\mathtt{R}}$ from $\mathbb{R}^2$ to the real numbers. Each function $\varphi_Y : \mathbb{R}^2 \to \mathbb{R}$ describes the grey level at each point of the topological space $\mathbb{R}^2$, 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).

A letter O

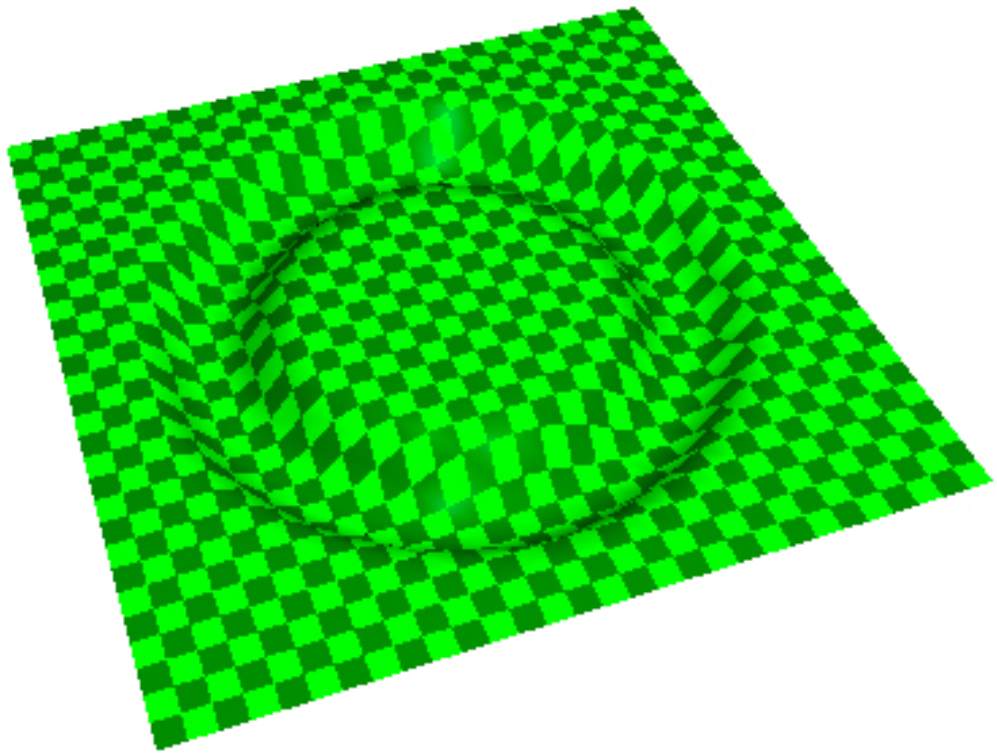

Figure: Part of the graph of a function representing a letter O.

Key observation

Persistent homology is invariant with respect to ANY homeomorphism!

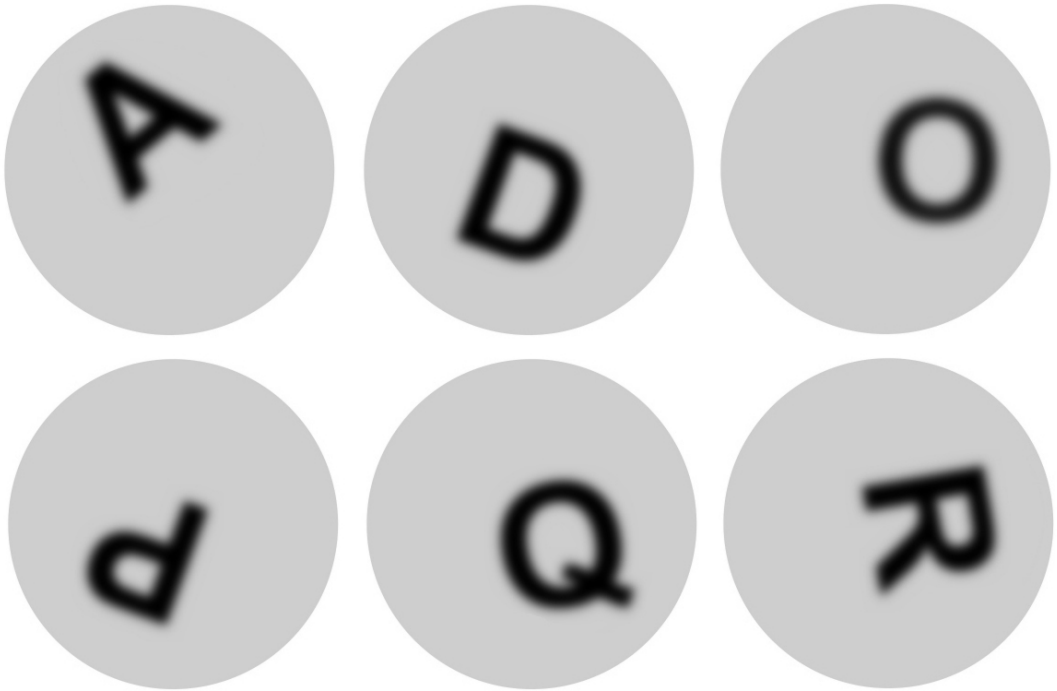

Figure: These functions share the same persistent homology.

Main question

How can we use persistent homology to distinguish these letters?

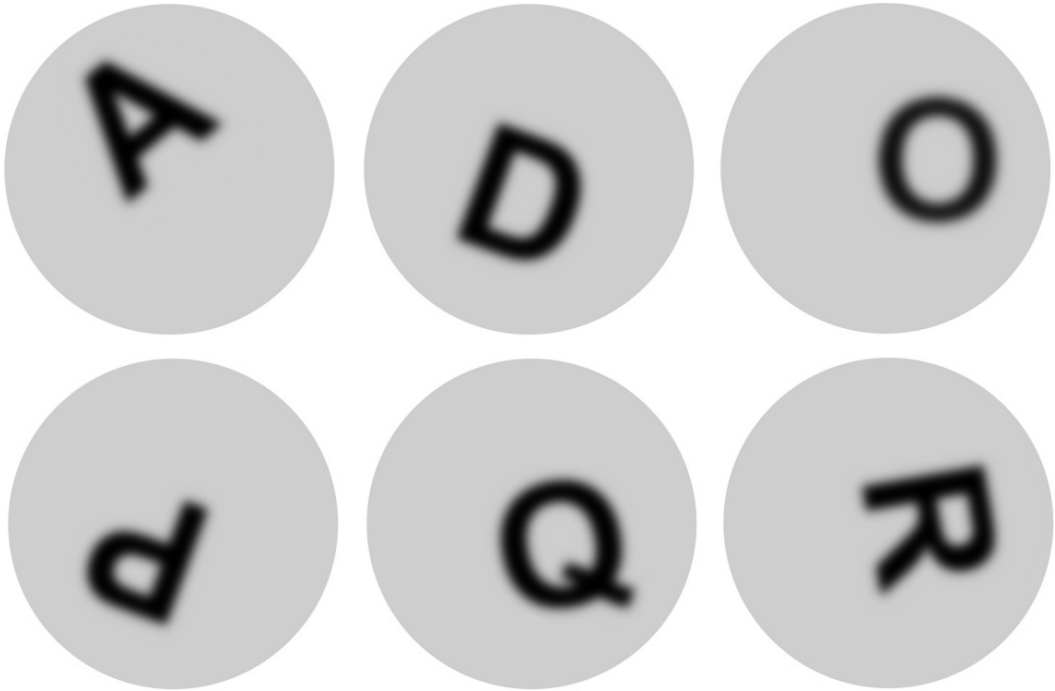

We have to restrict the invariance of persistent homology.

Couldn't we maintain classical persistent homology?

One could think of using other filtering functions, possibly defined on different topological spaces. For example, we could extract boundaries of letters and consider the distance from the center of mass of each boundary. This approach presents some drawbacks:

1. It "forgets" most of the information contained in the image $\varphi : \mathbb{R}^2 \to \mathbb{R}$ that we are considering, confining itself to examine the boundary of the letter represented by φ.
2. It usually requires an extra computational cost (e.g., to extract the boundaries of the letters).
3. It can produce a different topological space for each new filtering function (e.g., this happens for letters).
4. ABOVE ALL: 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.

Let us insert our goal in a general framework

Therefore, we need a more general method than just extracting boundaries and applying classical persistent homology.

Before proceeding in the generalization of persistent homology, let us extend our perspective and consider these two questions:

- What is our precise goal? Can we describe it formally?
- Can we measure in a metric way to which extent we have reached our objective?

In order to answer these questions, we need a geometrical model for shape comparison.

The key role of the observer

Shape comparison is based on comparing properties of perceptions.

Every comparison of properties involves the presence of

- an observer perceiving the properties
- a methodology to compare the properties

It follows that shape comparison is affected by subjectivity.

The key role of the observer

Truth often depends on the observer:

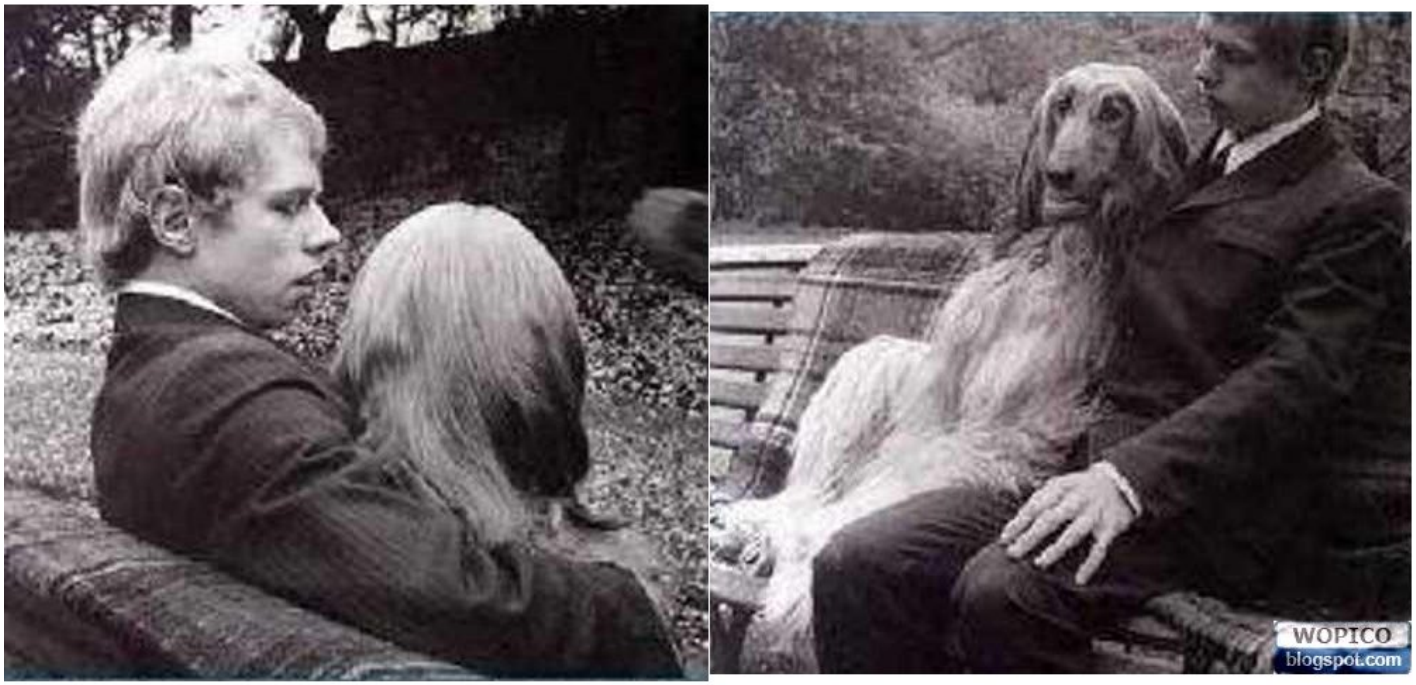

The key role of the observer

Truth often depends on the observer:

Impossible Ring and Pillars

by Guido Moretti

Our formal setting

- In shape comparison objects are not accessible directly, but only via measurements made by an observer.
- The comparison of two shapes is usually based on a family Φ of "measuring functions", which are defined on a set M (*set of measurements*) and take values in a set V (*set of measurement values*). Each function in Φ represents a measurement obtained via a measuring instrument.
- In most cases, the family Φ of measuring functions is invariant with respect to a given group G of transformations, that depends on the type of measurement we are considering.
- A G-invariant pseudo-metric d_G is usually available for the set Φ, so that we can quantify the difference between the measuring functions in Φ. *(pseudo-metric = metric without the property* $d(x, y) = 0 \implies x = y$*)*

The role of the observer in our formal setting

In our formal setting the choice of the observer is represented by the choice of the family Φ of functions (including its metric structure) and of the invariance given by the group G.

Indeed, different observers can

- get different perceptions/signals from the same phenomenon;
- see different equivalences between signals (i.e., refer to different invariance groups).

Example 1

A first way of "measuring" the vases: perceiving the color of each point

- M is the vase surface, $V = \mathbb{R}^3$ (the set of colors);
- Every $\varphi \in \Phi$ is a function associating each point of M with its color (represented by a triple (r, g, b) of real numbers);
- G is the set of rotations of M around the z-axis (we observe that $\Phi \circ G = \Phi$);
- We can set $d_G(\varphi_1, \varphi_2) = \inf_{g \in G} \sup_{x \in M} \|\varphi_1(x) - \varphi_2(g(x))\|$.

Example 2

A different way of "measuring" the vases: taking pictures

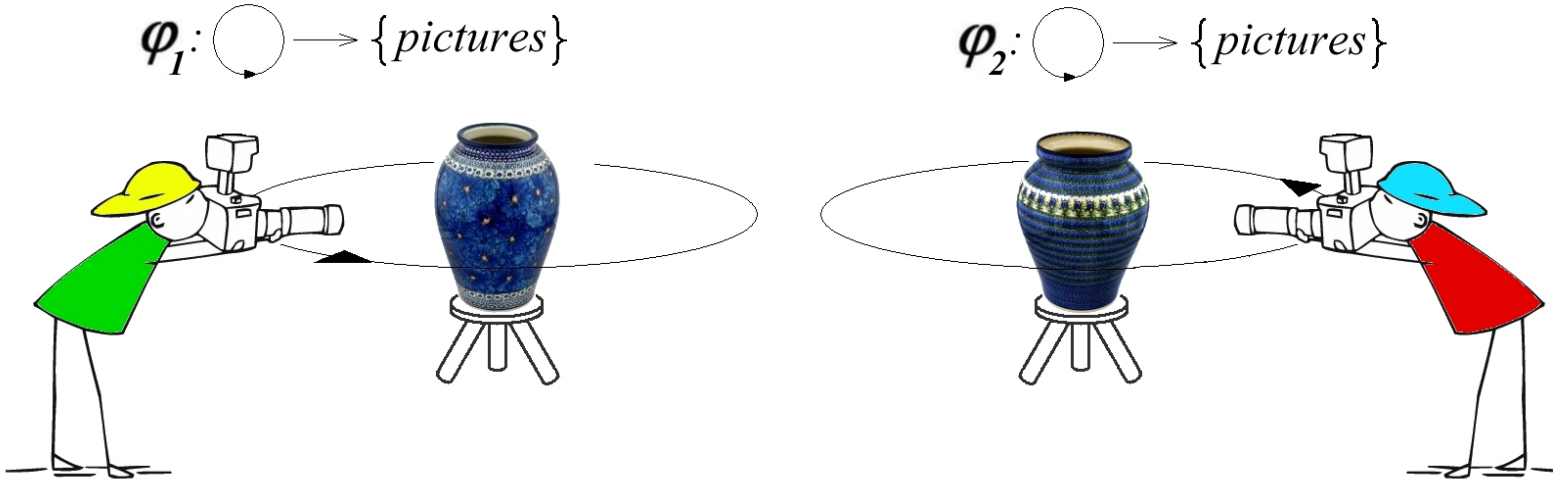

- $M = \mathbf{S}^1$, $V = \{\textit{the set of pictures}\}$
- Every $\varphi \in \Phi$ is a function associating each point of $\mathbf{S}^1$ with a picture.
- G is the set of rotations around the z-axis, and $\Phi \circ G = \Phi$
- We can set $d_G(\varphi_1, \varphi_2) = \inf_{g \in G} \sup_{x \in M} \|\varphi_1(x) - \varphi_2(g(x))\|$

Example 3

A third way of "measuring" the vases: weighing them

- $M =$ a singleton $\{\bar{x}\}$, $V = \mathbb{R}$
- Every $\varphi \in \Phi$ is a function associating the point in the singleton with a weight.
- G contains just the identity, and $\Phi \circ G = \Phi$
- We set

$$d_G(\varphi_1, \varphi_2) = \inf_{g \in G} \sup_{x \in M} \|\varphi_1(x) - \varphi_2(g(x))\| = |\varphi_1(\bar{x}) - \varphi_2(\bar{x})|$$

Our shape pseudo-distance d_G (formal definition)

Assume that the following objects are given:

- A set M, representing the set of measurements that we make.
- A set V, representing the values that can be taken by each measurement.
- A set Φ of functions from M to V. Each function $\varphi \in \Phi$ describes a possible set of results for all measurements in M.
- A group G acting on M, such that Φ is invariant with respect to G (i.e., for every $\varphi \in \Phi$ and every $g \in G$ we have that $\varphi \circ g \in \Phi$).
- A pseudo-metric d_G defined on the set Φ, that is invariant under the action of the group G (in other words, if $\varphi_1, \varphi_2 \in \Phi$ and $g \in G$ then $\varphi_2 \circ g \in \Phi$ and $d_G(\varphi_1, \varphi_2) = d_G(\varphi_1, \varphi_2 \circ g)$).

In some sense, the pair (Φ, d_G) represents the observer.

An interesting case

It often happens that M is a topological space and V is a metric space, endowed with a metric d_V. In this case the functions in F are assumed to be continuous, and the group G is assumed to be a subgroup of the group of all self-homeomorphisms of M. As an example, let us think of a CT scanning.

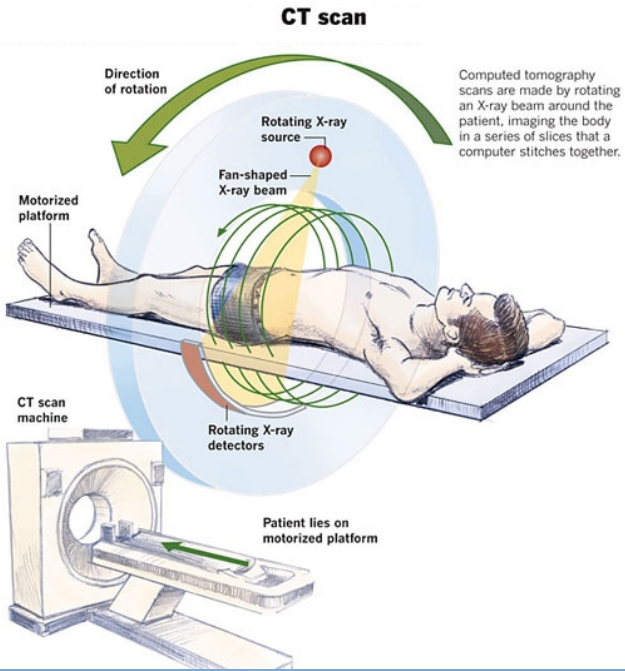

An interesting case

In this example

- $M = \mathbf{S}^1$ represents the topological space of all directions that are orthogonal to a given axis;
- $V = \mathbb{R}$ represents the metric space of all possible quantities of matter encountered by the X-ray beam in the considered direction.
- Every $\varphi \in \Phi$ is a function taking each direction in $\mathbf{S}^1$ to the quantity of matter encountered by the X-ray beam along that direction.
- G is the group of the rotations of $\mathbf{S}^1$ ($\Phi \circ G = \Phi$).

We can set $d_G(\varphi_1, \varphi_2) = \inf_{g \in G} \sup_{x \in M} |\varphi_1(x) - \varphi_2(g(x))|$ for $\varphi_1, \varphi_2 \in \Phi$.

Other interesting cases

If $V = \mathbb{R}^k$ we can use the pseudo-metric

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

for $\varphi_1, \varphi_2 \in \Phi$. The functional $\sup_{x \in M} \|\varphi_1(x) - \varphi_2(g(x))\|_\infty$ quantifies the change in the measurement induced by the transformation g.

The pseudo-metric d_G is produced by the attempt of minimizing this functional, varying the transformation g in the group G, and is called natural pseudo-distance.

Our assumptions

In the rest of this lecture we will assume that

1. M is a topological space X;
2. V is the metric space $\mathbb{R}$;
3. The functions $\varphi : X \to \mathbb{R}$ in Φ are continuous.

These assumptions allow us to require that if two measurements are close to each other (in some reasonable sense), then the values obtained by these measurements are close to each other, too. Without this kind of stability, our approach could not be of use in practical applications.

The group G will be assumed to be a subgroup of the group $\mathrm{Homeo}(M)$ of all self-homeomorphisms of M.

Natural pseudo-distance associated with a group G

We consider the case that our filtering functions are real-valued.

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.

Remark

PLEASE PAY ATTENTION:
THE TOPOLOGICAL SPACE X IS NOT "THE OBJECT" WE ARE STUDYING! IT IS JUST THE SPACE OF MEASUREMENTS.

Natural pseudo-distance associated with a group G

G-invariant persistent homology

G-invariant non-expansive operators

GIPHOD

We will not use Equivariant Homology

We need to apply persistent homology in a way that is invariant under the action of $G \subset \mathrm{Homeo}(X)$, but not under the one of $\mathrm{Homeo}(X)$.

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 (as happens for group of rotations of S^1), because in that case 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.

Filtering functions on a chain complex

Let (C, ∂) be a chain complex over a field $\mathbb{K}$ (so that each group of n-chains C_n is a vector space).

Definition

Assume a function $\bar{\varphi} : \bigcup_n C_n \to \mathbb{R} \cup \{-\infty\}$ is given, such that

i) $\bar{\varphi}$ takes the null chain $\mathbf{0} \in C_n$ to $-\infty$, for every $n \in \mathbb{Z}$;

ii) $\bar{\varphi}(\partial c) \leq \bar{\varphi}(c)$ for every $c \in \bigcup_n C_n$;

iii) $\bar{\varphi}(\lambda c) = \bar{\varphi}(c)$ for every $c \in \bigcup_n C_n$, $\lambda \in \mathbb{K}$, $\lambda \neq 0$;

iv) $\bar{\varphi}(c_1 + c_2) \leq \max(\bar{\varphi}(c_1), \bar{\varphi}(c_2))$ for every $c_1, c_2 \in C_n$ with $n \in \mathbb{Z}$.

We shall say that $\bar{\varphi}$ is a *filtering function on the chain complex* (C, ∂).

G-chain complexes

Let us recall the definition of G-chain complex.

Definition

Let us assume that a group G is given, such that G acts linearly on each vector space C_n and its action commutes with ∂, i.e., $\partial \circ g = g \circ \partial$ for every $g \in G$ (in particular, every $g \in G$ is a chain isomorphism from C to C). The chain complex (C, ∂) will be said a G-chain complex.

Chain subcomplex associated with a value $u \in \mathbb{R}$

Now, let us assume that (C, ∂) is a G-chain complex, endowed with a filtering function $\bar{\varphi}$.

For every $u \in \mathbb{R}$ we can consider the chain subcomplex $C^{\bar{\varphi} \leq u}$ of C defined by setting $C_n^{\bar{\varphi} \leq u} := \{c \in C_n : \bar{\varphi}(c) \leq u\}$ and restricting ∂ to $C^{\bar{\varphi} \leq u}$. $C^{\bar{\varphi} \leq u}$ is a subcomplex of C because of the properties in the definition of filtering function on a chain complex (in particular, $\partial(C_{n+1}^{\bar{\varphi} \leq u}) \subseteq C_n^{\bar{\varphi} \leq u}$). We observe that $C^{\bar{\varphi} \leq u}$ will not be a G-chain complex, since $g(C_n^{\bar{\varphi} \leq u}) \not\subseteq C_n^{\bar{\varphi} \leq u}$, in general.

Definition

The chain complex $\left(C^{\bar{\varphi} \leq u}, \partial\right)$ will be called the chain subcomplex of (C, ∂) associated with the value $u \in \mathbb{R}$, with respect to the filtering function $\bar{\varphi}$.

G-invariant persistent homology group of $\bar{\varphi}$

Definition

If $u, v \in \mathbb{R}$ and $u < v$, we can consider the inclusion i of the chain complex $C^{\bar{\varphi} \leq u}$ into the chain complex $C^{\bar{\varphi} \leq v}$. Such an inclusion induces a homomorphism $i_* : H_n\left(C^{\bar{\varphi} \leq u}\right) \to H_n\left(C^{\bar{\varphi} \leq v}\right)$. We shall call the group $PH_n^{\bar{\varphi}}(u, v) := i_*\left(H_n\left(C^{\bar{\varphi} \leq u}\right)\right)$ the n-th persistent homology group of the G-chain complex C, computed at the point (u, v) with respect to the filtering function $\bar{\varphi}$. The rank $\rho_n^{\bar{\varphi}}(u, v)$ of this group will be called *the n-th persistent Betti number function (PBNF) of the G-chain complex C, computed at the point (u, v) with respect to the filtering function $\bar{\varphi}$.*

$PH_n^{\bar{\varphi}}$ is invariant under the action of G

The key property of $PH_n^{\bar{\varphi}}$ is the invariance expressed by the following result.

Theorem

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

The previous theorem justifies the name G-invariant persistent homology, showing that the PBNFs of a G-chain complex do not change if we replace the filtering function $\bar{\varphi}$ with the function $\bar{\varphi} \circ g$, for $g \in G$.

A method to construct a suitable G-chain complex $(\bar{C}, \partial)$, endowed with a filtering function $\bar{\varphi}$

Let X and $(S(X), \partial)$ be a triangulable space and its singular chain complex over a field $\mathbb{K}$, respectively.

We recall that, by definition, every singular n-simplex in X is a continuous function from the standard n-simplex Δ_n into X.

Assume that a subgroup G of $\mathrm{Homeo}(X)$ and a continuous function $\varphi : X \to \mathbb{R}$ are chosen.

For every $u \in \mathbb{R}$, let us set $X^{\varphi \leq u} := \{x \in X : \varphi(x) \leq u\}$.

Let us consider the action of G on $S(X)$ defined by setting $g(\sigma) := g \circ \sigma$ for every $g \in G$ and every singular simplex σ in X, and extending this action linearly on $S(X)$.

$(S(X), \partial)$ is a G-chain complex.

Let us take a G-chain subcomplex $(\bar{C}, \partial)$ of the singular G-chain complex $(S(X), \partial)$

Now, let us assume that we are able to find a non-trivial G-chain subcomplex $(\bar{C}, \partial)$ of the singular G-chain complex $(S(X), \partial)$. This corresponds to select some **special** chains in $(S(X), \partial)$. This choice is a key point in our procedure.

In this case we can show that there is a natural way to endow $(\bar{C}, \partial)$ with a filtering function $\bar{\varphi}$.

We can endow $\bar{C}$ with a filtering function

Let $\{\sigma_j^n\}_{j\in J}$ be the set of all (distinct) singular n-simplexes in X. We can endow the chain complex $\bar{C}$ with a filtering function $\bar{\varphi}$ in the following way.

If c equals the null chain in $\bar{C}_n$, we set $\bar{\varphi}(c) := -\infty$.

If c is a non-null singular n-chain, we can write $c = \sum_{r=1}^{m} a^r \sigma_{j_r}^n \in \bar{C}_n$ with $a^r \in \mathbb{K}$, $a^r \neq 0$ for every index r, and $j_{r'} \neq j_{r''}$ for $r' \neq r''$. This representation is said to be reduced. Now, we set $\bar{\varphi}(c)$ equal to the smallest value u such that the corresponding sublevel set $X^{\varphi \leq u}$ contains the image of each singular simplex $\sigma_{j_r}^n$ involved in the reduced representation of c that we have considered. We observe that this representation is unique up to permutations of its summands, so that $\bar{\varphi}$ is well defined. We shall say that the function $\bar{\varphi}$ is induced by φ.

An important assumption

We observe that, for every topological subspace Y of X, $(\bar{C} \cap S(Y), \partial)$ is a chain complex over the field $\mathbb{K}$.
(The symbol $\bar{C} \cap S(Y)$ denotes the chain complex whose n-chains are the singular n-chains in Y that belong to $\bar{C}_n$.)

Before proceeding, we require that $\bar{C} \cap S(Y)$ verifies the property $(*)$ that is described in the following slide.
This property will allow us to prove the finiteness of the persistent Betti number functions of the G-chain complex $\bar{C}$.

An important assumption

In order to avoid "wild" chain complexes, we also make this assumption:

> $(*)$ If X' and X'' are two closed subsets of X with $X' \subseteq \mathrm{int}(X'')$, then a topological subspace $\hat{X}$ of X exists such that $X' \subseteq \hat{X} \subseteq X''$ and the homology group $H_n(\bar{C} \cap S(\hat{X}))$ is finitely generated for every non-negative integer n.

We wish to avoid chain complexes like the one where the 0-chains are all the usual singular 0-chains of X and the only 1-chain is the singular zero 1-chain of X.
We observe that $(*)$ is not as much an assumption about the regularity of the topological space X, but rather an assumption about the regularity of the G-chain complex.

Finiteness of the persistent Betti number functions

Property $(*)$ allows to prove the next result, which is analogous to the finiteness result proven for classical persistent homology.

Proposition

Assume $(*)$ holds. For every $n \in \mathbb{Z}$ the n-th persistent Betti number function $\rho_n^{\bar{\varphi}}(u,v)$ of the G-chain complex $(\bar{C}, \partial)$, endowed with the filtering function $\bar{\varphi}$, is finite at each point (u,v) in its domain.

The importance of assumption $(*)$

Remark

We stress the importance of the assumption $(*)$. It allows us to avoid chain complexes like the one where the 0-chains are all the usual singular 0-chains of X and the only 1-chain is the singular zero 1-chain of X. Obviously, this is a G-chain complex for any subgroup G of $\mathrm{Homeo}(X)$. In this case, for any pair (P_1, P_2) of distinct points of the topological space X, there is no singular 1-chain whose boundary is the singular 0-chain $P_2 - P_1$. Since the boundary homomorphism from 1-chains to 0-chains is zero, no non-zero 0-chain is a boundary. Hence the homology group $H_0(\bar{C})$ is not finitely generated, in general, and the property $(*)$ does not hold. For example, it does not hold for $X' = X'' = X$, independently of the regularity of the space X (unless X is a finite set). It is easy to check that persistent Betti numbers are not finite for the chain complex we have just described.

Stability of G-invariant persistent homology

From now on, in order to avoid technicalities that are not relevant in this lecture, we shall consider two PBNFs equivalent if they differ in a subset of their domain that has a vanishing measure.

A standard way of comparing two classical persistent Betti number functions is the matching distance d_{match}, a.k.a. bottleneck distance. (Here we identify persistent Betti number functions with the corresponding persistent diagrams.) It is important to observe that, in order to define it, we need the finiteness of the persistent Betti number functions.

This distance can be applied without any modification to the case of the persistent Betti number functions of the G-chain complex $\bar{C}$, because of the previously stated finiteness.

Stability of G-invariant persistent homology

The following theorem shows that the matching distance between persistent Betti number functions of the G-chain complex $\bar{C}$ is a lower bound for the natural pseudo-distance d_G. In other words, a small change of the filtering function with respect to d_G produces just a small change of the corresponding persistent Betti number function with respect to d_{match}. This property allows the use of PBNFs in real applications, where the presence of noise is unavoidable.

Theorem

For every $n \in \mathbb{Z}$, let us consider the n-th persistent Betti number functions $\rho_n^{\bar{\varphi}}$, $\rho_n^{\bar{\psi}}$ of the G-chain complex $(\bar{C}, \partial)$, endowed with the filtering functions $\bar{\varphi}$ and $\bar{\psi}$ induced by $\varphi : X \to \mathbb{R}$ and $\psi : X \to \mathbb{R}$, respectively. Then

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

An example illustrating our method

Let us consider an experimental setting where a robot is in the middle of a room, measuring its distance from the surrounding walls by a sensor, for each oriented direction. This measurement can be formalized by a function $\xi : \mathbf{S}^1 \to \mathbb{R}$, where $\xi(v)$ equals minus the distance from the wall in the oriented direction represented by the unit vector v, for each $v \in \mathbf{S}^1$.

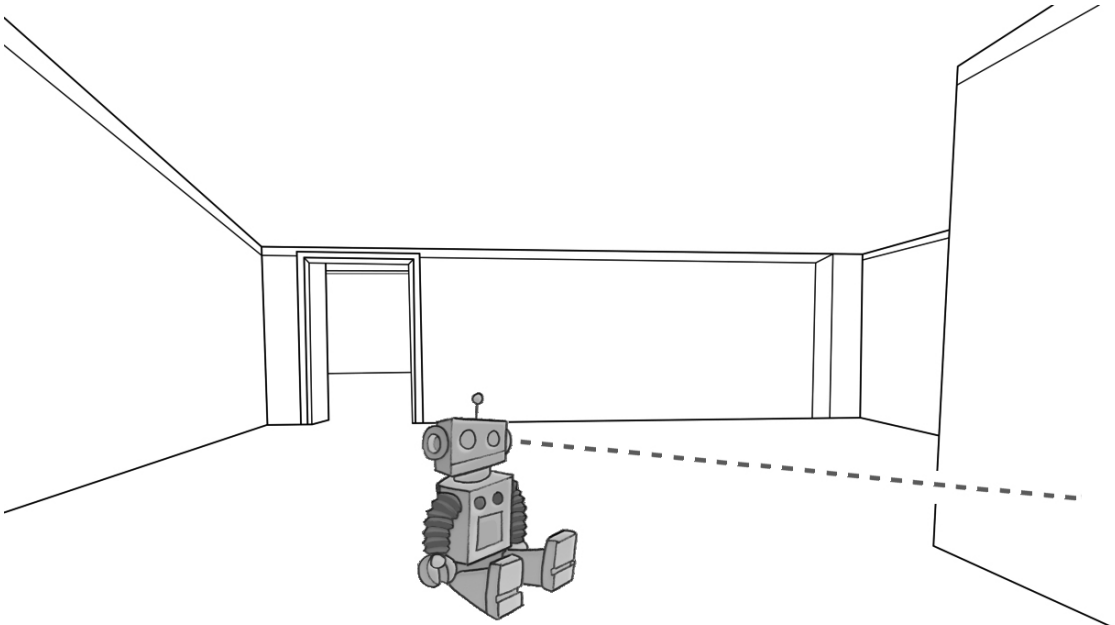

An example illustrating our method

Here are two instances φ and ψ of the function ξ for two different shapes of the room.

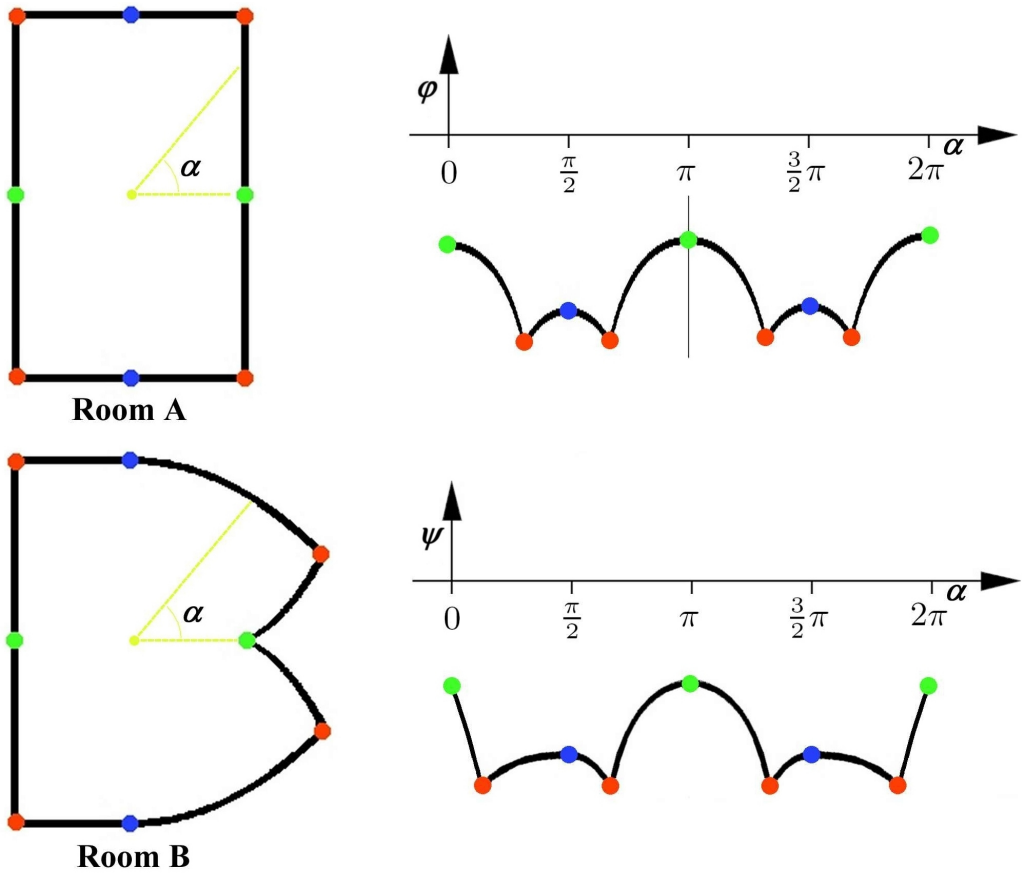

An example illustrating our method

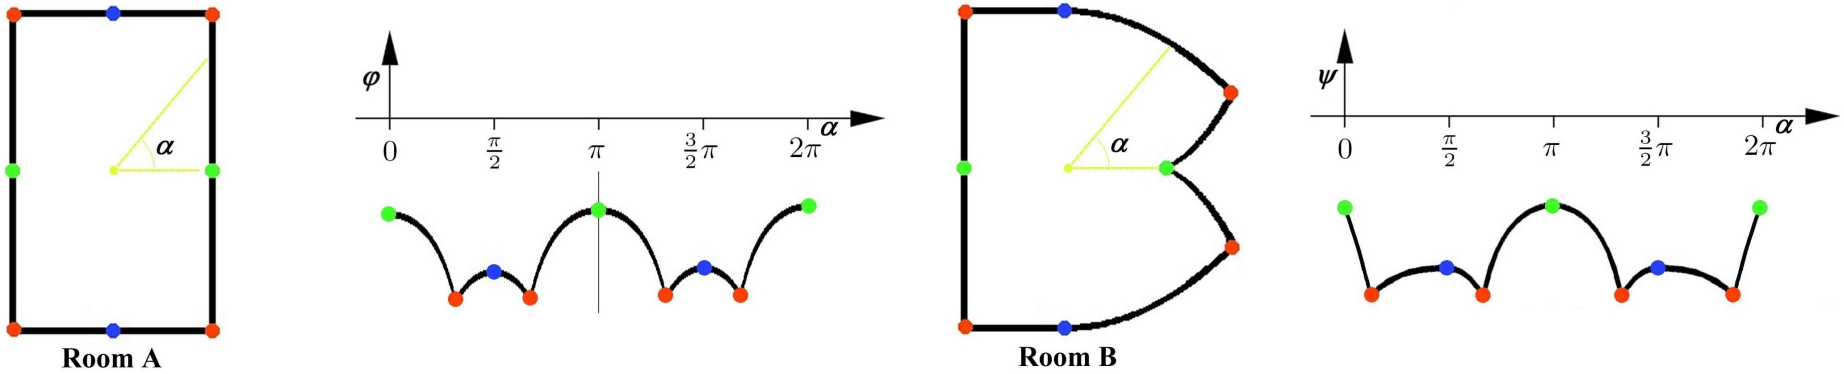

Let $R(\mathbf{S}^1)$ denote the group of orientation-preserving rigid motions of $\mathbf{S}^1 \subset \mathbb{R}^2$. We observe that a homeomorphism $f : \mathbf{S}^1 \to \mathbf{S}^1$ exists, such that $\varphi = \psi \circ f$ and $f \notin R(\mathbf{S}^1)$. It follows that $d_{\mathrm{Homeo}(\mathbf{S}^1)}(\varphi, \psi) = 0$, so that the direct application of classical persistent homology does not give a positive lower bound for $d_{R(\mathbf{S}^1)}(\varphi, \psi)$, while we will see that $d_{R(\mathbf{S}^1)}(\varphi, \psi) > 0$.

An example illustrating our method

We can consider the chain complex $\bar{C}$ whose n-chains are all the singular n-chains $c \in S_n(\mathbf{S}^1)$ for which the following property holds:

(P) If a singular simplex σ_i^n appears in a reduced representation of c with respect to the basis $\{\sigma_j^n\}$ of $S_n(\mathbf{S}^1)$, then the antipodal simplex $s \circ \sigma_i^n$ appears in that representation with the same multiplicity of σ_i^n, where s is the antipodal map $s : \mathbf{S}^1 \longrightarrow \mathbf{S}^1$.

In other words, in $\bar{C}$ we accept by definition only the singular chains in $\mathbf{S}^1$ that can be written in the form $\sum_{r=1}^{m} a^r \left(\sigma_{j_r}^n + s \circ \sigma_{j_r}^n\right)$.

It easy to check that $(\bar{C}, \partial)$ is a $R(\mathbf{S}^1)$-chain subcomplex of the complex $(S(\mathbf{S}^1), \partial)$. Every rotation $\rho \in R(\mathbf{S}^1)$ commutes with the antipodal map s and is a chain isomorphism from $\bar{C}$ to $\bar{C}$.
Property $(*)$ holds for the $R(\mathbf{S}^1)$-chain complex that we have defined.

An example illustrating our method

Referring to the two-rooms example, let us consider the birth of the first homology class in the homology groups $H_0\left(\bar{C}^{\bar{\varphi}\leq t}\right)$ and $H_0\left(\bar{C}^{\bar{\psi}\leq t}\right)$, respectively, when the parameter t increases.

While the group $H_0\left(\bar{C}^{\bar{\varphi}\leq t}\right)$ becomes non-trivial when t reaches the value $t_0 = \min\varphi = \min\psi$, the group $H_0\left(\bar{C}^{\bar{\psi}\leq t}\right)$ becomes non-trivial when t reaches a value $\bar{t} > \min\varphi = \min\psi$.

This is due to the fact that the sublevel set $\{x \in \mathbf{S}^1 : \varphi(x) \leq t_0\}$ contains two pairs of antipodal points, while the sublevel set $\{x \in \mathbf{S}^1 : \psi(x) \leq t_0\}$ contains no pair of antipodal points (see next figure).

An example illustrating our method

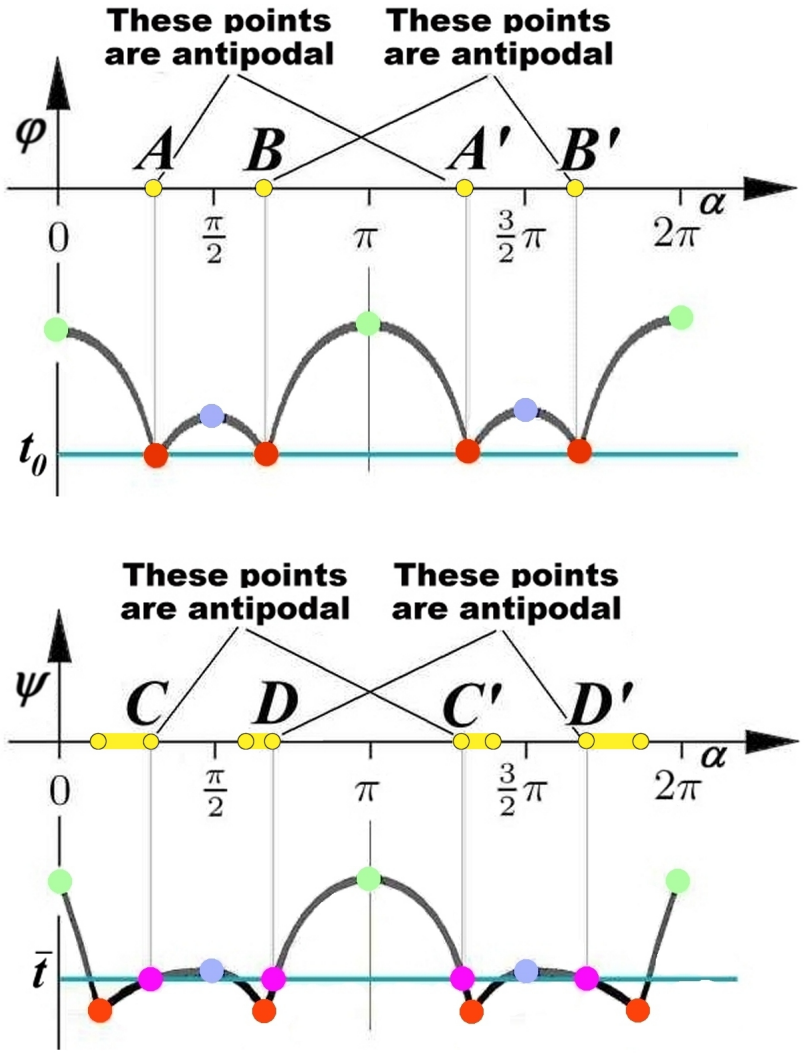

An example illustrating our method

Therefore, the only points at infinity in the persistence diagrams associated with the 0-th persistent homology groups of the G-chain subcomplex $\bar{C}$ of $S(\mathbf{S}^1)$ with respect to $\bar{\varphi}$ and $\bar{\psi}$ are (t_0, ∞) and $(\bar{t}, \infty)$, respectively.

It follows that the matching distance between the 0-th persistent Betti number functions of the $R(\mathbf{S}^1)$-chain complex $\bar{C}$ with respect to the filtering functions $\bar{\varphi}$ and $\bar{\psi}$ is at least $\bar{t} - t_0 > 0$.

By applying the Stability Theorem, we obtain the inequality $d_{R(\mathbf{S}^1)}(\varphi, \psi) \geq \bar{t} - t_0$. In other words, G-invariant persistent homology gives a non-trivial lower bound for $d_{R(\mathbf{S}^1)}(\varphi, \psi)$, while the matching distance between the classical persistent Betti number functions with respect to the filtering functions φ and ψ vanishes.

A general procedure to construct $\bar{C}$

The procedure to construct the chain complex $\bar{C}$ that we have illustrated can be generalized to triangulable spaces different from $\mathbf{S}^1$ and invariance groups G that are different from the group of rotations. The main idea consists in looking for another subgroup H of $\mathrm{Homeo}(X)$ 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 $h \in H$.

Due to the finiteness of H, the property 2 implies that the restriction to H of the conjugacy action of each $g \in G$ is a permutation of H. The legitimate n-chains in our chain complex $\bar{C}$ are defined to be the linear combinations of "elementary" singular chains c that can be written as $c = \sum_{i=1}^{r} h_i \circ \sigma$, where $\sigma : \Delta_n \to X$ is a singular n-simplex in X.

A general procedure to construct $\bar{C}$

Because of the property 2 and the linearity of the action of each $g \in G$, $g\left(\sum_{i=1}^{r} h_i \circ \sigma\right) = \sum_{i=1}^{r} g \circ h_i \circ \sigma = \sum_{i=1}^{r} (g \circ h_i \circ g^{-1}) \circ (g \circ \sigma) = \sum_{i=1}^{r} h_i \circ (g \circ \sigma)$ is another legitimate chain in our chain complex $\bar{C}$, so that $\bar{C}$ results to be a G-chain complex.

In the two-rooms example, we have chosen $H = \{id, s\} \subset G = R(\mathbf{S}^1)$, where s is the antipodal simmetry. We recall that the filtering function $\varphi : X \to \mathbb{R}$ induces a filtering function $\bar{\varphi}$ on the set of legitimate chains, where $\bar{\varphi}(c)$ is the smallest value u such that the corresponding sublevel set $X^{\varphi \leq u}$ contains the image of each singular simplex involved in a reduced representation of c, for every non-null chain $c \in \bar{C}_n$.

A general procedure to construct $\bar{C}$

If G is Abelian, a simple way to get a subgroup H of $\mathrm{Homeo}(X)$ verifying properties 1 and 2 consists in setting H equal to a finite subgroup of G. This is exactly what we did in the two-rooms example, setting $H = \{id, s\} \subset G = R(\mathbf{S}^1)$.
If G is finite, a trivial way to get a subgroup H of $\mathrm{Homeo}(X)$ verifying properties 1 and 2 consists in setting $H = G$. This choice leads to consider the quotient space X/G, provided that G acts freely on X. However, our approach is more general. Indeed, in the two-rooms example, if we set G equal to the (Abelian and finite) group generated by the reflections with respect to the coordinate axes, we could choose H equal to the group generated by the rotation of $2\pi/m$ radians. In this case, if the homeomorphism g reverses the orientation, then the conjugacy action $h \mapsto g \circ h \circ g^{-1}$ is not the identity, since it takes each h to its inverse h^{-1}. Furthermore, $H \not\subseteq G$.

Natural pseudo-distance associated with a group G

G-invariant persistent homology

G-invariant non-expansive operators

GIPHOD

G-invariant non-expansive operators

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 $\hat{\varphi} : \frac{X}{H} \longrightarrow \mathbb{R}$ 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.

G-invariant non-expansive operators

Fortunately, an alternative approach is available.

We will describe it in the next part of this lecture.

This part is based on an ongoing joint research project with

Grzegorz Jabłoński and Marc Ethier
Jagiellonian University - Kraków

G-invariant non-expansive 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.

F_1

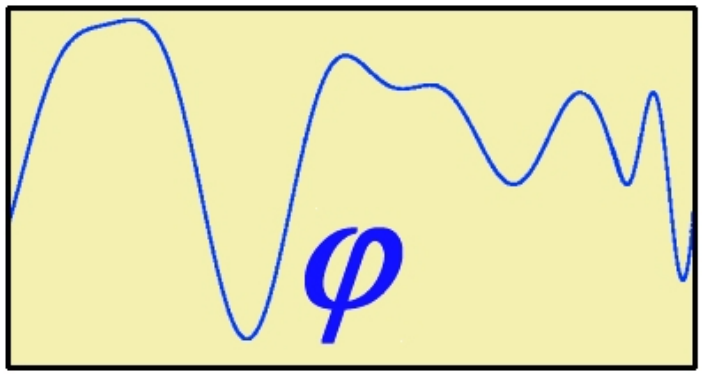

F_2

G-invariant non-expansive operators

Let us consider the following objects:

- A triangulable space X with nontrivial homology in degree k.
- A set Φ of continuous functions from X to $\mathbb{R}$, that contains the set of all constant functions.
- A topological subgroup G of $\mathrm{Homeo}(X)$ that acts on Φ by composition on the right.
- The natural pseudo-distance d_G on Φ 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 $\varphi_1, \varphi_2 \in \Phi$.
- The distance d_∞ on Φ, 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 $\mathscr{F}$ 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.

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 $\max \varphi$;
- $F(\varphi) := \max\left\{\varphi\left(x - \frac{\pi}{8}\right), \varphi\left(x + \frac{\pi}{8}\right)\right\}$;
- $F(\varphi) := \frac{1}{2}\left(\varphi\left(x - \frac{\pi}{8}\right) + \varphi\left(x + \frac{\pi}{8}\right)\right)$.

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

For every $\varphi_1, \varphi_2 \in \Phi$ we set

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

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

Proposition

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

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

An equivalence result

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

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

Our main idea

The previous theorem suggests to study $D^{\mathscr{F}}_{\text{match}}$ instead of d_G.

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

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

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

Obviously, $D^{\mathscr{F}^*}_{\text{match}} \leq D^{\mathscr{F}}_{\text{match}}$.

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

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, d_∞) is a compact metric space, 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 $\mathscr{F}^{\mathrm{all}}(\Phi, G)$

This statement follows:

Corollary

Assume that the metric space (Φ, d_∞) is compact. Let $\mathscr{F}$ be a subset of $\mathscr{F}^{\mathrm{all}}(\Phi, G)$. For every $\varepsilon > 0$, a finite subset $\mathscr{F}^*$ of $\mathscr{F}$ exists, such that

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

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

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

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 $\mathbf{S}^1$ 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.

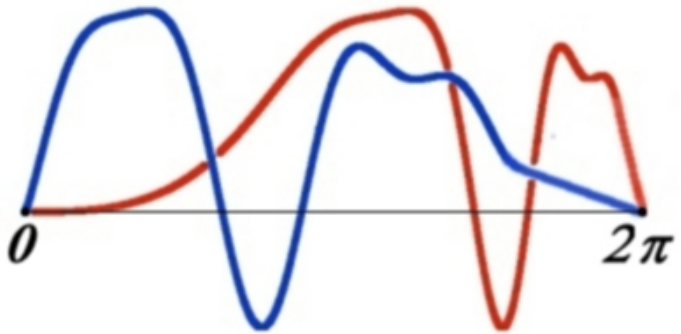

Let us check what happens in practice

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.

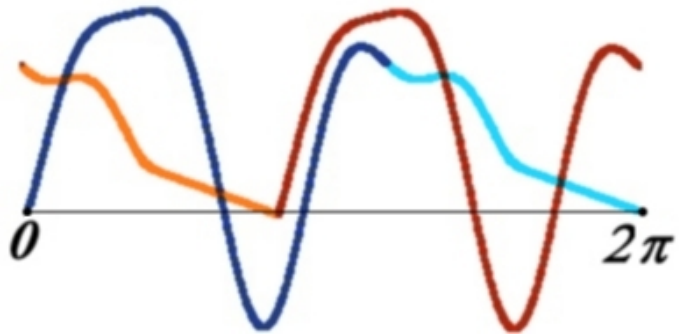

Finally, the choice of $\mathbf{I}(\mathbf{S}^1) = \{id\}$ as an invariance group means that two functions are considered equivalent if and only if they coincide everywhere.

Theoretical results

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

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

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

An important remark

It is important to use several operators. The use of just one operator still produces a pseudo-distance $D^{\mathscr{F}^*}_{match}$ 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:

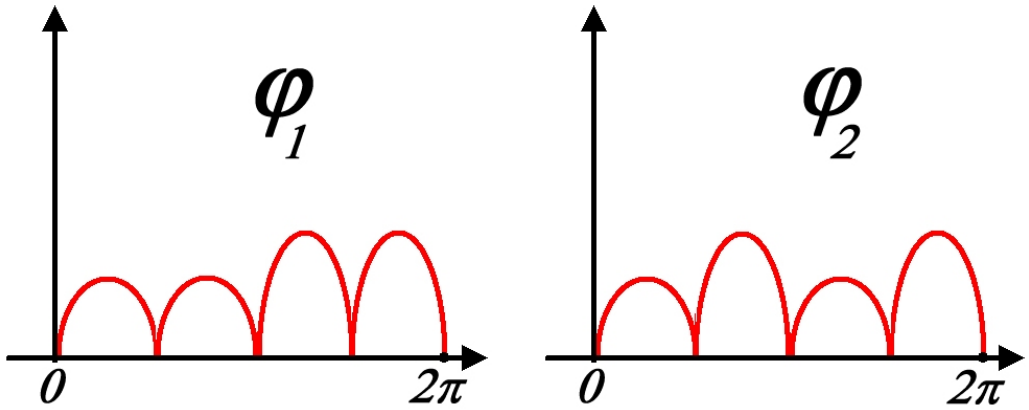

The results of an experiment: the group Homeo($\mathbf{S}^1$)

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
1.648	4.831

Standard deviation
0.934

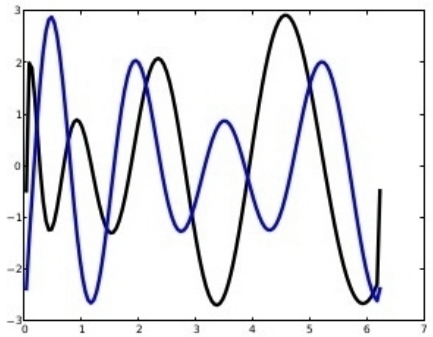

(a) φ_{516}, dist: 0.0465393

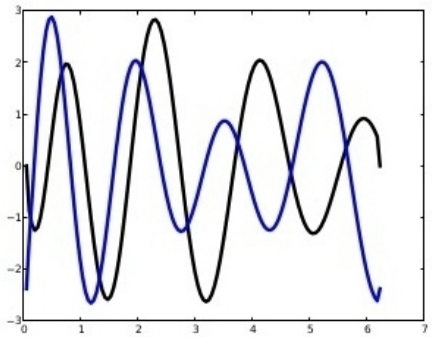

(b) φ_{381}, dist: 0.0541687

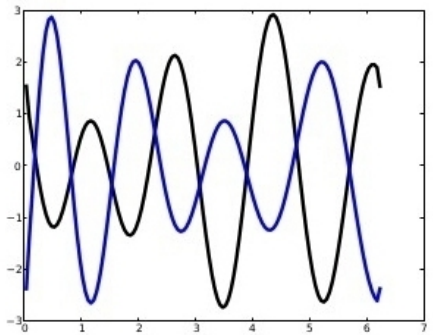

(c) φ_{7776}, dist: 0.0984192

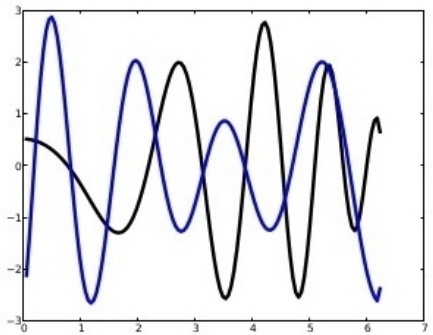

(d) φ_{6214}, dist: 0.10376

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

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

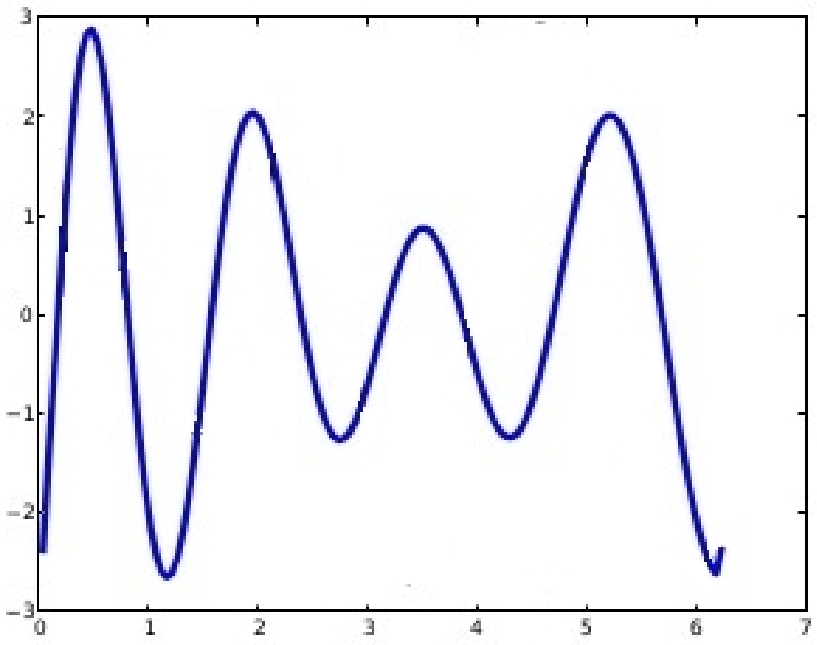

The results of an experiment: the group Homeo($\mathbf{S}^1$)

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

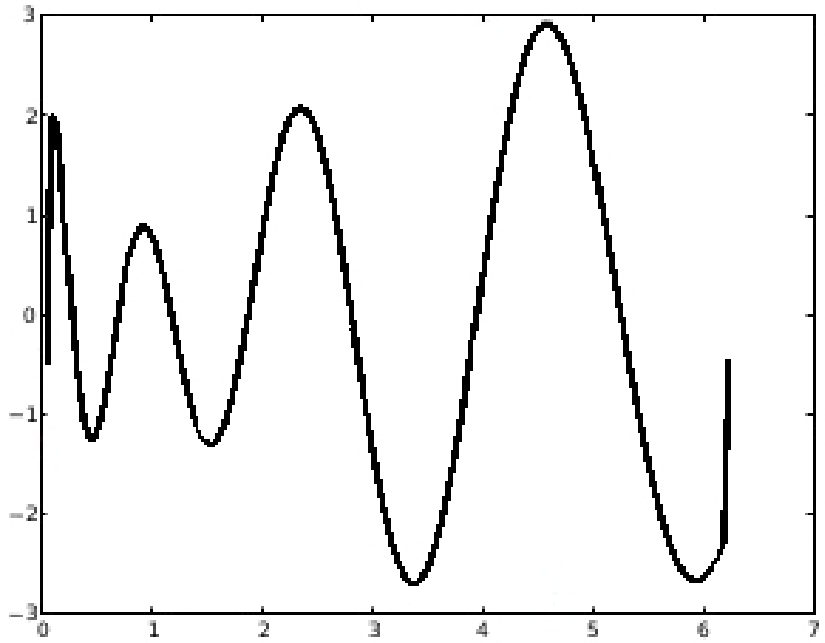

The results of an experiment: the group Homeo($\mathbf{S}^1$)

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.

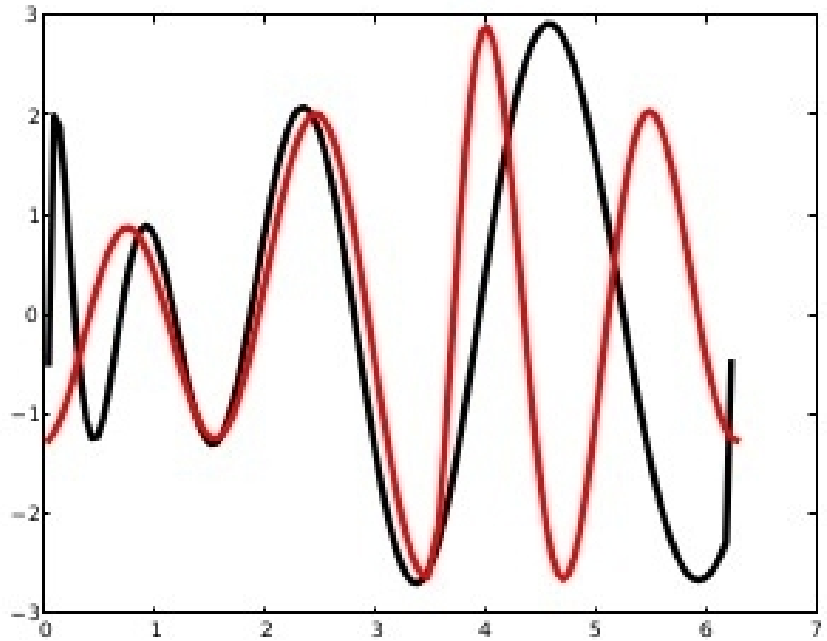

The results of an experiment: the group Homeo($\mathbf{S}^1$)

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**.

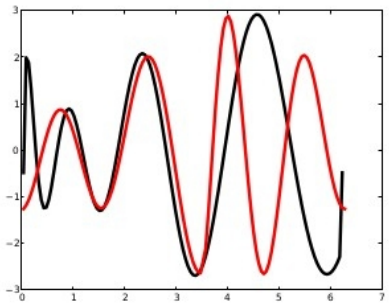

(a) φ_{516}, dist: 0.0465393

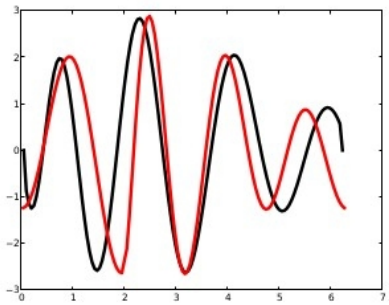

(b) φ_{381}, dist: 0.0541687

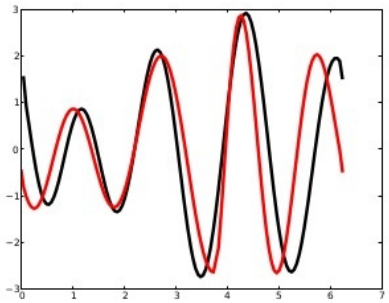

(c) φ_{7776}, dist: 0.0984192

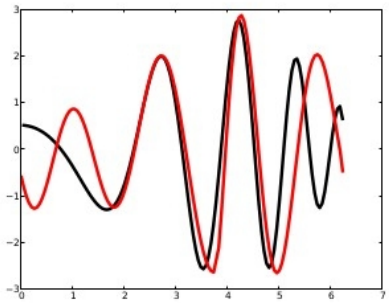

(d) φ_{6214}, dist: 0.10376

Theoretical results

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:

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

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

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

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

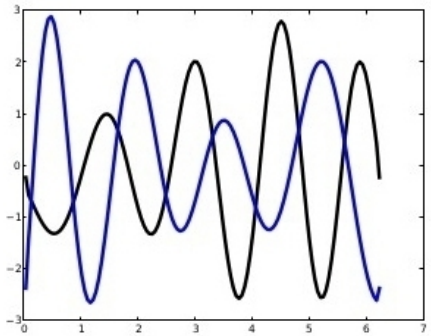

(a) φ_{5566}, dist: 0.333405

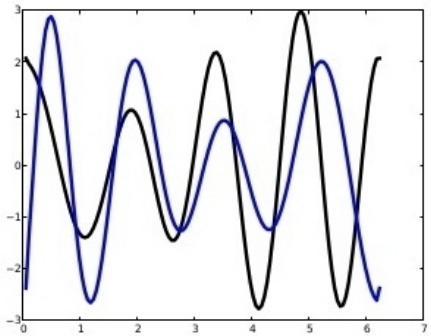

(b) φ_{8454}, dist: 0.422668

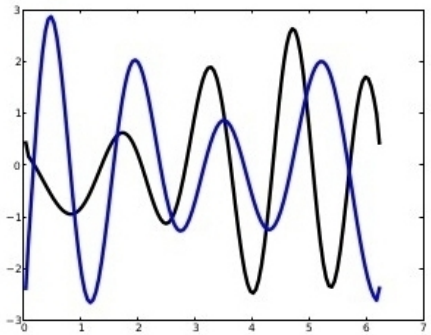

(c) φ_{8909}, dist: 0.453949

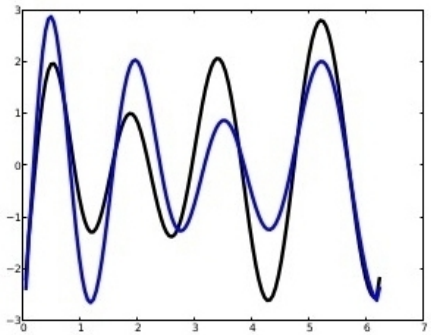

(d) φ_{4426}, dist: 0.46463

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:

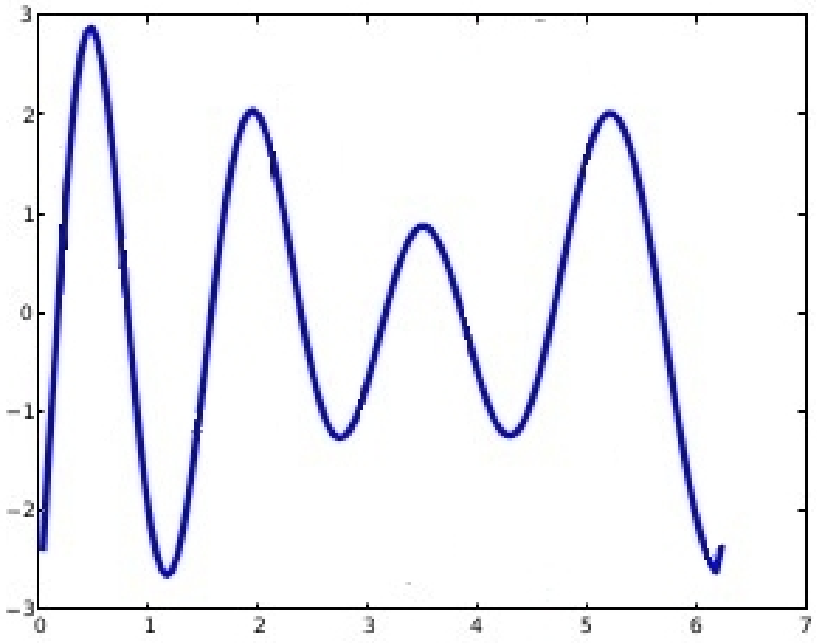

The results of an experiment: the group $R(\mathbf{S}^1)$

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

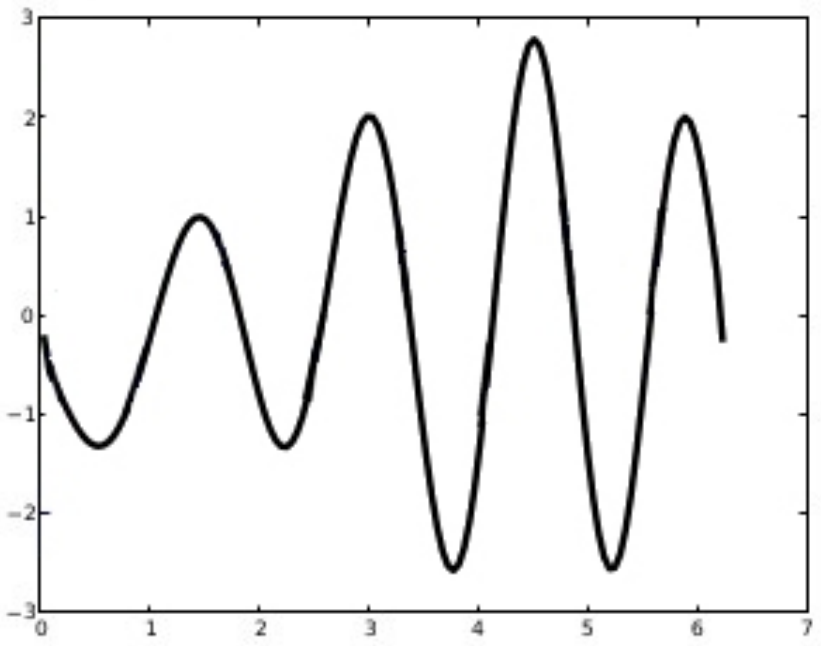

The results of an experiment: the group $R(\mathbf{S}^1)$

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.

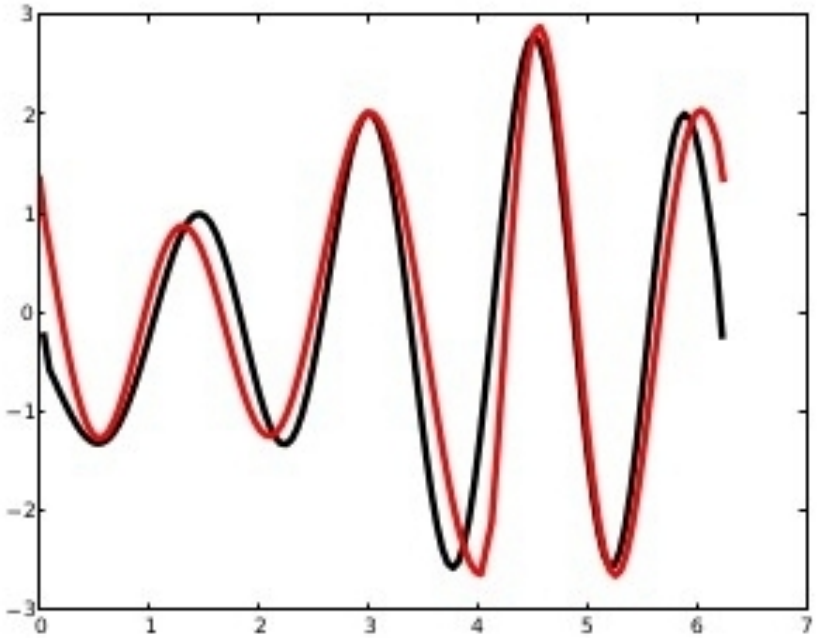

The results of an experiment: the group $R(\mathbf{S}^1)$

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**.

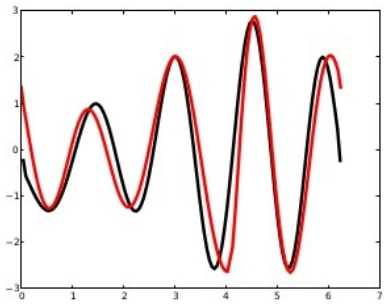

(a) φ_{5566}, dist: 0.333405

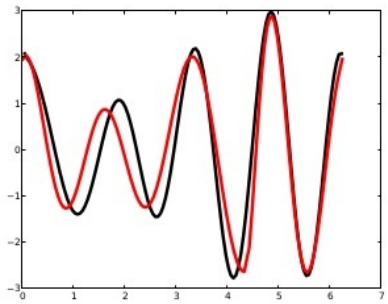

(b) φ_{8454}, dist: 0.422668

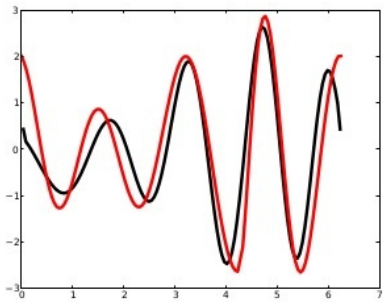

(c) φ_{8909}, dist: 0.453949

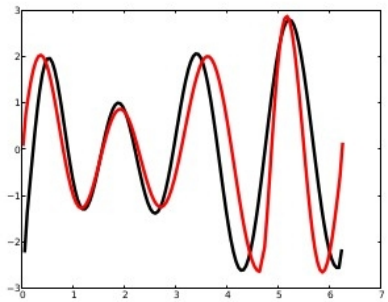

(d) φ_{4426}, dist: 0.46463

Theoretical results

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.

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 not $R(\mathbf{S}^1)$-operators):

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

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

This choice produces the pseudo-distance

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

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

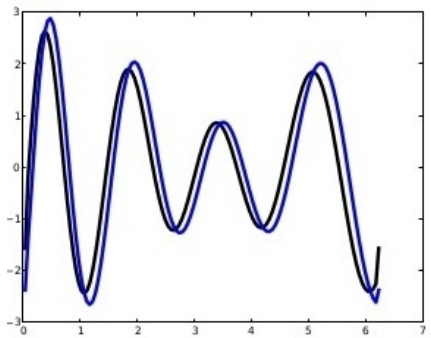

(a) φ_{7133}, dist: 0.415802

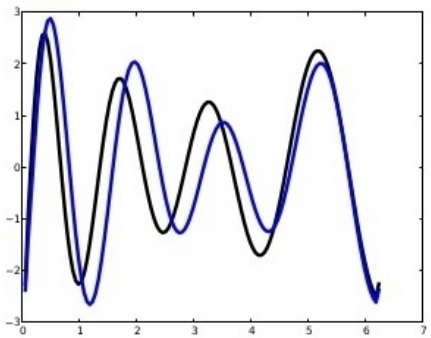

(b) φ_{7001}, dist: 0.598145

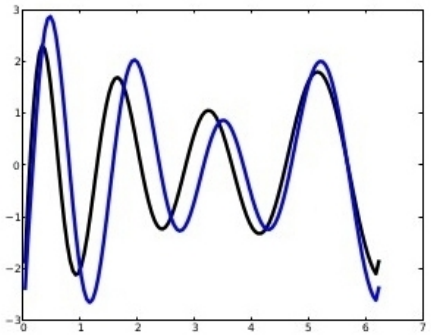

(c) φ_{389}, dist: 0.617218

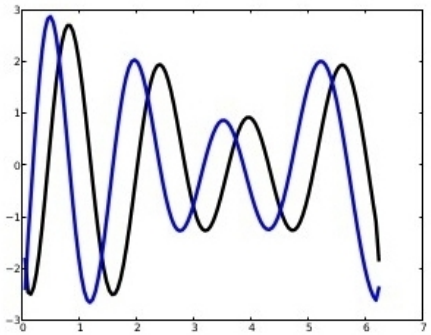

(d) φ_{5723}, dist: 0.617981

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:

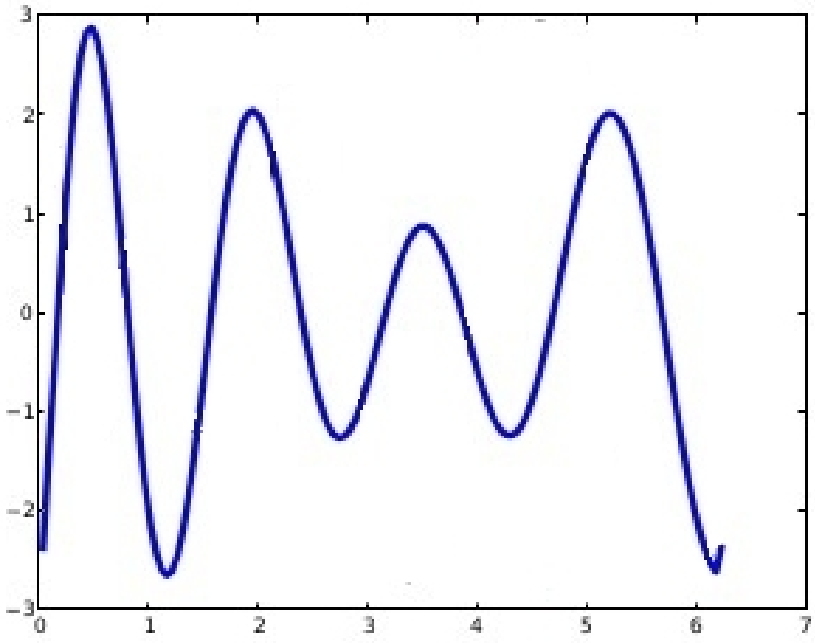

The results of an experiment: the group $\mathbf{I}(\mathbf{S}^1)$

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

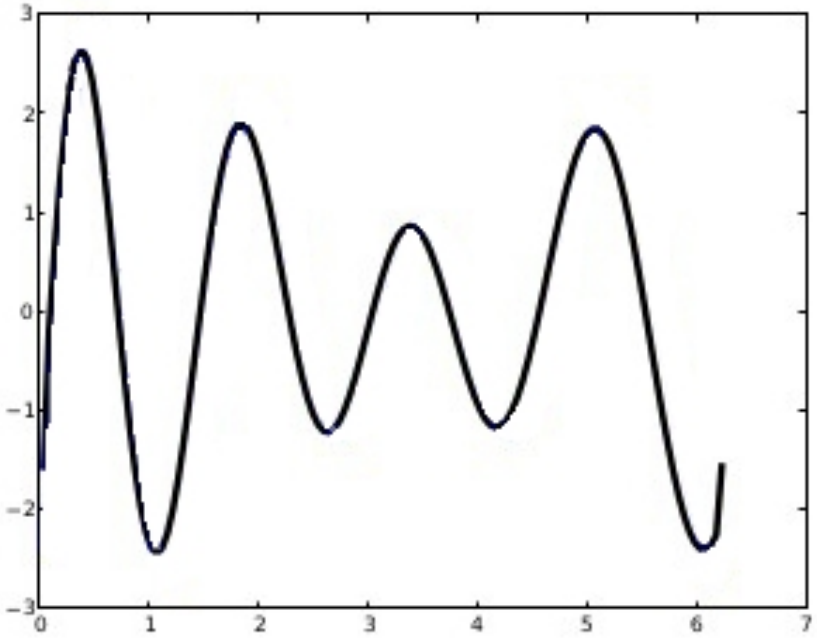

The results of an experiment: the group $\mathbf{I}(\mathbf{S}^1)$

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.

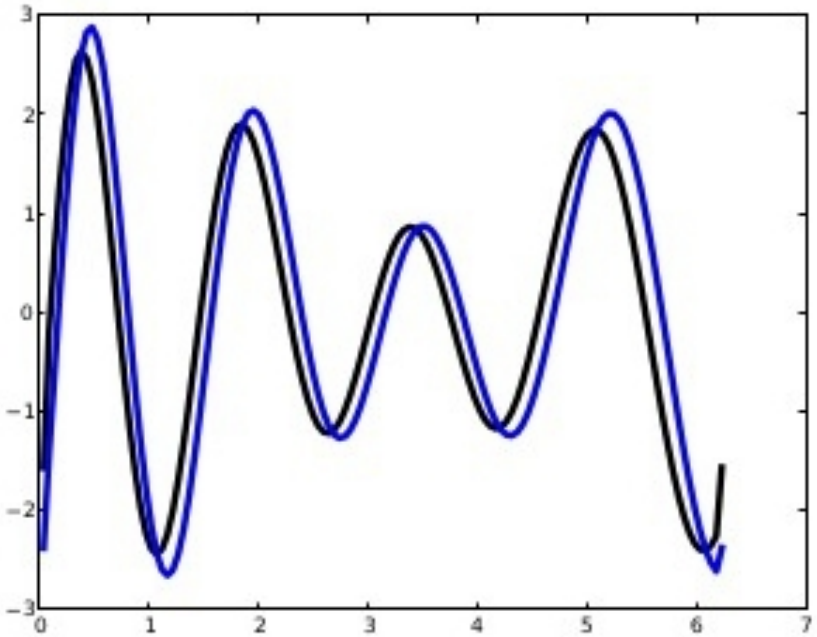

The results of an experiment: the group $\mathbf{I}(\mathbf{S}^1)$

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.

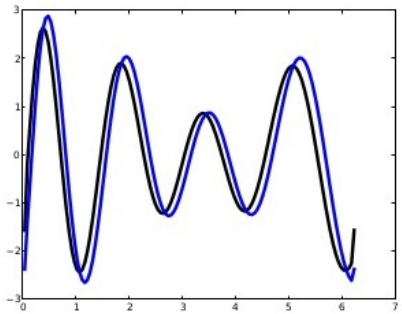

(a) φ_{7133}, dist: 0.415802

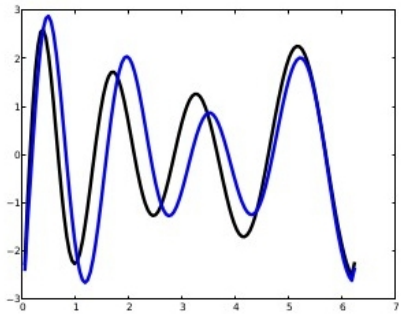

(b) φ_{7001}, dist: 0.598145

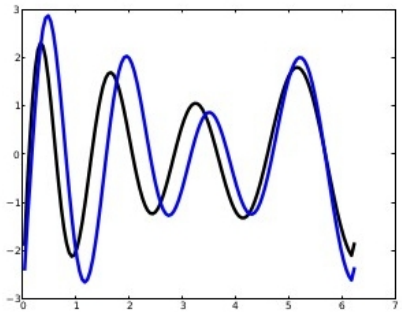

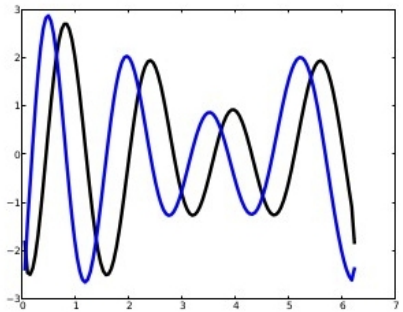

An open problem

We have proven that if Φ is compact, then $D^{\mathscr{F}}_{match}$ can be approximated computationally.

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

Further research is needed in this direction.

Natural pseudo-distance associated with a group G

G-invariant persistent homology

G-invariant non-expansive operators

GIPHOD

GIPHOD

GIPHOD: joint project with Grzegorz Jabłoński and Marc Ethier (Jagiellonian University - Kraków)

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 G-invariant persistent homology for image comparison.

Dataset: 10.000 grey-level images obtained by adding randomly chosen bell-shaped functions.

GIPHOD SHOULD BE AVAILABLE IN THE NEXT FEW MONTHS.

GIPHOD (Group Invariant Persistent Homology On-line Demonstrator)

We are going to show the results of an experiment where the invariance group G is the group of isometries:

Some data about the pseudo-metric $D^{\mathscr{F}}_{match}$ in this case:

- The images are coded as functions from $\mathbb{R}^2 \longrightarrow [0,1]$;
- Mean distance between images: 0.35752;
- Standard deviation of distance between images: 0.14881;
- Number of GINOs that have been used: 12.

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 $\int_{\mathbb{R}^2} \varphi(\mathbf{x})\ d\mathbf{x}$.
- $F(\varphi)$ defined by setting

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

  where $\beta : \mathbb{R} \to \mathbb{R}$ is an integrable function with $\int_{\mathbb{R}^2} \left|\beta\left(\|\mathbf{y}\|_2\right)\right|\ d\mathbf{y} \leq 1$. Four GINOs of this kind have been used.
- The opposite operators $-F$ of the six previous GINOs.

GIPHOD: Examples for the group of isometries

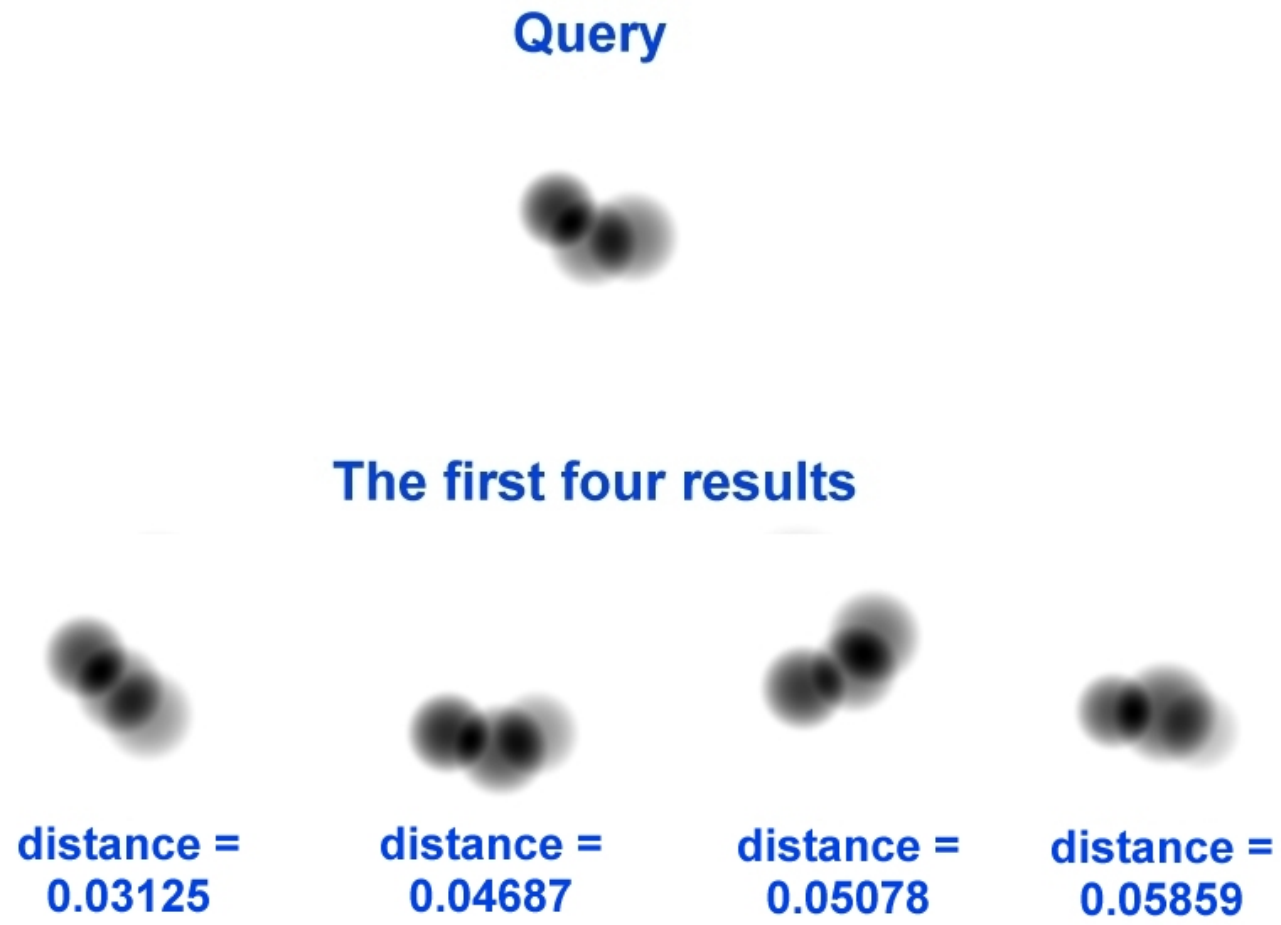

GIPHOD: Examples for the group of isometries

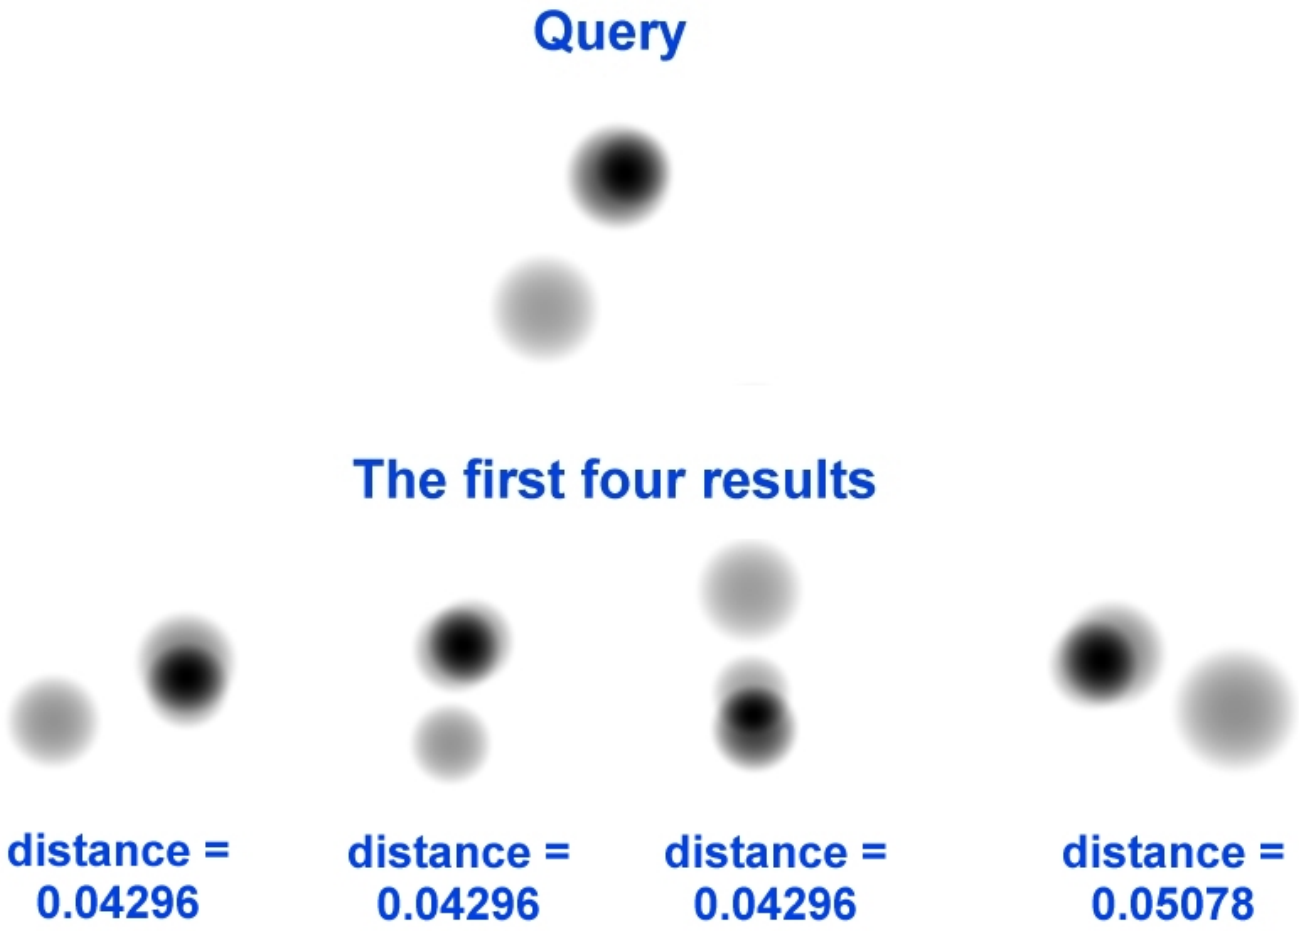

GIPHOD: Examples for the group of isometries

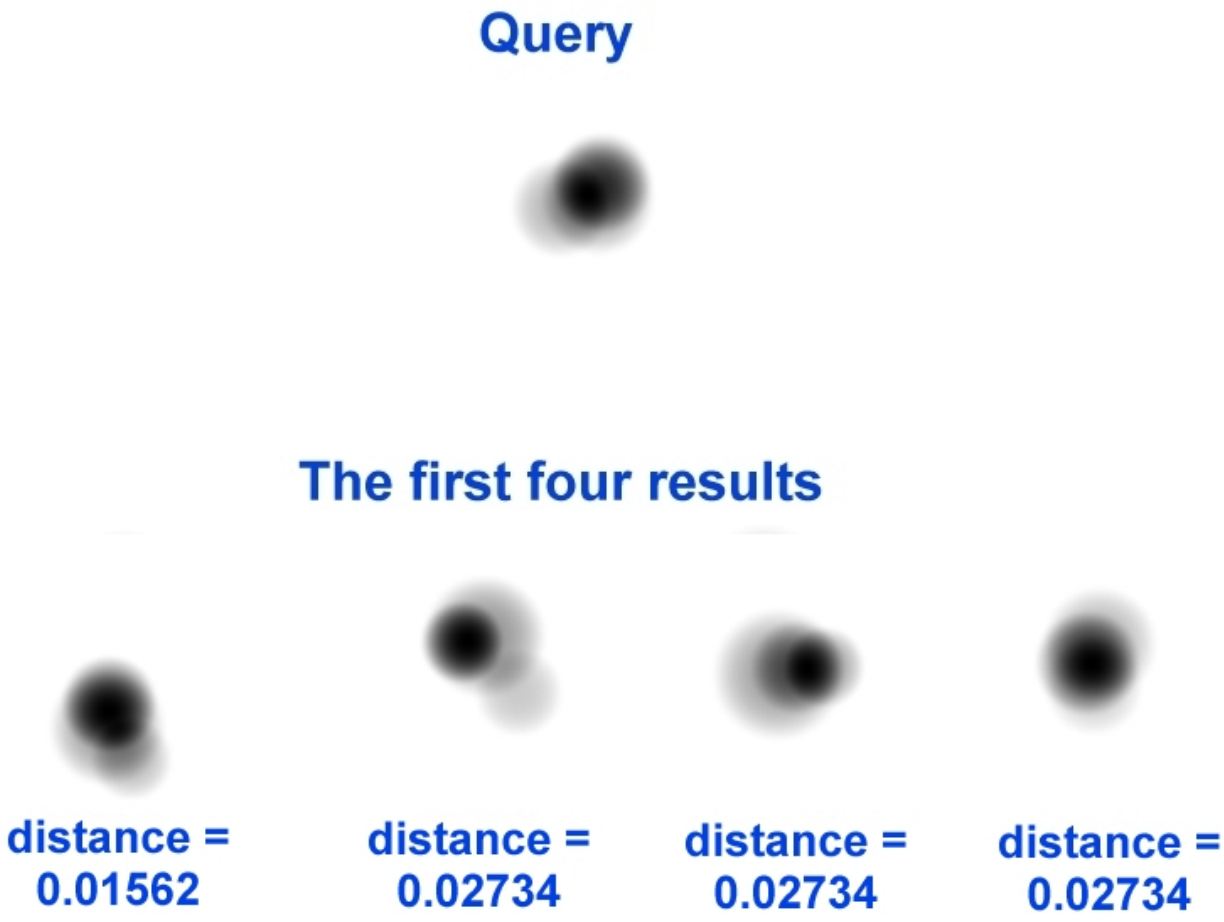

GIPHOD: Examples for the group of isometries

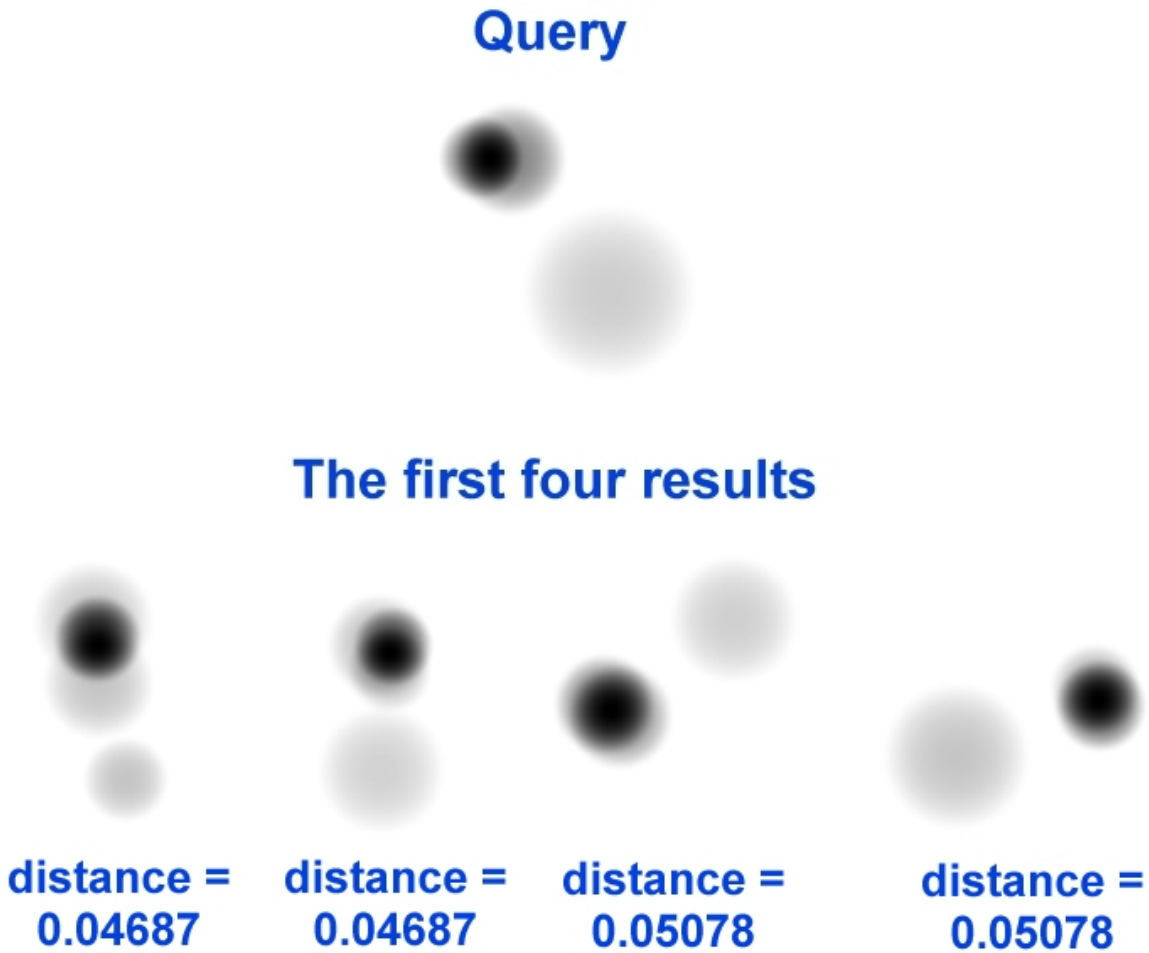

GIPHOD: Examples for the group of isometries

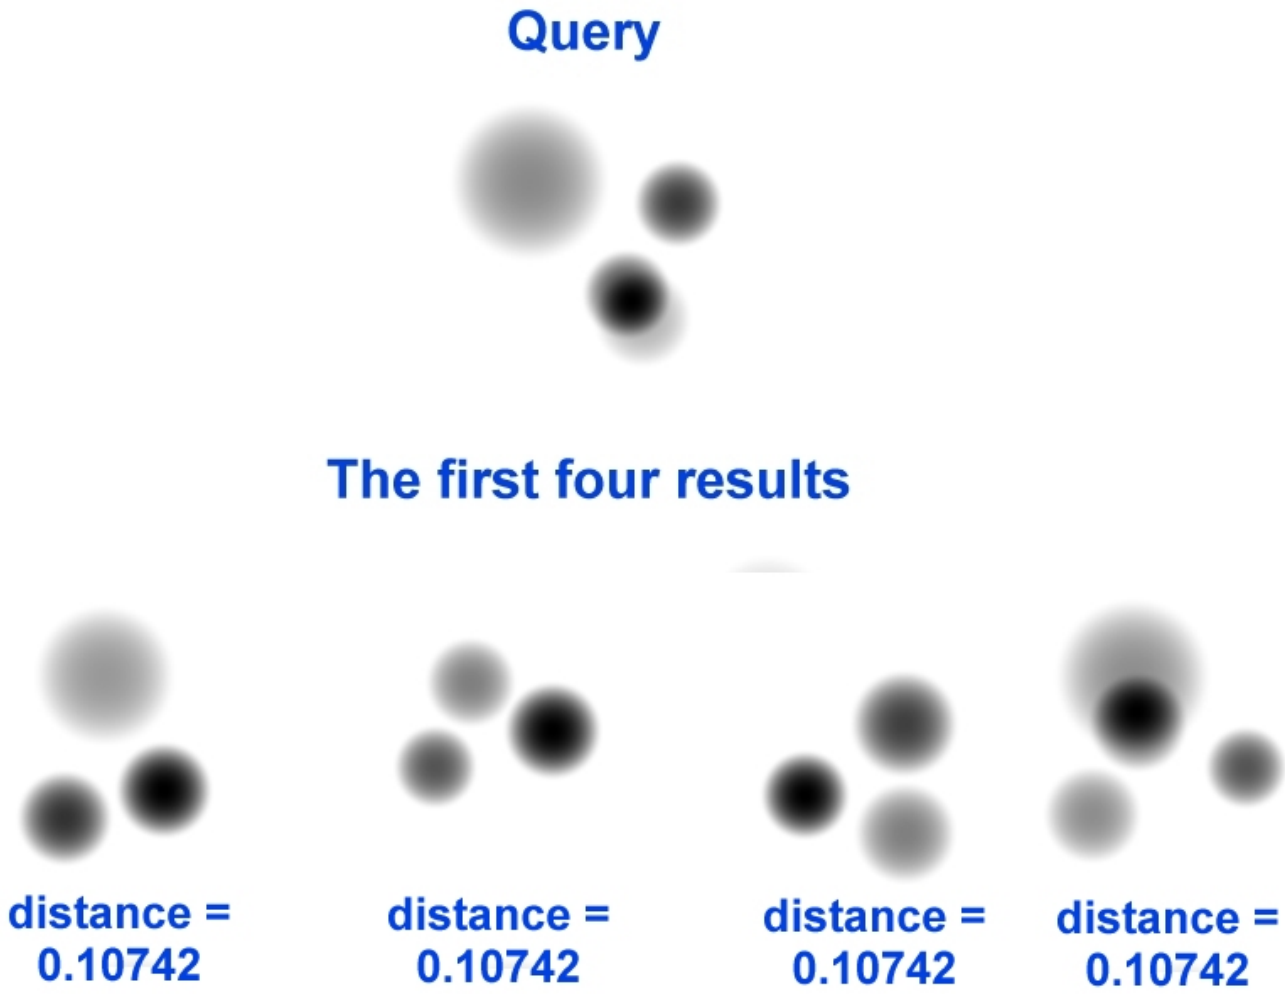

Conclusions

In this lecture 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 $\mathrm{Homeo}(\mathbf{S}^1)$. Two experiments concerning this method have been illustrated, showing the possible use of this approach for data retrieval.

References

- P. Donatini, P. Frosini, *Natural pseudodistances between closed manifolds*, Forum Mathematicum, vol. 16, n. 5, 695–715 (2004).
- P. Donatini, P. Frosini, *Natural pseudodistances between closed surfaces*, Journal of the European Mathematical Society, vol. 9, n. 2, 231–253 (2007).
- P. Frosini, *G-invariant persistent homology*, Mathematical Methods in the Applied Sciences, DOI: 10.1002/mma.3139. Preprint also available at arXiv:1212.0655.
- P. Frosini, G. Jabłoński, *Combining persistent homology and invariance groups for shape comparison*, AMS Acta, Università di Bologna, 4160 (2014), DOI: 10.6092/unibo/amsacta/4160. Preprint also available at `http://www.dm.unibo.it/~frosini/pdf/GINOs.pdf`

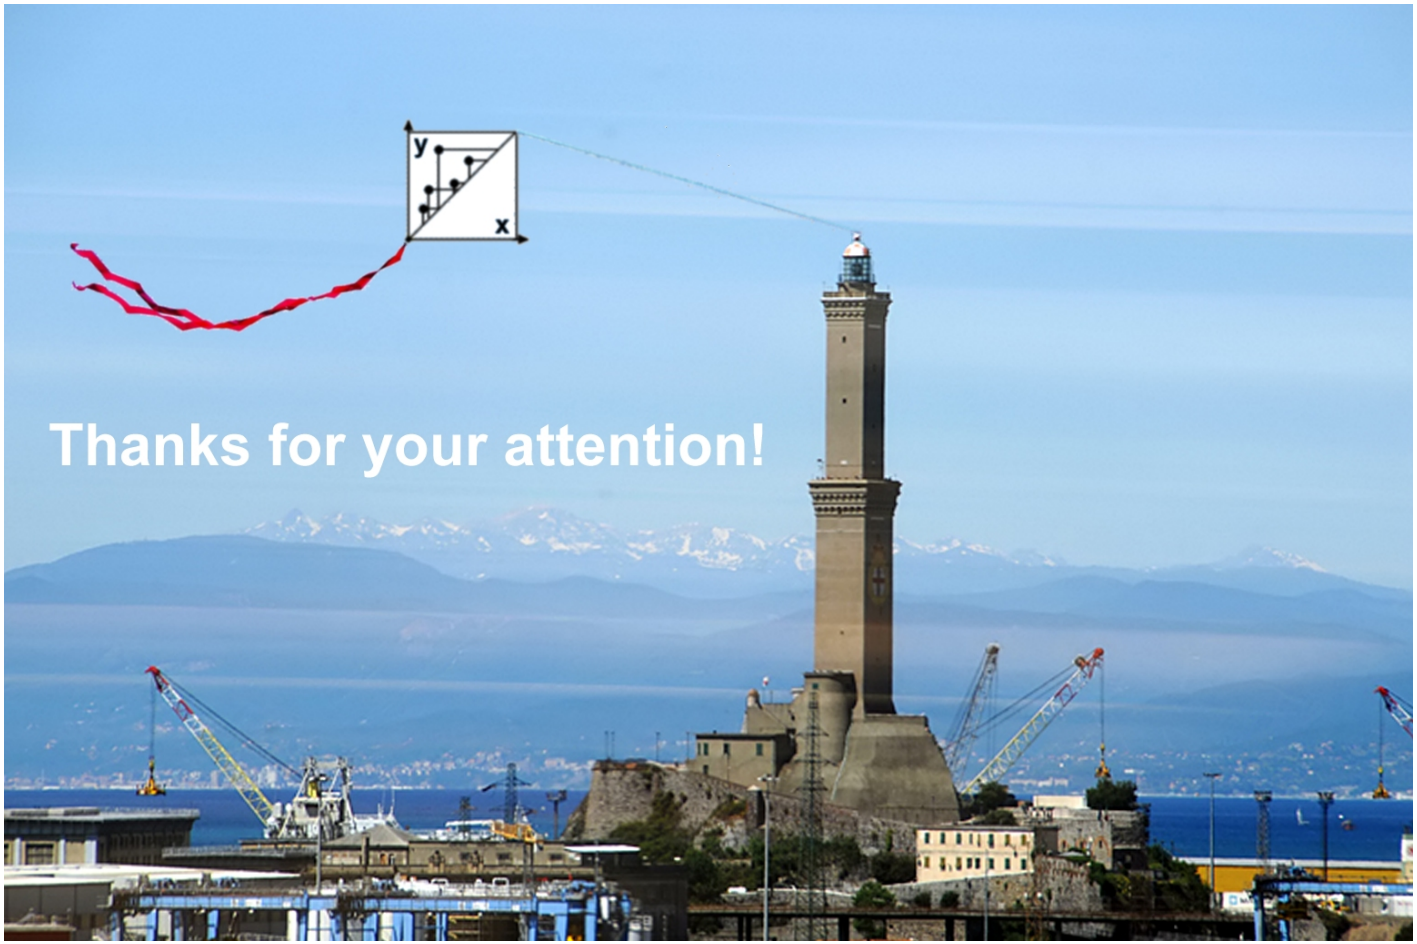
y
x
Thanks for your attention!