Combining Persistent Homology and Invariance Groups for Shape Comparison

Published Version:
Combining Persistent Homology and Invariance Groups for Shape Comparison / Frosini, Patrizio; Jabłoński, Grzegorz. - In: DISCRETE & COMPUTATIONAL GEOMETRY. - ISSN 0179-5376. - STAMPA. - 55:2(2016), pp. 373-409. [10.1007/s00454-016-9761-y]

Availability:
This version is available at: https://hdl.handle.net/11585/541704 since: 2021-03-13

Published:
DOI: http://doi.org/10.1007/s00454-016-9761-y

Patrizio Frosini, Grzegorz Jabłoński,
Combining persistent homology and invariance groups for shape comparison,
Discrete & Computational Geometry, vol. 55 (2016), n. 2, pages 373-409.

Noname manuscript No.
(will be inserted by the editor)

Combining persistent homology and invariance groups for shape comparison

(Second revision)

Patrizio Frosini · Grzegorz Jabłoński

This paper is dedicated to the memory of Marcello D'Orta and Jerry Essan Masslo.

the date of receipt and acceptance should be inserted later

Abstract Persistent homology has proven itself quite efficient in the topological and qualitative comparison of filtered topological spaces, when invariance with respect to every homeomorphism is required. However, we can make the following two observations about the use of persistent homology for application purposes. On the one hand, more restricted kinds of invariance are sometimes preferable (e.g., in shape comparison). On the other hand, in several practical situations filtering functions are not just auxiliary technical tools that can be exploited to study a given topological space, but instead the main aim of our analysis. Indeed, most of the data is usually produced by measurements, whose results are quite often functions defined on a topological space. As a simple example we can consider a 3D laser scanning of a surface, where the result of each measurement can be seen as a real-valued function defined on the manifold that describes the positions of the rangefinder measuring the distances. In fact, in many applications the dataset of interest is seen as a collection Φ of real-valued functions defined on a given topological space X, instead of a family of topological spaces. As a natural consequence, in these cases observers can be seen as collections of suitable operators on Φ.

Starting from these remarks, this paper proposes a way to combine persistent homology with the use of G-invariant non-expansive operators defined on Φ, where G is a group of self-homeomorphisms of X. Our goal is to give a method to study Φ in a way that is invariant with respect to G.

P. Frosini (corresponding author)
Department of Mathematics and ARCES
University of Bologna Piazza di Porta San Donato 5
40126, Bologna, Italy
E-mail: patrizio.frosini@unibo.it

G. Jabłoński
Institute of Computer Science
Jagiellonian University
Łojasiewicza 6, PL-30-348 Kraków, Poland
E-mail: grzegorz.jablonski@uj.edu.pl

Some theoretical results concerning our approach are proven, and two experiments are presented. An experiment illustrates the application of the proposed technique to compare 1D-signals, when the invariance is expressed by the group of affinities, the group of orientation-preserving affinities, the group of isometries, the group of translations and the identity group. Another experiment shows how our technique can be used for image comparison.

Keywords Natural pseudo-distance · filtering function · group action · persistent homology group · shape comparison

Mathematics Subject Classification (2000) Primary 55N35 · Secondary 47H09 54H15 57S10 68U05 65D18

1 Introduction

Persistent topology consists in the study of the properties of filtered topological spaces. From the very beginning, it has been applied to shape comparison [18, 25–27]. In this context, datasets are frequently represented by continuous $\mathbb{R}^m$-valued functions defined on a topological space X. As simple examples among many others, these functions can describe the coloring of a 3D object, the coordinates of the points in a planar curve, or the grey-levels in a x-ray CT image. Each continuous function $\varphi : X \to \mathbb{R}^m$ is called a *filtering function* and naturally induces a (multi)filtration on X, made by the sublevel sets of φ. Persistent topology allows to analyse the data represented by each filtering function by examining how much the topological properties of its sublevel sets "persist" when we go through the filtration. The main mathematical tool to perform this analysis is given by persistent homology [15]. This theory describes the birth and death of k-dimensional holes when we move along the considered filtration of the space X. When the filtering function takes its values in $\mathbb{R}$ we can look at it as a time, and the distance between the times of birth and death of a hole is defined to be its *persistence*. The more persistent is a hole, the more important it is for shape comparison, since holes with small persistence are usually due to noise.

An important property of classical persistent homology consists in the fact that if a self-homeomorphism $g : X \to X$ is given, then the filtering functions $\varphi, \varphi \circ g$ cannot be distinguished from each other by computing the persistent homology of the filtrations induced by φ and $\varphi \circ g$. As pointed out in [24], this is a relevant issue in the applications where the functions $\varphi, \varphi \circ g$ cannot be considered equivalent. This happens, e.g., when each filtering function $\varphi : X = \mathbb{R}^2 \to \mathbb{R}$ describes a grey-level image, since the images respectively described by φ and $\varphi \circ g$ may have completely different appearances. A simple instance of this problem is illustrated in Figure 1.

Therefore, a natural question arises: How can we adapt persistent homology in order to prevent invariance with respect to the group $\mathrm{Homeo}(X)$ of all self-homeomorphisms of the topological space X, maintaining just the invariance under the action of the self-homeomorphisms that belong to a proper subgroup

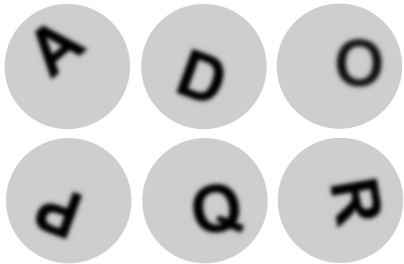

Fig. 1 Examples of letters $\mathtt{A}, \mathtt{D}, \mathtt{O}, \mathtt{P}, \mathtt{Q}, \mathtt{R}$ represented by functions $\varphi_{\mathtt{A}}, \varphi_{\mathtt{D}}, \varphi_{\mathtt{O}}, \varphi_{\mathtt{P}}, \varphi_{\mathtt{Q}}, \varphi_{\mathtt{R}}$ from $\mathbb{R}^2$ to the real numbers. Each 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). In spite of the differences between the shapes of the considered letters, the persistent homology of the functions $\varphi_{\mathtt{A}}, \varphi_{\mathtt{D}}, \varphi_{\mathtt{O}}, \varphi_{\mathtt{P}}, \varphi_{\mathtt{Q}}, \varphi_{\mathtt{R}}$ is the same in every degree.

of $\mathrm{Homeo}(X)$? For example, the comparison of the letters illustrated in Figure 1 should require just the invariance with respect to the group of similarities of $\mathbb{R}^2$, since they all are equivalent with respect to the group $\mathrm{Homeo}(\mathbb{R}^2)$. We point out that depicted letters are constructed from thick lines and therefore have some width in opposite to the concept of geometrical lines.

One could think of solving the previous problem by using other filtering functions, possibly defined on different topological spaces. For example, we could extract the boundaries of the letters in Figure 1 and consider the distance from the center of mass of each boundary as a new filtering function. 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. If the boundary is computed by taking a single level of φ, this is also in contrast with the general spirit of persistent homology.
2. It usually requires an extra computational cost (e.g., to extract the boundaries of the letters in our previous example).
3. It can produce a different topological space for each new filtering function (e.g., the letters of the alphabet can have non-homeomorphic boundaries). Working with several topological spaces instead of just one can be a disadvantage.
4. 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.

The purpose of this paper is to present a possible solution for the previously described problem. It is based on a dual approach to the invariance with respect to a subgroup G of $\mathrm{Homeo}(X)$, and consists in changing the direct study of the group G into the study of how the operators that are invariant under the action of G act on classical persistent homology. This change of perspective reveals interesting mathematical properties, allowing to treat G as a variable in our applications. According to this method, the shape properties and the invariance group can be determined separately, depending on our task. The operators that we consider in this paper act on the space of admissible filtering functions and, in some sense, can be interpreted as the "glasses" we use to look at the data. Their use allows to combine persistent homology and the invariance with respect to the group G, extending the range of application of classical persistent homology to the cases in which we are interested in G-invariance rather than in $\mathrm{Homeo}(X)$-invariance.

The idea of applying operators to filtering functions before computing persistent homology has been already considered in previous papers. For example, in [7] convolutions have been used to get a bound for the norm of persistence diagrams of a diffusing function. Furthermore, in [24] scale space persistence has been shown useful to detect critical points of a function by examining the evolution of their homological persistence values through the scale space. As for combining persistent homology and transformation groups, the interest in measuring the invariance of a signal with respect to a group of translations (i.e. the study of its periodicity or quasi-periodicity) has been studied in [9, 23], using embedding operators. However, our approach requires to consider just a particular kind of operators (i.e. G-invariant non-expansive operators on the set of admissible filtering functions), and faces the more general problem of adapting persistent homology to *any* group of self-homeomorphisms of a topological space.

For another approach to this problem, using quite a different method, we refer the reader to [17].

1.1 Our main idea in a nutshell

After choosing a set Φ of admissible filtering functions from the topological space X to $\mathbb{R}$, and a subgroup G of $\mathrm{Homeo}(X)$, we consider the set $\boldsymbol{\mathcal{F}}(\Phi, G)$ of all G-invariant non-expansive operators $F : \Phi \to \Phi$. Basically, our idea consists in comparing two functions $\varphi_1, \varphi_2 \in \Phi$ by computing the supremum of the bottleneck distances between the classical persistence diagrams of the filtering functions $F \circ \varphi_1$ and $F \circ \varphi_2$, varying F in $\boldsymbol{\mathcal{F}}(\Phi, G)$. In our paper we prove that this approach is well-defined, G-invariant, stable and computable (under suitable assumptions).

1.2 Outline of the paper

Our paper is organized as follows. In Section 2 we introduce some concepts that will be used in the paper and recall some basic facts about persistent homology. In Section 3 we prove our main results concerning the theoretical properties of our method (Theorems 15, 16 and 17). In Section 4 we illustrate the application of our technique to an experiment concerning 1D-signals. In Section 5 a possible application to image retrieval is outlined. A short discussion concludes the paper.

2 Mathematical setting

Let us consider a (non-empty) metric space X, triangulated by a finite (and hence compact) simplicial complex. We also assume that X has nontrivial homology in degree k. This last assumption is always satisfied for $k = 0$ and unrestrictive for $k \geq 1$, since we can embed X in a larger (finitely) triangulable space Y_k with nontrivial homology in degree k, and substitute X with Y_k. Let $C^0(X, \mathbb{R})$ be the set of all continuous functions from X to $\mathbb{R}$, endowed with the topology induced by the sup-norm $\| \cdot \|_\infty$. The compactness of X implies that the functions in $C^0(X, \mathbb{R})$ are uniformly continuous. Let Φ be a topological subspace of $C^0(X, \mathbb{R})$, containing at least the set of all constant functions. The functions in the topological space Φ will be called *admissible filtering functions on* X.

We assume that a subgroup G of the group $\mathrm{Homeo}(X)$ of all homeomorphisms from X onto X is given, acting on the set Φ by composition on the right. In other words, we assume that G has been chosen in such a way that if $\varphi \in \Phi$ and $g \in G$ then $\varphi \circ g \in \Phi$. At least the identity group verifies this property. The action of $g \in G$ takes each function $\varphi \in \Phi$ to the function $\varphi \circ g \in \Phi$. We do not require G to be a proper subgroup of $\mathrm{Homeo}(X)$, so the equality $G = \mathrm{Homeo}(X)$ can possibly hold. It is easy to check that G is a topological group with respect to the topology of uniform convergence, and that the right action of $G \subseteq \mathrm{Homeo}(X)$ on the set Φ is continuous.

If S is a subset of $\mathrm{Homeo}(X)$, the set $\{\varphi \circ s : \varphi \in \Phi, s \in S\}$ will be denoted by the symbol $\Phi \circ S$. Since we are assuming that the action of every $g \in G$ takes each function $\varphi \in \Phi$ to a function $\varphi \circ g \in \Phi$, we have $\Phi \circ G = \Phi$. We observe that in practice, in our experiments, we will work with a finite subset $\Phi_{ds} \subset \Phi$. In this case the inclusion $\Phi_{ds} \circ G \supseteq \Phi_{ds}$ will hold.

We can consider the natural pseudo-distance d_G on the space Φ (cf. [20, 11–13, 2]):

Definition 1 *The pseudo-distance $d_G : \Phi \times \Phi \to \mathbb{R}$ is defined by setting*

$$d_G(\varphi_1, \varphi_2) = \inf_{g \in G} \max_{x \in X} |\varphi_1(x) - \varphi_2(g(x))| .$$

It is called the (1-dimensional) natural pseudo-distance associated with the group G acting on Φ.

The term “1-dimensional” refers to the fact that the filtering functions are real-valued. The concepts considered in this paper can be easily extended to the case of $\mathbb{R}^m$-valued filtering functions, by substituting the absolute value in $\mathbb{R}$ with the max-norm $\|(u_1, \ldots, u_m)\| := \max_i |u_i|$ in $\mathbb{R}^m$. However, the use of $\mathbb{R}^m$-valued filtering functions would require the introduction of a technical machinery that is beyond the purposes of our research (cf., e.g., [5]), in order to adapt the bottleneck distance to the new setting. Therefore, for the sake of simplicity, in this paper we will just consider the 1-dimensional case.

It follows directly from the definition of the natural pseudo-distance d_G that when G is the trivial group Id, then d_G equals the sup-norm distance d_∞

on Φ, defined by setting $d_\infty(\varphi_1, \varphi_2) := \|\varphi_1 - \varphi_2\|_\infty$. Moreover, the definition of d_G immediately implies that if G_1 and G_2 are subgroups of $\mathrm{Homeo}(X)$ acting on Φ and $G_1 \subseteq G_2$, then $d_{G_2}(\varphi_1, \varphi_2) \le d_{G_1}(\varphi_1, \varphi_2)$ for every $\varphi_1, \varphi_2 \in \Phi$. As a consequence, the following double inequality holds, for every subgroup G of $\mathrm{Homeo}(X)$ and every $\varphi_1, \varphi_2 \in \Phi$ (see also Theorem 5.2 in [5]):

$$d_{\mathrm{Homeo}(X)}(\varphi_1, \varphi_2) \le d_G(\varphi_1, \varphi_2) \le d_\infty(\varphi_1, \varphi_2).$$

Remark 2 *The proof that d_G is a pseudo-metric* does *use the assumption that G is a group, and one can give examples of subsets S of* Homeo(X) *for which the function $\mu_S(\varphi_1, \varphi_2) := \inf_{s \in S} \|\varphi_1 - \varphi_2 \circ s\|_\infty$ is not a pseudo-distance on Φ.*

The rationale of using the natural pseudo-distance is that pattern recognition is usually based on comparing properties that are described by functions defined on a topological space. These properties are often the only accessible data, implying that every discrimination should be based on them. The fundamental assumption is that two objects cannot be distinguished if they share the same properties with respect to a given observer (cf. [1]).

In order to proceed, we consider the set $\mathcal{F}(\Phi, G)$ of all operators that verify the following properties:

1. F is a function from Φ to Φ;
2. $F(\varphi \circ g) = F(\varphi) \circ g$ for every $\varphi \in \Phi$ and every $g \in G$;
3. $\|F(\varphi_1) - F(\varphi_2)\|_\infty \le \|\varphi_1 - \varphi_2\|_\infty$ for every $\varphi_1, \varphi_2 \in \Phi$ (i.e. F is non-expansive).

Obviously, $\mathcal{F}(\Phi, G)$ is not empty, since it contains at least the identity operator.

Properties 1 and 2 show that F is a G-operator, referring to the right action of G on Φ.

Remark 3 *The operators that we are considering are not required to be linear. However, due to the non-expansivity property, the operators in $\mathcal{F}(\Phi, G)$ are* 1-*Lipschitz and hence are continuous.*

In this paper, we shall say that a pseudo-metric $\bar{d}$ on Φ is *strongly G-invariant* if it is invariant under the action of G with respect to each variable, i.e., if $\bar{d}(\varphi_1, \varphi_2) = \bar{d}(\varphi_1 \circ g, \varphi_2) = \bar{d}(\varphi_1, \varphi_2 \circ g) = \bar{d}(\varphi_1 \circ g, \varphi_2 \circ g)$ for every $\varphi_1, \varphi_2 \in \Phi$ and every $g \in G$.

Remark 4 *It is easily seen that the natural pseudo-distance d_G is strongly G-invariant.*

Example 5 *Take $X = S^1$, G equal to the group $R(S^1)$ of all rotations of S^1, and Φ equal to the set $C^0(S^1, \mathbb{R})$ of all continuous functions from S^1 to $\mathbb{R}$. As an example of an operator in $\mathcal{F}(\Phi, G)$ we can consider the operator F_α defined by setting $F_\alpha(\varphi)(x) := \frac{1}{2} \cdot (\varphi(x) + \varphi(x_\alpha))$ for every $\varphi \in C^0(S^1, \mathbb{R})$ and every $x \in S^1$, where x_α denotes the point obtained from x by rotating S^1 of a* *fixed angle* α. *It is easy to check that* F_α *is an* $R(S^1)$*-invariant non-expansive (linear) operator defined on* $C^0(S^1,\mathbb{R})$. *An example of an* $R(S^1)$*-invariant non-expansive non-linear operator defined on* $C^0(S^1,\mathbb{R})$ *is given by the operator* $\bar{F}$ *defined by setting* $\bar{F}(\varphi)(x) = \varphi(x) + 1$ *for every* $\varphi \in C^0(S^1,\mathbb{R})$ *and every* $x \in S^1$.

This simple statement holds (the symbol $\mathbf{0}$ denotes the function taking the value 0 everywhere):

Proposition 6 $\|F(\varphi)\|_\infty \leq \|\varphi\|_\infty + \|F(\mathbf{0})\|_\infty$ *for every* $F \in \boldsymbol{\mathcal{F}}(\Phi, G)$ *and every* $\varphi \in \Phi$.

Proof $\|F(\varphi)\|_\infty = \|F(\varphi) - F(\mathbf{0}) + F(\mathbf{0})\|_\infty \leq \|F(\varphi) - F(\mathbf{0})\|_\infty + \|F(\mathbf{0})\|_\infty \leq \|\varphi - \mathbf{0}\|_\infty + \|F(\mathbf{0})\|_\infty = \|\varphi\|_\infty + \|F(\mathbf{0})\|_\infty$, since F is non-expansive. □

If $\mathcal{F} \neq \emptyset$ is a subset of $\boldsymbol{\mathcal{F}}(\Phi, G)$ and Φ is bounded with respect to d_∞, then we can consider the function

$$d_{\mathcal{F}}(F_1, F_2) := \sup_{\varphi \in \Phi} \|F_1(\varphi) - F_2(\varphi)\|_\infty$$

from $\mathcal{F} \times \mathcal{F}$ to $\mathbb{R}$.

Proposition 7 *If* $\mathcal{F}$ *is a non-empty subset of* $\boldsymbol{\mathcal{F}}(\Phi, G)$ *and* Φ *is bounded then the function* $d_{\mathcal{F}}$ *is a distance on* $\mathcal{F}$.

Proof See Appendix A. □

Remark 8 *The* sup *in the definition of* $d_{\mathcal{F}}$ *cannot be replaced with* max. *As an example, consider the case* $X = [0,1]$, $\Phi = C^0([0,1],[0,1])$, G *equal to the group containing just the identity and the homeomorphism taking each point* $x \in [0,1]$ *to* $1-x$, $F_1(\varphi)$ *equal to the constant function taking everywhere the value* $\max \varphi$, *and* $F_2(\varphi)$ *equal to the constant function taking everywhere the value* $\int_0^1 \varphi(x)\ dx$. *Both* F_1 *and* F_2 *are non-expansive* G*-operators. We have that* $d_{\mathcal{F}}(F_1, F_2) = 1$, *but no function* $\psi \in \Phi = C^0([0,1],[0,1])$ *exists, such that* $\|F_1(\psi) - F_2(\psi)\|_\infty = 1$.

2.1 Persistent homology

Before proceeding, we recall some basic definitions and facts in persistent homology. For a more detailed and formal treatment, we refer the interested reader to [15,1,3,6]. Roughly speaking, persistent homology describes the changes of the homology groups of the sub-level sets $X_t = \varphi^{-1}((-\infty, t])$ varying t in $\mathbb{R}$, where φ is a real-valued continuous function defined on a topological space X. The parameter t can be seen as an increasing time, whose change produces the birth and death of k-dimensional holes in the sub-level set X_t. For $k = 0, 1, 2$, the expression "k-dimensional holes" refers to connected components, tunnels and voids, respectively.

Persistent homology can be introduced in several different settings, including the one of simplicial complexes and simplicial homology, and the one of topological spaces and singular homology. As for the link between the discrete and the topological settings, we refer the interested reader to [4,10]. In this paper we will consider the topological setting and the singular homology functor H. An elementary introduction to singular homology can be found in [21].

The concept of persistence can be formalized by the definition of persistent homology group with respect to the function $\varphi : X \to \mathbb{R}$:

Definition 9 *If $u, v \in \mathbb{R}$ and $u < v$, we can consider the inclusion i of X_u into X_v. Such an inclusion induces a homomorphism $i_k : H_k(X_u) \to H_k(X_v)$ between the homology groups of X_u and X_v in degree k. The group $PH_k^{\varphi}(u,v) := i_k(H_k(X_u))$ is called the* k-th persistent homology group with respect to the function $\varphi : X \to \mathbb{R}$, computed at the point (u,v)*. The rank $r_k(\varphi)(u,v)$ of this group is said* the k-th persistent Betti number function with respect to the function $\varphi : X \to \mathbb{R}$, computed at the point (u,v).

Remark 10 *It is easy to check that the persistent homology groups (and hence also the persistent Betti number functions) are invariant under the action of* $\mathrm{Homeo}(X)$*. For further discussion see Appendix B.*

A classical way to describe persistent Betti number functions (up to subsets of measure zero of their domain) is given by multisets named *persistence diagrams*. Another equivalent description is given by *barcodes* (cf. [3]). The k-th persistence diagram is the multiset of all pairs $p_j = (b_j, d_j)$, where b_j and d_j are the times of birth and death of the j-th k-dimensional hole, respectively. When a hole never dies, we set its time of death equal to ∞. The multiplicity $m(p_j)$ says how many holes share both the time of birth b_j and the time of death d_j. For technical reasons, the points (t,t) are added to each persistence diagram, each one with infinite multiplicity.

Persistence diagrams can be compared by a metric δ_{match}, called *bottleneck distance* or *matching distance*. We recall here its formal definition, taking into account that each persistence diagram D can contain an infinite number of points, and that each point $p \in D$ has a multiplicity $m(p) \geq 1$. For every $q \in \Delta^*$, the equality $m(q) = 0$ means that q does not belong to the persistence diagram D. In our exposition we will set $\Delta := \{(x,y) \in \mathbb{R}^2 : x = y\}$, $\Delta^+ := \{(x,y) \in \mathbb{R}^2 : x < y\}$, $\bar{\Delta}^+ := \{(x,y) \in \mathbb{R}^2 : x \leq y\}$, $\Delta^* := \Delta^+ \cup \{(x,\infty) : x \in \mathbb{R}\}$ and $\bar{\Delta}^* := \bar{\Delta}^+ \cup \{(x,\infty) : x \in \mathbb{R}\}$.

We start by endowing the set $\bar{\Delta}^*$ with the pseudo-metric

$$d^*((x,y),(x',y')) := \min\left\{\max\{|x-x'|,|y-y'|\}, \max\left\{\frac{y-x}{2}, \frac{y'-x'}{2}\right\}\right\}$$

by agreeing that $\infty - y = \infty$, $y - \infty = -\infty$ for $y \neq \infty$, $\infty - \infty = 0$, $\frac{\infty}{2} = \infty$, $|\pm\infty| = \infty$, $\min\{\infty, c\} = c$, $\max\{\infty, c\} = \infty$.

The pseudo-metric d^* between two points $p = (x, y)$ and $p' = (x', y')$ compares the cost of moving p to p' and the cost of moving p and p' onto Δ and takes the smaller. We observe that $d^*(p, p') = 0$ for every $p, p' \in \Delta$. If $p \in \Delta^+$ and $p \in \Delta$, then $d^*(p, p')$ equals the distance, induced by the max-norm, between p and Δ. Points at infinity have a finite distance only to other points at infinity, and their distance equals the Euclidean distance between their abscissas.

Definition 11 *Let D_1, D_2 be two persistence diagrams. We define the bottleneck distance δ_{match} between D_1 and D_2 by setting*

$$\delta_{match}(D_1, D_2) := \inf_{\sigma} \sup_{x \in D_1} d^*(x, \sigma(x))$$

where $\sigma : D_1 \to D_2$ is a bijection.

For further details about the concepts of persistence diagram and bottleneck distance, we refer the reader to [8] and [15]. Algorithms to compute the bottleneck distance δ_{match} can be found in [16].

Since each persistent Betti number function is associated with exactly one persistence diagram, it follows that the metric δ_{match} naturally induces a pseudo-metric d_{match} on the sets of the persistent Betti number functions. We recall that a pseudo-metric is just a metric without the property assuring that if two points have a null distance then they must coincide. For more details about the existence of pairs of persistent Betti number functions that differ from each other at the points of a set of measure zero but are associated with the same persistence diagram, we refer the interested reader to [5].

A key property of the distance d_{match} is its stability with respect to d_∞ and $d_{\mathrm{Homeo}(X)}$, stated in the following result.

Theorem 12 *If k is a natural number and $\varphi_1, \varphi_2 \in \Phi = C^0(X, \mathbb{R})$, then*

$$d_{match}(r_k(\varphi_1), r_k(\varphi_2)) \leq d_{\mathrm{Homeo}(X)}(\varphi_1, \varphi_2) \leq d_\infty(\varphi_1, \varphi_2).$$

The proof of the inequality $d_{match}(r_k(\varphi_1), r_k(\varphi_2)) \leq d_\infty(\varphi_1, \varphi_2)$ in Theorem 12 can be found in [8] (Main Theorem) for the case of tame filtering functions and in [5] (Theorem 3.13) for the general case of continuous functions. The first inequality in Theorem 12 follows from Theorem 5.2 in [5] (which is in turn a corollary of the Multidimensional Stability Theorem 4.4 in [5], following from one-dimensional stability). The other inequality is a trivial consequence of the definition of $d_{\mathrm{Homeo}(X)}$. Theorem 12 also shows that the natural pseudo-distance d_G allows to obtain a stability result for persistence diagrams that is better than the classical one, involving d_∞. Figure 2 illustrates this fact, displaying two filtering functions $\varphi_1, \varphi_2 : [0, 1] \to \mathbb{R}$ such that $d_{match}(r_k(\varphi_1), r_k(\varphi_2)) = d_{\mathrm{Homeo}(X)}(\varphi_1, \varphi_2) = 0 < \|\varphi_1 - \varphi_2\|_\infty = 1$. In other words, d_G is closer than d_∞ to d_{match} and examples exist where the knowledge of d_G allows us to infer that two given persistence diagrams are close to each other, while the knowledge of d_∞ cannot give the same information.

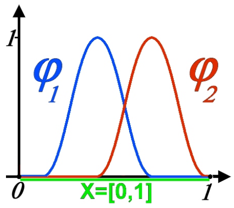

Fig. 2 These two functions have the same persistent homology ($d_{match}(r_k(\varphi_1), r_k(\varphi_2)) = 0$), but $\|\varphi_1 - \varphi_2\|_\infty = 1$. However, they are equivalent with respect to the group $G = \mathrm{Homeo}([0,1])$, hence $d_{\mathrm{Homeo}(X)}(\varphi_1, \varphi_2) = 0$. As a consequence, $d_{\mathrm{Homeo}(X)}(\varphi_1, \varphi_2)$ gives an upper bound of $d_{match}(r_k(\varphi_1), r_k(\varphi_2))$ that is better than the one given by the sup-norm $\|\varphi_1 - \varphi_2\|_\infty$, via the classical *Bottleneck Stability Theorem* for persistence diagrams (cf. [8]).

2.2 Strongly G-invariant comparison of filtering functions via persistent homology

Let us fix a non-empty subset $\mathcal{F}$ of $\boldsymbol{\mathcal{F}}(\Phi, G)$. For every fixed k, we can consider the following pseudo-metric $D_{match}^{\mathcal{F},k}$ on Φ:

$$D_{match}^{\mathcal{F},k}(\varphi_1, \varphi_2) := \sup_{F \in \mathcal{F}} d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2)))$$

for every $\varphi_1, \varphi_2 \in \Phi$, where $r_k(\varphi)$ denotes the k-th persistent Betti number function with respect to the function $\varphi : X \to \mathbb{R}$. We will usually omit the index k, when its value is clear from the context or not influential.

Proposition 13 *$D_{match}^{\mathcal{F}}$ is a strongly G-invariant pseudo-metric on Φ.*

Proof Theorem 12 and the non-expansivity of every $F \in \mathcal{F}$ imply that

$$\begin{aligned}
& d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2))) \leq \\
& \|F(\varphi_1) - F(\varphi_2)\|_\infty \leq \\
& \|\varphi_1 - \varphi_2\|_\infty .
\end{aligned}$$

Therefore $D_{match}^{\mathcal{F}}$ is a pseudo-metric, since it is the supremum of a family of pseudo-metrics that are bounded at each pair (φ_1, φ_2). Moreover, for every $\varphi_1, \varphi_2 \in \Phi$ and every $g \in G$

$$\begin{aligned}
& D_{match}^{\mathcal{F}}(\varphi_1, \varphi_2 \circ g) := \\
& \sup_{F \in \mathcal{F}} d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2 \circ g))) = \\
& \sup_{F \in \mathcal{F}} d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2) \circ g)) = \\
& \sup_{F \in \mathcal{F}} d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2))) = \\
& D_{match}^{\mathcal{F}}(\varphi_1, \varphi_2)
\end{aligned}$$

because of Property 2 in the definition of $\boldsymbol{\mathcal{F}}(\Phi, G)$ and the invariance of persistent homology under the action of homeomorphisms (Remark 10). Due to the fact that the function $D^{\mathcal{F}}_{match}$ is symmetric, this is sufficient to guarantee that $D^{\mathcal{F}}_{match}$ is strongly G-invariant. $\square$

2.3 Approximating $D^{\mathcal{F}}_{match}$

A method to approximate $D^{\mathcal{F}}_{match}$ is given by the next proposition.

Proposition 14 *Assume Φ bounded. Let $\mathcal{F}^* = \{F_1, \ldots, F_m\}$ be a finite subset of $\mathcal{F}$. If for every $F \in \mathcal{F}$ at least one index $i \in \{1, \ldots, m\}$ exists, such that $d_{\mathcal{F}}(F_i, F) \le \epsilon$, then*

$$\left| D^{\mathcal{F}^*}_{match}(\varphi_1, \varphi_2) - D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2) \right| \le 2\epsilon$$

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

Proof Let us assume $F \in \mathcal{F}$ and $d_{\mathcal{F}}(F_i, F) \le \epsilon$. Because of the definition of $d_{\mathcal{F}}$, for any $\varphi_1, \varphi_2 \in \Phi$ we have the inequalities $\|F_i(\varphi_1) - F(\varphi_1)\|_\infty \le \epsilon$ and $\|F_i(\varphi_2) - F(\varphi_2)\|_\infty \le \epsilon$. Hence

$$d_{match}(r_k(F_i(\varphi_1)), r_k(F(\varphi_1))) \le \epsilon \text{ and } d_{match}(r_k(F_i(\varphi_2)), r_k(F(\varphi_2))) \le \epsilon$$

because of the stability of persistent homology (Theorem 12). It follows that

$$|d_{match}(r_k(F_i(\varphi_1)), r_k(F_i(\varphi_2))) - d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2)))| \le 2\epsilon.$$

The thesis of our proposition immediately follows from the definitions of $D^{\mathcal{F}}_{match}$ and $D^{\mathcal{F}^*}_{match}$. $\square$

Therefore, if we can cover $\mathcal{F}$ by a finite set of balls of radius ϵ, centered at points of $\mathcal{F}$, the approximation of $D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2)$ can be reduced to the computation of the maximum of a finite set of bottleneck distances between persistence diagrams, which are well-known to be computable by means of efficient algorithms.

This fact leads us to study the properties of the topological space $\boldsymbol{\mathcal{F}}(\Phi, G)$. We will do that in the next section.

3 Main theoretical results

We start by proving that the pseudo-metric $D^{\mathcal{F}}_{match}$ is stable with respect to both the natural pseudo-distance associated with the group G and the sup-norm.

Theorem 15 *If $\emptyset \ne \mathcal{F} \subseteq \boldsymbol{\mathcal{F}}(\Phi, G)$, then $D^{\mathcal{F}}_{match} \le d_G \le d_\infty$.*

Proof For every $F \in \boldsymbol{\mathcal{F}}(\Phi, G)$, every $g \in G$ and every $\varphi_1, \varphi_2 \in \Phi$, we have that

$$\begin{aligned}
&d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2))) = \\
&d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2) \circ g)) = \\
&d_{match}(r_k(F(\varphi_1)), r_k(F(\varphi_2 \circ g))) \leq \\
&\|F(\varphi_1) - F(\varphi_2 \circ g)\|_\infty \leq \|\varphi_1 - \varphi_2 \circ g\|_\infty.
\end{aligned}$$

The first equality follows from the invariance of persistent homology under the action of Homeo(X) (Remark 10), and the second equality follows from the fact that F is a G-operator. The first inequality follows from the stability of persistent homology (Theorem 12), while the second inequality follows from the non-expansivity of F.

It follows that, if $\mathcal{F} \subseteq \boldsymbol{\mathcal{F}}(\Phi, G)$, then for every $g \in G$ and every $\varphi_1, \varphi_2 \in \Phi$

$$D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2) \leq \|\varphi_1 - \varphi_2 \circ g\|_\infty.$$

Hence,

$$\begin{aligned}
&D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2) \leq \inf_{g \in G} \|\varphi_1 - \varphi_2 \circ g\|_\infty \leq \\
&\|\varphi_1 - \varphi_2\|_\infty = d_\infty(\varphi_1, \varphi_2)
\end{aligned}$$

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

The natural pseudo-distance d_G and the pseudo-distance $D^{\mathcal{F}}_{match}$ are defined in completely different ways. The former is based on a variational approach involving the set of all homeomorphisms in G, while the latter refers only to a comparison of persistent homologies depending on a family of G-invariant operators. Therefore, the next result may appear unexpected.

Theorem 16 $D^{\boldsymbol{\mathcal{F}}(\Phi,G)}_{match} = d_G$.

Proof For every $\psi \in \Phi$ let us consider the operator $F_\psi : \Phi \to \Phi$ defined by setting $F_\psi(\varphi)$ equal to the constant function taking everywhere the value $d_G(\varphi, \psi)$, for every $\varphi \in \Phi$ (i.e., $F_\psi(\varphi)(x) = d_G(\varphi, \psi)$ for any $x \in X$).

We observe that

i) F_ψ is a G-operator on Φ, because the strong invariance of the natural pseudo-distance d_G with respect to the group G (Remark 4) implies that if $\varphi \in \Phi$ and $g \in G$, then $F_\psi(\varphi \circ g)(x) = d_G(\varphi \circ g, \psi) = d_G(\varphi, \psi) = F_\psi(\varphi)(g(x)) = (F_\psi(\varphi) \circ g)(x)$, for every $x \in X$.

ii) F_ψ is non-expansive, because

$$\begin{aligned}
&\|F_\psi(\varphi_1) - F_\psi(\varphi_2)\|_\infty = |d_G(\varphi_1, \psi) - d_G(\varphi_2, \psi)| \leq \\
&d_G(\varphi_1, \varphi_2) \leq \|\varphi_1 - \varphi_2\|_\infty.
\end{aligned}$$

Therefore, $F_\psi \in \boldsymbol{\mathcal{F}}(\Phi, G)$.

For every $\varphi_1, \varphi_2, \psi \in \Phi$ we have that

$$d_{match}(r_k(F_\psi(\varphi_1)), r_k(F_\psi(\varphi_2))) = |d_G(\varphi_1, \psi) - d_G(\varphi_2, \psi)|.$$

Indeed, apart from the trivial points on the line $\{(u, v) \in \mathbb{R}^2 : u = v\}$, the persistence diagram associated with $r_k(F_\psi(\varphi_1))$ contains only the point $(d_G(\varphi_1, \psi), \infty)$, while the persistence diagram associated with $r_k(F_\psi(\varphi_2))$ contains only the point $(d_G(\varphi_2, \psi), \infty)$. Both the points have the same multiplicity, which equals the (non-null) k-th Betti number of X.

Setting $\psi = \varphi_2$, we have that

$$d_{match}(r_k(F_{\varphi_2}(\varphi_1)), r_k(F_{\varphi_2}(\varphi_2))) = d_G(\varphi_1, \varphi_2).$$

As a consequence, we have that

$$D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2) \geq d_G(\varphi_1, \varphi_2).$$

By applying Theorem 15, we get $D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2) = d_G(\varphi_1, \varphi_2)$ for every $\varphi_1, \varphi_2 \in \Phi$. $\square$

The following two results (Theorem 17 and Corollary 19) hold, when the metric space (Φ, d_∞) is compact.

Theorem 17 *If the metric space (Φ, d_∞) is compact, then also the metric space $\left(\boldsymbol{\mathcal{F}}(\Phi, G), d_{\boldsymbol{\mathcal{F}}(\Phi,G)}\right)$ is compact.*

Proof Since Φ is bounded, Proposition 7 guarantees that the function $d_{\boldsymbol{\mathcal{F}}(\Phi,G)}$ is defined and $\left(\boldsymbol{\mathcal{F}}(\Phi, G), d_{\boldsymbol{\mathcal{F}}(\Phi,G)}\right)$ is a metric space. Therefore it will suffice to prove that $\boldsymbol{\mathcal{F}}(\Phi, G)$ is sequentially compact. In order to do this, let us assume that a sequence (F_i) in $\boldsymbol{\mathcal{F}}(\Phi, G)$ is given.

Given that (Φ, d_∞) is a compact (and hence separable) metric space, we can find a countable and dense subset $\Phi^* = \{\varphi_j\}_{j \in \mathbb{N}}$ of Φ. We can extract a subsequence (F_{i_h}) from (F_i), such that for every fixed index j the sequence $(F_{i_h}(\varphi_j))$ converges to a function in Φ with respect to the sup-norm. (This follows by recalling that $F_i : \Phi \to \Phi$ for every index i, with (Φ, d_∞) compact, and by applying a classical diagonalization argument.)

Now, let us consider the operator $\bar{F} : \Phi \to \Phi$ defined in the following way.

We define $\bar{F}$ on Φ^* by setting $\bar{F}(\varphi_j) := \lim_{h \to \infty} (F_{i_h}(\varphi_j))$ for each $\varphi_j \in \Phi^*$.

Then we extend $\bar{F}$ to Φ as follows. For each $\varphi \in \Phi$ we choose a sequence (φ_{j_r}) in Φ^*, converging to φ in Φ, and set $\bar{F}(\varphi) := \lim_{r \to \infty} \bar{F}(\varphi_{j_r})$. We claim that such a limit exists in Φ and does not depend on the sequence that we have chosen, converging to φ in Φ. In order to prove that the previous limit exists, we observe that for every $r, s \in \mathbb{N}$

$$\begin{aligned}
&\left\|\bar{F}\left(\varphi_{j_r}\right) - \bar{F}\left(\varphi_{j_s}\right)\right\|_\infty = \\
&\left\|\lim_{h\to\infty}\left(F_{i_h}\left(\varphi_{j_r}\right)\right) - \lim_{h\to\infty}\left(F_{i_h}\left(\varphi_{j_s}\right)\right)\right\|_\infty = \\
&\lim_{h\to\infty}\left\|F_{i_h}\left(\varphi_{j_r}\right) - F_{i_h}\left(\varphi_{j_s}\right)\right\|_\infty \le \\
&\lim_{h\to\infty}\left\|\varphi_{j_r} - \varphi_{j_s}\right\|_\infty = \\
&\left\|\varphi_{j_r} - \varphi_{j_s}\right\|_\infty
\end{aligned} \tag{3.1}$$

because each operator F_{i_h} is non-expansive.

Since the sequence (φ_{j_r}) converges to φ in Φ, it follows that $\left(\bar{F}\left(\varphi_{j_r}\right)\right)$ is a Cauchy sequence. The compactness of Φ implies that $\left(\bar{F}\left(\varphi_{j_r}\right)\right)$ converges in Φ.

If another sequence (φ_{k_r}) is given in Φ^*, converging to φ in Φ, then for every index r

$$\left\|\bar{F}\left(\varphi_{j_r}\right) - \bar{F}\left(\varphi_{k_r}\right)\right\|_\infty \le \left\|\varphi_{j_r} - \varphi_{k_r}\right\|_\infty$$

and the proof goes as in (3.1) with φ_{j_s} replaced by φ_{k_r}.

Since both (φ_{j_r}) and (φ_{k_r}) converge to φ, it follows that $\lim_{r\to\infty}\bar{F}\left(\varphi_{j_r}\right) = \lim_{r\to\infty}\bar{F}\left(\varphi_{k_r}\right)$. Therefore the definition of $\bar{F}(\varphi)$ does not depend on the sequence (φ_{j_r}) that we have chosen, converging to φ.

Now we have to prove that $\bar{F} \in \boldsymbol{\mathcal{F}}(\Phi, G)$, i.e., that $\bar{F}$ verifies the three properties defining this set of operators.

We have already seen that $\bar{F} : \Phi \to \Phi$.

For every $\varphi, \varphi' \in \Phi$ we can consider two sequences $(\varphi_{j_r}), (\varphi_{k_r})$ in Φ^*, converging to φ and φ' in Φ, respectively. Due to the fact that the operators F_{i_h} are non-expansive, we have that

$$\begin{aligned}
&\left\|\bar{F}\left(\varphi\right) - \bar{F}\left(\varphi'\right)\right\|_\infty = \\
&\left\|\lim_{r\to\infty}\bar{F}\left(\varphi_{j_r}\right) - \lim_{r\to\infty}\bar{F}\left(\varphi_{k_r}\right)\right\|_\infty = \\
&\left\|\lim_{r\to\infty}\lim_{h\to\infty}\left(F_{i_h}\left(\varphi_{j_r}\right)\right) - \lim_{r\to\infty}\lim_{h\to\infty}\left(F_{i_h}\left(\varphi_{k_r}\right)\right)\right\|_\infty = \\
&\lim_{r\to\infty}\lim_{h\to\infty}\left\|F_{i_h}\left(\varphi_{j_r}\right) - F_{i_h}\left(\varphi_{k_r}\right)\right\|_\infty \le \\
&\lim_{r\to\infty}\lim_{h\to\infty}\left\|\varphi_{j_r} - \varphi_{j_r}\right\|_\infty = \\
&\lim_{r\to\infty}\left\|\varphi_{j_r} - \varphi_{j_r}\right\|_\infty = \\
&\left\|\varphi - \varphi'\right\|_\infty .
\end{aligned}$$

Therefore, the operator $\bar{F}$ is non-expansive. As a consequence, it is also continuous.

Now we can prove that the sequence (F_{i_h}) converges to $\bar{F}$ with respect to $d_{\boldsymbol{\mathcal{F}}(\Phi,G)}$. Let us consider an arbitrarily small $\epsilon > 0$. Since Φ is compact and

Φ^* is dense in Φ, we can find a finite subset $\{\varphi_{j_1},\ldots,\varphi_{j_n}\}$ of Φ^* such that for each $\varphi\in\Phi$ an index $r\in\{1,\ldots,n\}$ exists, such that $\|\varphi-\varphi_{j_r}\|_\infty\le\epsilon$. Since the sequence (F_{i_h}) converges pointwise to $\bar F$ on the set Φ^*, an index $\bar h$ exists, such that $\|\bar F(\varphi_{j_r})-F_{i_h}(\varphi_{j_r})\|_\infty\le\epsilon$ for any $h\ge\bar h$ and any $r\in\{1,\ldots,n\}$.

Therefore, for every $\varphi\in\Phi$ we can find an index $r\in\{1,\ldots,n\}$ such that $\|\varphi-\varphi_{j_r}\|_\infty\le\epsilon$ and the following inequalities hold for every index $h\ge\bar h$, because of the non-expansivity of $\bar F$ and F_{i_h}:

$$
\begin{aligned}
&\left\|\bar F(\varphi)-F_{i_h}(\varphi)\right\|_\infty\le\\
&\left\|\bar F(\varphi)-\bar F(\varphi_{j_r})\right\|_\infty+\left\|\bar F(\varphi_{j_r})-F_{i_h}(\varphi_{j_r})\right\|_\infty+\|F_{i_h}(\varphi_{j_r})-F_{i_h}(\varphi)\|_\infty\le\\
&\|\varphi-\varphi_{j_r}\|_\infty+\left\|\bar F(\varphi_{j_r})-F_{i_h}(\varphi_{j_r})\right\|_\infty+\|\varphi_{j_r}-\varphi\|_\infty\le 3\epsilon.
\end{aligned}
$$

We observe that $\bar h$ does not depend on φ, but only on ϵ and the set $\{\varphi_{j_1},\ldots,\varphi_{j_n}\}$.

It follows that $\left\|\bar F(\varphi)-F_{i_h}(\varphi)\right\|_\infty\le 3\epsilon$ for every $\varphi\in\Phi$ and every $h\ge\bar h$.

Hence, $\sup_{\varphi\in\Phi}\left\|\bar F(\varphi)-F_{i_h}(\varphi)\right\|_\infty\le 3\epsilon$ for every $h\ge\bar h$. Therefore, the sequence (F_{i_h}) converges to $\bar F$ with respect to $d_{\boldsymbol{\mathcal F}(\Phi,G)}$.

The last thing that we have to prove is that $\bar F$ is a G-operator. Let us consider a $\varphi\in\Phi$, a sequence (φ_{j_r}) in Φ^* converging to φ in Φ, and a $g\in G$. Obviously, the sequence $(\varphi_{j_r}\circ g)$ converges to $\varphi\circ g$ in Φ. We recall that the right action of G on Φ is continuous, $\bar F$ is continuous and each F_{i_h} is a G-operator. Hence, given that the sequence (F_{i_h}) converges to $\bar F$ with respect to $d_{\boldsymbol{\mathcal F}(\Phi,G)}$,

$$
\begin{aligned}
&\bar F(\varphi\circ g)=\\
&\bar F\left(\lim_{r\to\infty}\varphi_{j_r}\circ g\right)=\\
&\lim_{r\to\infty}\bar F(\varphi_{j_r}\circ g)=\\
&\lim_{r\to\infty}\lim_{h\to\infty}\left(F_{i_h}(\varphi_{j_r}\circ g)\right)=\\
&\lim_{r\to\infty}\lim_{h\to\infty}\left(F_{i_h}(\varphi_{j_r})\circ g\right)=\\
&\lim_{r\to\infty}\left(\left(\lim_{h\to\infty}\left(F_{i_h}(\varphi_{j_r})\right)\right)\circ g\right)=\\
&\left(\lim_{r\to\infty}\lim_{h\to\infty}\left(F_{i_h}(\varphi_{j_r})\right)\right)\circ g=\\
&\left(\lim_{r\to\infty}\bar F(\varphi_{j_r})\right)\circ g=\\
&\bar F(\varphi)\circ g.
\end{aligned}
$$

This proves that $\bar F$ is a G-operator.

In conclusion, $\bar F\in\boldsymbol{\mathcal F}(\Phi,G)$.

From the fact that the sequence (F_{i_h}) converges to $\bar F$ with respect to $d_{\boldsymbol{\mathcal F}(\Phi,G)}$, it follows that $\left(\boldsymbol{\mathcal F}(\Phi,G),d_{\boldsymbol{\mathcal F}(\Phi,G)}\right)$ is sequentially compact.

Example 18 *As a simple example of a case where the previous Theorem 17 can be applied, we can consider $X = S^1 \subset \mathbb{R}^2$, Φ equal to the set of all 1-Lipschitz functions from S^1 to $[0,1]$, and G equal to the topological group of all isometries of S^1. The topological space Φ can be easily shown to be compact by applying the Ascoli-Arzelà Theorem.*

Corollary 19 *Assume that the metric space (Φ, d_∞) is compact. Let $\mathcal{F}$ be a non-empty subset of $\boldsymbol{\mathcal{F}}(\Phi, G)$. For every $\epsilon > 0$, a finite subset $\mathcal{F}^*$ of $\mathcal{F}$ exists, such that*

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

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

Proof Let us consider the closure $\bar{\mathcal{F}}$ of $\mathcal{F}$ in $\boldsymbol{\mathcal{F}}(\Phi, G)$. Let us also consider the covering $\mathcal{U}$ of $\bar{\mathcal{F}}$ obtained by taking all the open $\frac{\epsilon}{2}$-balls centered at points of $\mathcal{F}$. Theorem 17 guarantees that $\boldsymbol{\mathcal{F}}(\Phi, G)$ is compact, hence also $\bar{\mathcal{F}}$ is compact. Therefore we can extract a finite covering $\{B_1, \ldots, B_m\}$ of $\bar{\mathcal{F}}$ from $\mathcal{U}$. We can set $\mathcal{F}^*$ equal to the set of centers of the balls $B_1, \ldots, B_m$. The statement of our corollary immediately follows from Proposition 14.

The previous Corollary 19 shows that, under suitable hypotheses, the computation of $D^{\mathcal{F}}_{match}(\varphi_1, \varphi_2)$ can be reduced to the computation of the maximum of a finite set of bottleneck distances between persistence diagrams, for every $\varphi_1, \varphi_2 \in \Phi$.

3.1 Comments on the use of $D^{\mathcal{F}}_{match}$

The goal of this paper is to propose $D^{\mathcal{F}}_{match}$ as a comparison tool that shares with the natural pseudo-distance d_G the property of being invariant under the action of a given group of homeomorphisms, but is more suitable than d_G for computation and applications. As for this subject, two observations are important.

On the one hand, the reader could think of the direct approximation of d_G as a valid alternative to the use of $D^{\mathcal{F}}_{match}$. This approach would lead to consider a *finite* subgroup H of G and to compute the pseudo-metric $d_H(\varphi_1, \varphi_2) = \min_{h \in H} \|\varphi_1 - \varphi_2 \circ h\|_\infty$ as an approximation of d_G. Unfortunately, in many cases we cannot obtain a good approximation of the topological group G by means of a finite subgroup H, even if G is compact. As a simple example, we can consider the group $G = SO(3)$ of all orientation-preserving isometries of $\mathbb{R}^3$ that take the point $(0,0,0)$ to itself. Obviously, $SO(3)$ is a compact topological group with respect to the topology of uniform convergence. Now, we restrict the homeomorphisms in $SO(3)$ to the 2-sphere $S^2 \subset \mathbb{R}^3$, and endow S^2 with the Euclidean metric. With a little abuse of notation, we maintain the symbol $SO(3)$ for this new group of homeomorphisms. The group $SO(3)$ acts on the set Φ of the 1-Lipschitz functions from S^2 to $\mathbb{R}$, by composition on the right. It can be easily shown that Φ is a compact topological space with respect to the topology induced by the sup-norm by applying the Ascoli-Arzelà Theorem. Moreover, it is not difficult to prove that a positive constant c exists such that for every finite subgroup H of $SO(3)$ we can find two functions $\varphi_H, \psi_H \in \Phi$ with $d_{SO(3)}(\varphi_H, \psi_H) = 0$ and $d_H(\varphi_H, \psi_H) \geq c$. It follows that the approximation error $\|d_H - d_{SO(3)}\|_\infty$ is greater than a positive constant c for any finite subgroup H of $SO(3)$, where c does depend on H.

We recall that the attempt of approximating G by a set S instead of a group H appears inappropriate, because if S is not a group then the function $\mu_S(\varphi_1, \varphi_2) := \min_{s\in S} \|\varphi_1 - \varphi_2 \circ s\|_\infty$ is not a pseudo-metric (see Remark 2). This makes the use of μ_S impractical for data retrieval.

It follows that, in general, the idea of a direct approximation of d_G seems unsuitable for applications. For the general problem of sampling $SO(3)$, we refer the interested reader to the paper [22].

On the other hand, $D^{\mathcal{F}^*}_{match}$ is always a strongly G-invariant pseudo-metric giving a lower bound for d_G, for any subset $\mathcal{F}^*$ of $\boldsymbol{\mathcal{F}}(\Phi, G)$. Moreover, if $\mathcal{F}^* \subseteq \mathcal{F}$ is an ϵ-approximation of $\mathcal{F}$ and Φ is bounded, the pseudo-metric $D^{\mathcal{F}^*}_{match}$ is a 2ϵ-approximation of $D^{\mathcal{F}}_{match}$ (Proposition 14). We have also seen (Corollary 19) that the existence of a finite ϵ-approximation $\mathcal{F}^* \subseteq \mathcal{F}$ of any $\mathcal{F} \subseteq \boldsymbol{\mathcal{F}}(\Phi, G)$ is always guaranteed in the case that Φ is compact. Therefore, at least in this case, there is no obstruction to obtain a finite set $\mathcal{F}^*$ for which the pseudo-metric $D^{\mathcal{F}^*}_{match}$ is an arbitrarily good approximation of $D^{\mathcal{F}}_{match}$, contrary to what happens for the pseudo-distance d_G. Indeed, we have shown that no finite subgroup H exists for which the pseudo-distance d_H is an arbitrarily good approximation of d_G, in general. In other words, $D^{\mathcal{F}}_{match}$ has better properties than d_G with respect to approximation. Furthermore, the results of the experiments described in Sections 4 and 5 show that the use of some small family of simple operators may produce a pseudo-metric $D^{\mathcal{F}^*}_{match}$ that is not far from d_G and can be efficiently used for data retrieval, even if $\mathcal{F}^*$ is not a good approximation of $\boldsymbol{\mathcal{F}}(\Phi, G)$.

These observations justify the use of $D^{\mathcal{F}}_{match}$ in place of d_G, for practical purposes.

We also wish to underline the dual nature of our approach. When G becomes "larger and larger" the associated family $\boldsymbol{\mathcal{F}}(\Phi, G)$ of G-invariant non-expansive operators becomes "smaller and smaller", so making the computation of $D^{\boldsymbol{\mathcal{F}}(\Phi,G)}_{match}$ easier and easier, contrarily to what happens for the direct computation of d_G. In other words, the approach based on $D^{\boldsymbol{\mathcal{F}}(\Phi,G)}_{match}$ seems to be of use exactly when d_G is difficult to compute in a direct way. Moreover, assuming that $\mathcal{F}$ is a finite subset of $\boldsymbol{\mathcal{F}}(\Phi, G)$ and H is a finite subgroup of G, the duality in the definitions of $D^{\boldsymbol{\mathcal{F}}(\Phi,G)}_{match}$ and d_G causes another important difference in the use of $D^{\mathcal{F}}_{match}$ and d_H as respective approximations. It consists in the fact that while $D^{\mathcal{F}}_{match}$ is a *lower* bound for $D^{\boldsymbol{\mathcal{F}}(\Phi,G)}_{match} = d_G$, d_H is an *upper* bound for d_G. As a consequence, if we take the pseudo-metric d_G as the ground truth, the retrieval errors associated with the use of $D^{\mathcal{F}}_{match}$ are just false positive, while the ones associated with the use of d_H are just false negative.

Remark 20 *The main purpose of this paper is not to pursue the approximation of d_G by using $D^{\mathcal{F}}_{match}$ via Theorem 16. Indeed, on the one hand that theorem does not say anything about the way of choosing a suitable set of operators. On the other hand it could be that the use of $D^{\mathcal{F}}_{match}$ to approximate d_G requires a family of operators whose complexity equals the one of directly approximating d_G via brute force. This would not be strange, because the current state of development of research does not allow to estimate d_G from a practical point of view, generally speaking. We highlight that the problem of quickly approximating the natural pseudo-distance is unsolved also in the case of G equal to* $\mathrm{Homeo}(X)$*, to the best of the authors' knowledge. Moreover, the only result we know concerning the approximation of d_G via persistent homology is limited to filtering functions from S^1 to $\mathbb{R}^2$ [19], and its relevance is purely theoretical.*

Therefore, our main purpose is to introduce a new and easily computable pseudo-metric that is a lower bound for d_G. Nevertheless, we can make two relevant observations. First of all, the path to the approximation of d_G via $D^{\mathcal{F}}_{match}$ is not closed, even if it probably requires to develop further ideas. Indeed, Theorem 17 states the compactness of the set of all non-expansive G-operators in the case that Φ is compact, so laying the groundwork for the study of new approximation schemes. Secondly, even if no theoretical approach to the choice of our operators is presently available, it can happen that the use of some small family of simple operators produces a pseudo-metric $D^{\mathcal{F}}_{match}$ that is not far from d_G. We shall devote Section 4 to check this possibility in an experiment concerning data represented by functions from $\mathbb{R}$ to $\mathbb{R}$.

Remark 21 *The pseudo-distance $D^{\mathcal{F}}_{match}$ is based on the set $\mathcal{F}$. The smallest set $\mathcal{F}$ of non-expansive G-operators such that $D^{\mathcal{F}}_{match}$ coincides with the natural pseudo-distance d_G is the one containing just the operators F_ψ defined in the proof of Theorem 16. However, this trivial set of operators is completely useless from the point of view of applications, since computing F_ψ for every $\psi \in \Phi$ is equivalent to computing the natural pseudo-distance d_G. As for the applications to shape comparison, we need the operators in $\mathcal{F}$ to be simple to compute and $\mathcal{F}$ to be small, but still large enough to guarantee that $D^{\mathcal{F}}_{match}$ is not too far from d_G.*

4 Experiments

In the previous section we have seen that our approach to shape comparison via non-expansive G-operators applied to persistent homology allows to get invariance with respect to arbitrary subgroups G of $\mathrm{Homeo}(X)$. However, some assumptions are required, concerning G and the set Φ of admissible filtering functions. A natural question arises about what happens in practical applications, when our assumptions are not always guaranteed to hold. To answer this question, we provide numerical results for some experiments concerning piecewise linear functions. Our experiments may be described as the construction of a dataset that provides functionality to retrieve the most similar functions with respect to a given "query" function, after arbitrarily choosing an invariance group G.

The goal of this section is to show that our approximation of the natural pseudo-distance d_G via the use of a finite subset of operators is well behaving. A motivating factor is Corollary 19, stating that if the set Φ of admissible filtering functions is compact, then for every set of operators $\mathcal{F}$ there exists a finite set of operators $\mathcal{F}^*$ such that the pseudo-distance induced by $\mathcal{F}^*$ is ϵ-close to the pseudo-distance induced by the set $\mathcal{F}$, even if $\mathcal{F}$ is an infinite set. While it is impractical to use the proof of Corollary 19 to build the finite set $\mathcal{F}^*$, we show that a small subset of $\mathcal{F}$ is sufficient in several applications, both in the compact and in the non-compact case.

Our ground truth for shape comparison is the pseudo-distance d_G, approximated by brute force methods (when possible). On the one hand, the approximation of d_G usually has a large computational cost, as we shall see in this section. On the other hand, $D^{\mathcal{F}}_{match}$ allows to get a simple and easy-implementable approximation of d_G. This fact justifies our approach.

In our experiments we have set $\Phi := C^0(\mathbb{R}, \mathbb{R})$. We have chosen to work with a dataset $\Phi_{ds} \subset \Phi$ of 20.000 piecewise linear functions $\varphi : \mathbb{R} \to [-1, 1]$, with support contained in the closed interval $[-1, 2]$. In order to obtain the graph of each function we have randomly chosen eight points $\{(x_i, y_i)\}_{i\in\{0,..,7\}}$ in the rectangle $[-1, 2] \times [-1, 1]$ such that y-coordinate of the first and the last point equal to 0 and $-1 < x_0 < x_1 < \ldots < x_7 < 2$. The graph of the function is obtained by connecting (x_i, y_i) to (x_{i+1}, y_{i+1}) by a segment for $0 \le i \le 6$.

Additionally, for computational reasons we require that all functions are Lipschitz, with a given Lipschitz constant C, hence each function with a Lipschitz constant greater than C is filtered out. An example of a randomly generated function is presented in Fig. 3.

The choice of our dataset yields useless persistent Betti number functions in degree $k > 0$. Therefore, in our experiments we will use just persistent Betti number functions in degree $k = 0$.

Invariance groups and operators To evaluate the approach described in this paper, we use five invariance groups G_i, for $i = 1, .., 5$. Each group G_i induces a strongly G_i-invariant pseudo-metric d_{G_i}. Then we define a set $\mathcal{F}^*_i$ of non-expansive G_i-operators for each group G_i. Here is the list of the groups we have used in our experiments:

1. G_1: the group of all affine transformations from $\mathbb{R}$ to $\mathbb{R}$;
2. G_2: the group of all orientation-preserving affine transformations from $\mathbb{R}$ to $\mathbb{R}$;
3. G_3: the group of all isometries of $\mathbb{R}$;
4. G_4: the group of all translations of $\mathbb{R}$;
5. G_5: the trivial group containing just the identity map $id : \mathbb{R} \to \mathbb{R}$.

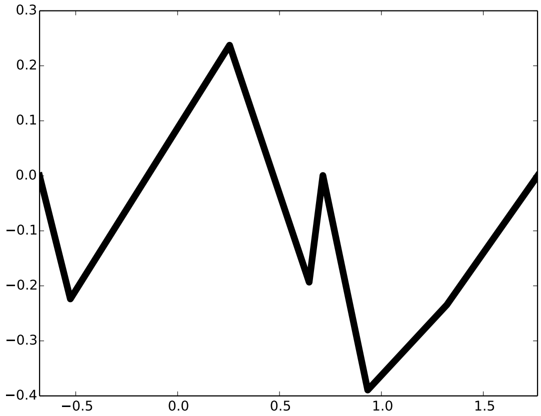

Fig. 3 One of the functions used in our experiments. The function is zero outside the closed interval $[-1, 2]$.

In our experiments we have tested the use of the pseudo-distances $D^{\mathcal{F}_i^*}_{match}$ to decide if two functions in our dataset are similar with respect to the invariance groups G_i. This approach avoids the computation of the natural pseudo-distances d_{G_i}, which can be hard to approximate.

We have selected our G-invariant non-expansive operators just relying on two practical rules of thumb: 1) Their apparent link with perceptual and geometric properties; 2) The attempt to choose them as different as possible from each other.

Finding the most similar function with respect to the chosen invariance group After constructing the dataset Φ_{ds} that we have previously described, we compute the 0-th persistence diagram of $F(\varphi)$, for every $\varphi \in \Phi_{ds}$ and every $F \in \mathcal{F}_i^*$, varying the index i. Afterwards, we choose a "query" function φ_q in our dataset, which will be compared with all functions in Φ_{ds}. Finally, we compute $D^{\mathcal{F}_i^*}_{match}(\varphi_q, \varphi)$ for every $\varphi \in \Phi_{ds}$ and $i = 1, \ldots, 5$. In Figures 4–8 we show the most similar functions with respect to $D^{\mathcal{F}_1^*}_{match}$, ..., $D^{\mathcal{F}_5^*}_{match}$ (i.e. the functions φ minimizing $D^{\mathcal{F}_i^*}_{match}(\varphi_q, \varphi)$, with $\varphi \neq \varphi_q$).

In the next subsections we describe the operators that we have used in our experiments, for each of the invariance groups G_i.

4.1 Invariance with respect to the group G_1 of all affinities of the real line

The first group that we consider, denoted by G_1, consists of all affinities of the real line (i.e. maps $x \mapsto \alpha x + \beta$ with $\alpha \neq 0$, $\beta \in \mathbb{R}$). Intuitively, we can squeeze, stretch, horizontally reflect and translate the graph of the function.

In order to define the non-expansive G_1-operators that we will use in this section, we introduce the operator $F_{\hat{w},\hat{c}}$ defined as:

$$F_{\hat{w},\hat{c}}(\varphi)(x) := \sup_{r\in\mathbb{R}} \sum_{i=1}^{n} w_i \cdot \varphi\left(x + rc_i\right), \tag{4.1}$$

where $\hat{w}$ and $\hat{c}$ are two vectors $\hat{w} := (w_1, \dots, w_n)$, $\hat{c} := (c_1, \dots, c_n)$ in $\mathbb{R}^n$. We additionally require $\sum_{i=1}^{n} |w_i| = 1$, so we can easily check that $F_{\hat{w},\hat{c}}$ is a non-expansive G_1-operator. From the computational point of view, the operator $F_{\hat{w},\hat{c}}$ can be approximated by substituting the supremum in its definition with a maximum for r belonging to a finite set. For more details about this, see Appendix C.

In order to apply our method to approximate d_{G_1}, we will consider the set $\mathcal{F}_1^*$, consisting of the following non-expansive G_1-operators:

- The identity operator;
- F_1^b, defined by setting $F_1^b(\varphi)(x) = -\varphi(x)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_1^c to F_1^n based on (4.1) with the following arbitrarily chosen values of $\hat{w}, \hat{c}$.
  $F_1^c : \hat{w} = (0.3, 0.4, 0.3), \hat{c} = (0.3, 0.6, 0.9)$
  $F_1^d : \hat{w} = (-0.3, 0.4, -0.3), \hat{c} = (0.3, 0.6, 0.9)$
  $F_1^e : \hat{w} = (-0.3, 0.4, -0.3), \hat{c} = (0.1, 0.5, 0.6)$
  $F_1^f : \hat{w} = (0.3, 0.4, 0.3), \hat{c} = (0.1, 0.5, 0.9)$
  $F_1^g : \hat{w} = (-0.3, 0.4, -0.3), \hat{c} = (0.1, 0.2, 0.5)$
  $F_1^h : \hat{w} = (-0.3, -0.4, -0.3), \hat{c} = (0.1, 0.3, 0.3)$
  $F_1^i : \hat{w} = (0.3, -0.4, 0.3), \hat{c} = (0.3, 0.5, 0.6)$
  $F_1^j : \hat{w} = (0.5, 0.5), \hat{c} = (0.1, 0.5)$
  $F_1^k : \hat{w} = (0.5, -0.5), \hat{c} = (0.1, 0.5)$
  $F_1^l : \hat{w} = (-0.5, -0.5), \hat{c} = (0.1, 0.5)$
  $F_1^m : \hat{w} = (0.3, -0.4, 0.3), \hat{c} = (0.1, 0.6, 0.9)$
  $F_1^n : \hat{w} = (0.3, -0.4, 0.3), \hat{c} = (0.1, 0.2, 0.9)$.

In Fig. 4 we show an example of retrieval in our dataset. The two functions that are most similar to a given query function are displayed. We also show the alignments of the retrieved functions to the query function. One can notice that if we restrict ourselves to consider a finite set S_1 of affinities from $\mathbb{R}$ to $\mathbb{R}$, the inclusion $S_1 \subset G_1$ implies that for every φ_1, φ_2 in our dataset

$$D_{match}^{\mathcal{F}_1^*}(\varphi_1, \varphi_2) \leq d_{G_1}(\varphi_1, \varphi_2) \leq \min_{g\in S_1} \|\varphi_1 - \varphi_2 \circ g\|_\infty.$$

This is due to the stability of $D_{match}^{\mathcal{F}}$ with respect to the natural pseudo-distance d_G associated with the group G (Theorem 15), and to the definition of natural pseudo-distance. It follows that if

$$\left| D_{match}^{\mathcal{F}_1^*}(\varphi_1, \varphi_2) - \min_{g\in S_1} \|\varphi_1 - \varphi_2 \circ g\|_\infty \right| \leq \epsilon$$

then

$$\left| D_{match}^{\mathcal{F}_1^*}(\varphi_1, \varphi_2) - d_{G_1}(\varphi_1, \varphi_2) \right| \leq \epsilon.$$

This inequality suggests a method to evaluate the approximation of d_{G_1} that we obtain by means of $D_{match}^{\mathcal{F}_1^*}$, via an estimate of $\min_{g\in S_1}\|\varphi_1-\varphi_2\circ g\|_\infty$. We choose $c := 10C$ and discretize the domains for α and β in the affinity $x\mapsto \alpha x+\beta$, by considering two sets $\{\alpha_1,\ldots,\alpha_r\}$ and $\{\beta_1,\ldots,\beta_s\}$, with $1/c\le|\alpha_i|\le c$ and $-c\le\beta_j\le c$ for $1\le i\le r$ and $1\le j\le s$. Then we compute $\min_{i,j}\|\varphi_1-\varphi_2\circ g_{ij}\|_\infty$, where $g_{ij}(x):=\alpha_i x+\beta_j$. In practice, we set $C=5$ and discretize the intervals $[\frac{1}{50},50]$ and $[-50,50]$ by choosing equidistant points with the distance between neighboring points equal to 0.01. This approach requires the computation of the sup-distance d_∞ between φ_1 and $\varphi_2\circ g_{ij}$ for a total number of rs functions g_{ij}. The overall computation of $\min_{i,j}\|\varphi_1-\varphi_2\circ g_{ij}\|_\infty$ is slow, and we performed it just to find an upper bound for the distance between $D_{match}^{\mathcal{F}_1^*}$ and d_{G_1}, in order to evaluate our method. Actually, the purpose of our approach is to avoid the computation of $\min_{i,j}\|\varphi_1-\varphi_2\circ g_{ij}\|_\infty$ and d_{G_1}, by substituting d_{G_1} with $D_{match}^{\mathcal{F}_1^*}$.

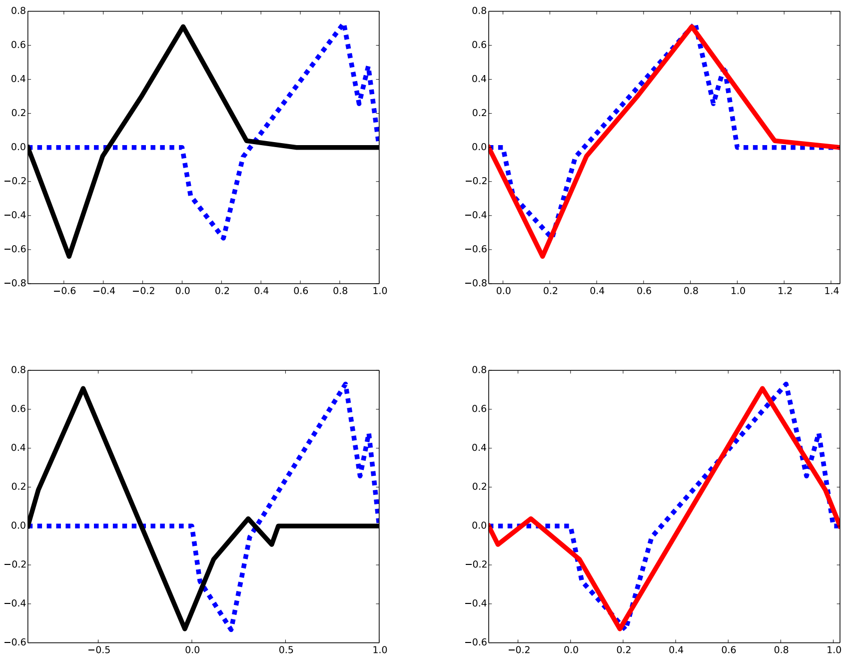

Fig. 4 Output of the experiment concerning the invariance group G_1: the most similar (solid black line, left upper plot) and the second most similar function (solid black line, left lower plot) with respect to the query function (dotted blue line) are displayed. In these two cases the natural pseudo-distance d_{G_1} takes the respective values 0.35 and 0.2. On the right side of the figure the results of alignment of the retrieved function to the query function are displayed (solid red lines). These alignments are obtained via brute force computation, by approximating the affine transformations in G_1. They are added to allow a visual and qualitative comparison.

We encourage the readers to analyze the presented results. We point out that in Fig. 4 the graphs of the retrieved functions (solid black lines) are similar to the graph of the "query" function (dotted blue line), with respect to the group G_1. The red graphs show how we can get good alignments of the retrieved functions to the query function by applying affine transformations.

4.2 Invariance with respect to the group G_2 of all orientation-preserving affinities of the real line

The second group that we consider, denoted by G_2, consists of all affinities of the real line that preserve the orientation (i.e. maps $x \mapsto \alpha x + \beta$ with $\alpha > 0$ and $\beta \in \mathbb{R}$). This group is smaller than G_1 - one cannot use reflections to align the functions. With reference to the operator $F_{\hat{w},\hat{c}}$ defined in (4.1), it is easy to check that after changing the condition $r \in \mathbb{R}$ to $r > 0$ in (4.1), the operator is invariant under the action of the group of all orientation-preserving affinities. Let us denote this new operator as $\bar{F}_{\hat{w},\hat{c}}$. With reference to the other operators used in previous subsection 4.1, we know that the identity operator and F_1^b are non-expansive G_2-operators. As a consequence we use them also for the invariance group G_2, adding the operators F_2^a, F_2^b, F_2^c, F_2^d and F_2^e defined for the following arbitrarily chosen values of $\hat{w}, \hat{c}$ in $\bar{F}_{\hat{w},\hat{c}}$.
$F_2^a : \hat{w} = (0.3, 0.4, 0.3), \hat{c} = (0.3, 0.6, 0.9)$
$F_2^b : \hat{w} = (-0.3, 0.4, -0.3), \hat{c} = (0.3, 0.6, 0.9)$
$F_2^c : \hat{w} = (-0.2, 0.2, -0.2, 0.2, -0.2), \hat{c} = (0.2, 0.4, 0.6, 0.8, 1.0)$
$F_2^d : \hat{w} = (0.2, -0.2, 0.2, -0.2, 0.2), \hat{c} = (0.2, 0.4, 0.6, 0.8, 1.0)$
$F_2^e : \hat{w} = (-0.1, 0.2, -0.4, 0.2, -0.1), \hat{c} = (0.2, 0.4, 0.6, 0.8, 1.0)$.

In plain words, the group G_2 does not allow reflections, but only squeezing/stretching and translations. In Fig. 5 we show an example of retrieval in our dataset. The two functions (solid black lines) that are most similar to a given query function (dotted blue line) are displayed. The red graphs show how we can get good alignments of the retrieved functions to the query function by applying orientation-preserving affine transformations. These alignments have been obtained by approximating the transformations in G_2.

4.3 Invariance with respect to the group of all isometries of the real line

The third group that we consider, denoted by G_3, consists of all isometries of the real line (i.e. maps $x \mapsto \alpha x + \beta$ with $\alpha = \pm 1$ and $\beta \in \mathbb{R}$).

Since $G_3 \subseteq G_1$, the operators that we have used for comparison with respect to the group G_1 are also G_3-operators. As a consequence we can use them also for the invariance group G_3, adding the operators F_3^a, F_3^b, F_3^c, F_3^d and F_3^e defined as follows:

- F_3^a, defined by setting $F_3^a(\varphi)(x) = \max\left(\varphi\left(x - \frac{1}{4}\right), \varphi(x), \varphi\left(x + \frac{1}{4}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;

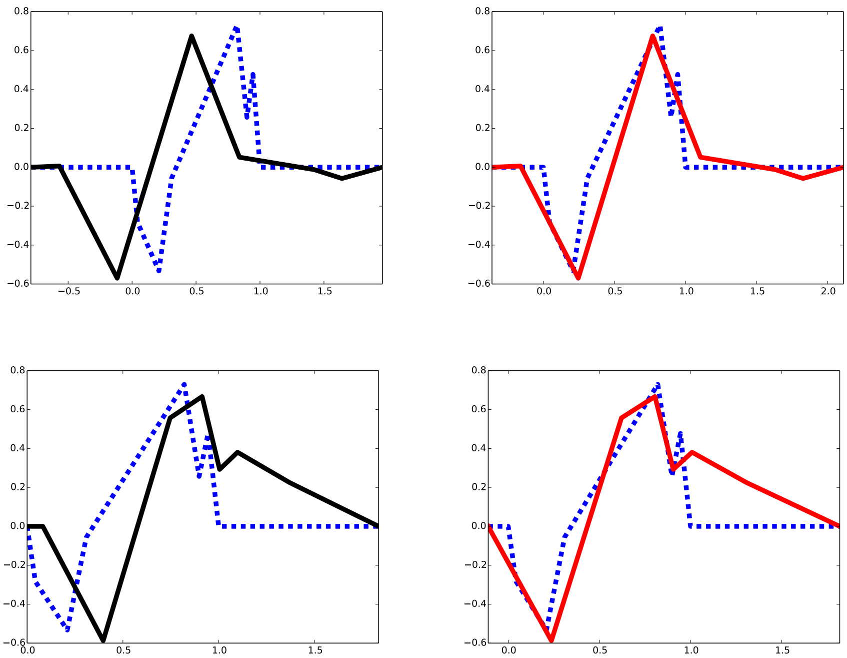

Fig. 5 Output of the experiment concerning the invariance group G_2. Color and type of plots are the same as in Fig. 4. In both the displayed cases the natural pseudo-distance d_{G_2} takes the value 0.4.

- F_3^b, defined by setting $F_3^b(\varphi)(x) = \frac{1}{3}\left(\varphi\left(x-\frac{1}{4}\right)+\varphi(x)+\varphi\left(x+\frac{1}{4}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_3^c, defined by setting $F_3^c(\varphi)(x) = \frac{1}{3}\left(\varphi\left(x-\frac{1}{3}\right)+\varphi(x)+\varphi\left(x+\frac{1}{3}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_3^d, defined by setting $F_3^d(\varphi)(x) = \frac{1}{5}\left(\varphi\left(x-\frac{1}{3}\right)+\varphi\left(x-\frac{1}{4}\right)+\varphi(x)+\varphi\left(x+\frac{1}{4}\right)+\varphi\left(x+\frac{1}{3}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_3^e, defined by setting $F_3^e(\varphi)(x) = \max\left(\varphi\left(x-\frac{1}{3}\right), \varphi\left(x-\frac{1}{4}\right), \varphi(x), \varphi\left(x+\frac{1}{4}\right), \varphi\left(x+\frac{1}{3}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$.

In Fig. 6 we show an example of retrieval in our dataset. The two functions (solid black lines) that are most similar to a given query function (dotted blue line) are displayed.

4.4 Invariance with respect to the group of all translations of the real line

The fourth group that we consider, denoted by G_4, consists of all translations of the real line (i.e. maps $x \mapsto x + \beta$ with $\beta \in \mathbb{R}$).

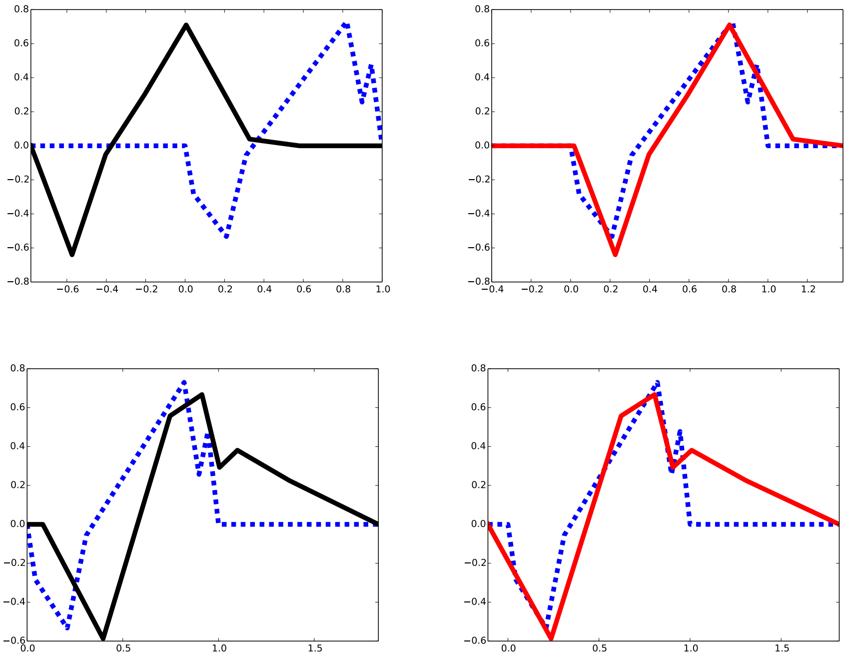

Fig. 6 Output of the experiment concerning the invariance group G_3. Color and type of plots are the same as in Fig. 4. In both the displayed cases the natural pseudo-distance d_{G_3} takes the value 0.4.

Since $G_4 \subseteq G_3$, the operators that we have used for comparison with respect to the group G_3 are also G_4-operators. As a consequence we can use them also for the invariance group G_4, adding the operators F_4^a, F_4^b, F_4^c, F_4^d and F_4^e defined as follows:

- F_4^a, defined by setting $F_4^a(\varphi)(x) = \max\left(\varphi(x), \varphi\left(x+\frac{1}{4}\right), \varphi\left(x+\frac{1}{3}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_4^b, defined by setting $F_4^b(\varphi)(x) = \frac{1}{4}\varphi\left(x-\frac{1}{4}\right) + \frac{3}{4}\varphi(x)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_4^c, defined by setting $F_4^c(\varphi)(x) = \frac{3}{4}\varphi(x) + \frac{1}{4}\varphi\left(x+\frac{1}{4}\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_4^d, defined by setting $F_4^d(\varphi)(x) = \frac{1}{3}\varphi(x) + \frac{1}{6}\varphi\left(x+\frac{1}{5}\right) + \frac{1}{2}\varphi\left(x+\frac{3}{5}\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_4^e, defined by setting $F_4^e(\varphi)(x) = \max\left(\varphi(x), \varphi\left(x+\frac{1}{5}\right), \varphi\left(x+\frac{3}{5}\right)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$.

In Fig. 7 we show an example of retrieval in our dataset. The two functions (solid black lines) that are most similar to a given query function (dotted blue line) are displayed.

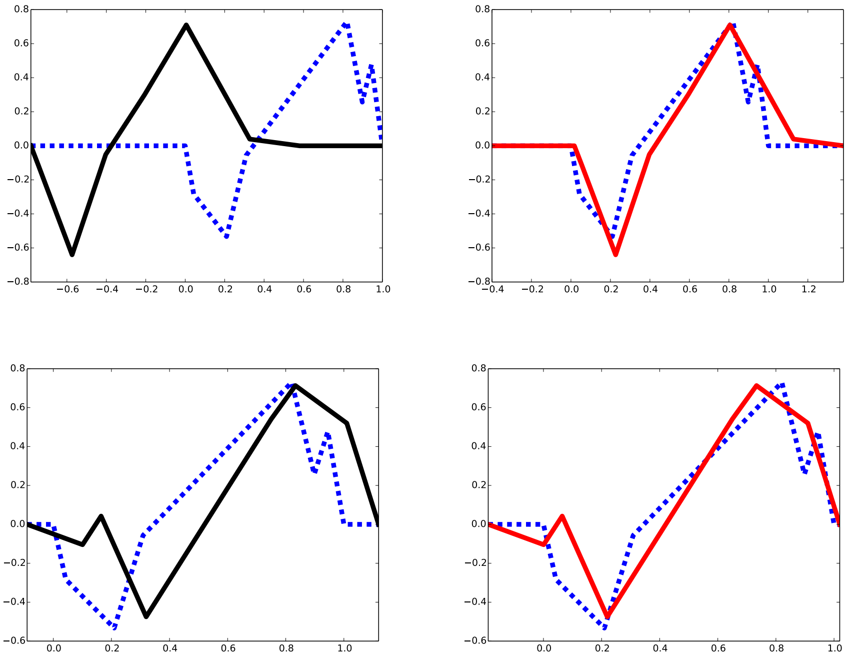

Fig. 7 Output of the experiment concerning the invariance group G_4. Color and type of plots are the same as in Fig. 4. In these two cases the natural pseudo-distance d_{G_4} takes the value 0.4.

4.5 Invariance with respect to the trivial group

The fifth (and last) group that we consider, denoted by G_5, is the trivial group Id containing just the identity. We observe that the concept of Id-operator coincides with the concept of operator.

Since $G_5 \subseteq G_4$, the operators that we have used for comparison with respect to the group G_4 are also G_5-operators. As a consequence we can use them also for the invariance group $G_5 = Id$, adding the operators F_5^a, F_5^b, F_5^c, F_5^d and F_5^e defined as follows:

- F_5^a, defined by setting $F_5^a(\varphi)(x) = \varphi(x)\sin(5\pi x)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_5^b, defined by setting $F_5^b(\varphi)(x) = \varphi(x)\sin(9\pi x)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_5^c, defined by setting $F_5^c(\varphi)(x) = (\varphi(x) + 2) \cdot f_{\frac{1}{4}}(x)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_5^d, defined by setting $F_5^d(\varphi)(x) = (\varphi(x) + 2)) \cdot f_{\frac{1}{2}}(x)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$;
- F_5^e, defined by setting $F_5^e(\varphi)(x) = \frac{1}{2}(\varphi(x) + 2) \cdot \left(f_{\frac{3}{8}}(x) + f_{\frac{5}{8}}(x)\right)$ for every $\varphi \in \Phi$ and every $x \in \mathbb{R}$

where $f_\mu(x) = e^{-\left(\frac{x-\mu}{0.1}\right)^2}$.

In Fig. 8 we show an example of retrieval in our dataset. The two functions (solid black lines) that are most similar to a given query function (dotted blue line) are displayed. No alignment is necessary here, since the unique allowed transformation is the identity, and d_{G_5} equals the sup-norm.

From the practical point of view, the computation of d_{G_5} can be done directly, without using the pseudo-distance $D^{\mathcal{F}_5^*}_{match}$ as an approximation. However, we decided to include this last experiment for the sake of completeness, in order to show how our method behaves also in this trivial case.

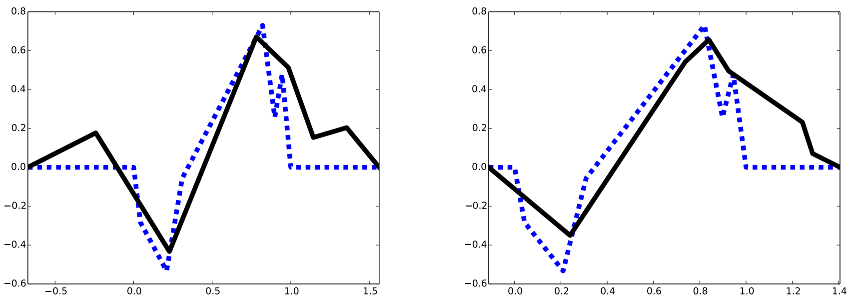

Fig. 8 Output of the experiment concerning the invariance group $G_5 = Id$. The most similar and the second most similar function (solid black lines) with respect to the query function (dotted blue line) are displayed. In both the displayed cases the natural pseudo-distance d_{G_5} takes the value around 0.4.

4.6 Quantitative results

The purpose of this section is to give a quantitative estimate of the approximation of the natural pseudo-distance d_{G_i} via the pseudo-distance $D^{\mathcal{F}_i^*}_{match}$. Due to the time-consuming nature of the computation of d_{G_i}, we use only part of the set Φ_{ds}.

In Table 1 we show the mean value of d_{G_i} for $i = 1, \ldots, 5$ and statistics of the error made by substituting d_{G_i} with $D^{\mathcal{F}_i^*}_{match}$: mean absolute error (MAE) and mean relative error (MRE).

On average, the relative error results are around 0.20, with the best results for the group G_1. The results displayed in Table 1 show that a small set of operators is sufficient to produce a relatively good approximation of the pseudo-distances d_{G_i} that we have considered. The most important question and natural next step is to find heuristics or optimal methods to decide which operators bring most information.

In our opinion it is surprising that a set of just a few operators is sufficient to get a good approximation of any natural pseudo-distance d_{G_i}. This fact

Group	Mean d_{G_i}	MAE	MRE
G_1	0.54	0.08	0.18
G_2	0.56	0.09	0.20
G_3	0.61	0.10	0.19
G_4	0.63	0.11	0.19
G_5	0.91	0.22	0.24

Table 1 Mean values for d_{G_i}, together with the mean absolute error (MAE) and the mean relative error (MRE) made by substituting d_{G_i} with $D_{match}^{\mathcal{F}_i^*}$. These values have been computed on 1000 functions from Φ_{ds} (5×10^5 pairs) in case of G_3, G_4, G_5 and on 100 functions (5×10^3 pairs) in case of G_1 and G_2. The reason of not using the whole set of functions was computation time for the brute force approximation.

seems quite promising, since it opens the way to an alternative approach to approximate the natural pseudo-distance, besides the one based on brute-force computation.

5 Towards an image retrieval system

The experiment in the previous section was prepared to show quantitative results. In this section we present another experiment, whose goal is to show qualitative results. We demonstrate how our approach can be used in a simple image retrieval task, where the invariance group consists of all isometries of $\mathbb{R}^2$. We compute the pseudodistance $D_{match}^{\mathcal{F}^*}$ using ten operators, and referring just to homology in degree 0. While in the previous experiment there was no reason to use homology in degree 1, this could be of use in the image case. Nevertheless, we do not include it, in order to speed up computations.

We keep the notation consistent with the one in the previous section. The dataset of objects consists of 10.000 grey-level images with three to six spots. Each spot is generated by adding a 2D bump function with randomly chosen size at a randomly chosen position, whereas images are represented as functions from $\mathbb{R}^2$ to the interval $[0, 1]$ with support in the square $[0, 1]^2$. However, in this experiment the set Φ of admissible filtering functions will consists of all continuous functions from $\mathbb{R}^2$ to $[-1, 1]$ with compact support. This choice will allow us to use a wider range of operators.

In order to skip unnecessary technical details, we confine ourselves to give a concise description of the operators that we have used. We have chosen ten operators, divided into two families of five. The first family consists of the following operators:

1. Identity operator;
2. Four operators based on convolution of the image φ with different kernels. These operators are formally defined in the Appendix D.

The second family consists of the operators that we can obtain by reversing the sign of the previous five operators. Overall, we have ten operators.

Fig. 9 Examples of grey-level images from the dataset used in the second experiment. Black and white represent the values 1 and 0, respectively.

The reader can easily verify that our ten operators are non-expansive and G-invariant, when G is the group of all isometries of $\mathbb{R}^2$.

We look for the most similar images to three images from the dataset. In each of the Figures 10, 11 and 12 we present the three images that minimize the pseudo-distance $D^{\mathcal{F}^*}_{match}$ from the query image.

Remark 22 *Due to the randomness of the method we used to construct our images, it is quite unlikely that two images in our dataset are identical (or even nearly identical), especially when the number of bumps that appear in them is* 5 *or* 6*. We would like to underline that our goal is not to find images that are equal to each other, but images that resemble each other with respect to the group of isometries.*

Remark 23 *We intentionally decided to build our dataset by producing images that do not encode meaningful information for humans. Otherwise, the qualitative results would be biased by a priori knowledge of the image content. In our research we focus on topological and geometrical properties, and neglect the perceptual aspects of image comparison. Therefore, we decided not to use standard datasets from image comparison and retrieval projects.*

Discussion and future work

In our paper we have described a method to combine persistent homology and the invariance with respect to a given group G of homeomorphisms, acting on a set Φ of filtering functions. This technique allows us to treat G as a variable

(a) Primary image

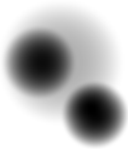

(b) Result 1

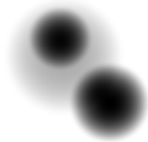

(c) Result 2

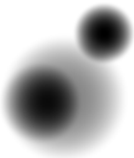

(d) Result 3

Fig. 10 Most similar images to the query image from a dataset of 2D artificial images. Computed pseudo-distances $D^{\mathcal{F}^*}_{match}$ are respectively 0.035, 0.042 and 0.050 for images 10(b), 10(c) and 10(d). The second image can be approximately obtained from the first one via reflection and rotation, which are both isometries. The third and the fourth images require only rotation.

in our problem, and to distinguish functions that are not directly distinguishable in the classical setting. Our approach is based on a new pseudo-distance depending on a set of G-invariant non-expansive operators, that approximates the natural pseudo-distance in the limit.

Some relevant questions remain open:

- How can we choose the G-invariant operators in order to get the best possible results, depending on the set Φ and the group G? How large should the set of operators be? Is it possible to build a dictionary of operators to be used for a specific group?
- How could our theoretical results be applied to problems in shape comparison?
- What is the best way to adapt the use of G-invariant non-expansive operators to multidimensional persistence?

As for the last question, we observe that it would be interesting to extend our approach to cases involving signals that are naturally described by $\mathbb{R}^k$-valued functions, such as color images or surfaces in $\mathbb{R}^3$. This kind of data is common in many applications, so that such an advancement would greatly enhance the use of our method in topological data analysis.

Our first experiments suggest that the introduced method is pretty robust, taking advantage from the stability of persistent homology. We hope that this

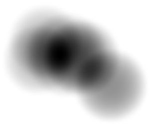

(a) Primary image

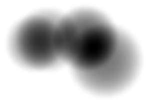

(b) Result 1

(c) Result 2

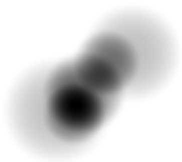

(d) Result 3

Fig. 11 Most similar images to the query image from a dataset of 2D artificial images. Computed pseudo-distances $D_{match}^{\mathcal{F}^*}$ are respectively 0.039, 0.046 and 0.050 for images 11(b), 11(c) and 11(d).

property can open the way to new applications of the concept of persistence, in presence of constraints concerning the invariance of our data.

The reader can note that our approach could be extended to the use of invariants that do not arise from persistent homology. It is sufficient that each invariant I that we consider is stable and invariant under the action of every homeomorphism. In other words, we require that I is a function from Φ to a metric space (D, d_D) of descriptors, such that $d_D(I(\varphi_1), I(\varphi_2)) \leq \|\varphi_1 - \varphi_2\|_\infty$ for every $\varphi_1, \varphi_2 \in \Phi$, and $I(\varphi \circ g) = I(\varphi)$ for every $\varphi \in \Phi$ and $g \in G$. As three examples among many, we could use the invariants taking each φ to its size homotopy group (the analogue of persistent homology group in homotopy theory) [20], its Reeb graph [14], or also its maximum, together with suitable metrics. The reason for which we have confined ourselves to the case of persistent homology is our opinion that this theory is both endowed with pretty fast and efficient algorithms and still sufficiently powerful for shape comparison. However, we think that also the application of our method to other invariants should be explored.

In conclusion, we would also like to consider the problem of formalizing the framework that we have described in this paper in a categorical setting. In our approach, each object belonging to a given dataset is seen as a collection $\{\varphi_i : X_i \to \mathbb{R}\}$ of continuous functions, where each φ_i belongs to an admissible space Φ_i. These functions represent the measurements made on the object. No

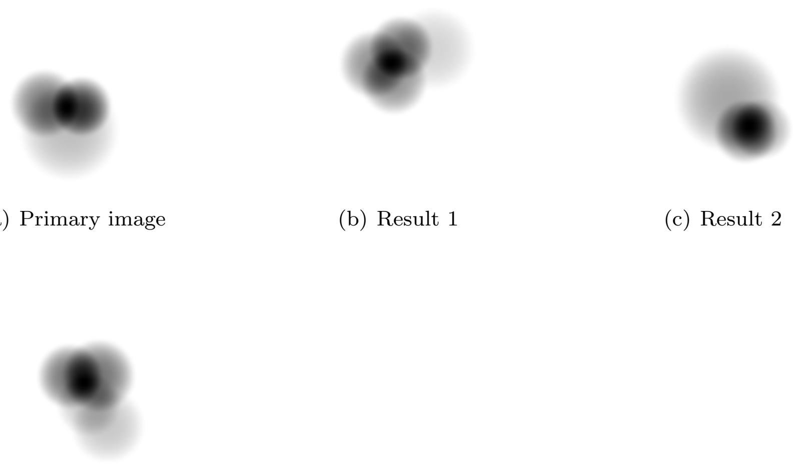

(a) Primary image (b) Result 1 (c) Result 2

(d) Result 3

Fig. 12 Most similar images to the query image from a dataset of 2D artificial images. Computed pseudo-distances $D^{\mathcal{F}^*}_{match}$ are respectively 0.011, 0.015 and 0.019 for images 12(b), 12(c) and 12(d).

attempt is made to define the objects in a direct way, according to the idea that each object is accessible just via acts of measurement (cf. [1]).

However, the measurements $\{\varphi_i : X_i \to \mathbb{R}\}$ are not directly used by the observer that has to judge about similarity and dissimilarity. Indeed, perception usually changes the signals $\{\varphi_i : X_i \to \mathbb{R}\}$ into several new (and usually simpler) collections $\{\psi_i^j : X_i \to \mathbb{R}\}$ of data. This passage is given by some operators F_i^j, taking each function φ_i into a new function ψ_i^j. In the approach that we have presented, these operators are supposed to be G-invariant and non-expansive, because perception usually benefits of some invariance and quantitative constraint. In other words, the observer is represented by an ordered family $\{F_i^j : \Phi_i \to \Phi_i\}$ of suitable operators, each one acting on a set Φ_i of admissible signals. As a consequence, two objects in the dataset can be distinguished if and only if the observer is endowed with an operator F, changing the corresponding signals into two new signals that are not equivalent with respect to the invariance group.

We think that this approach could benefit of a precise categorical formalization, and we plan to devote our research to this topic in the future. Far from being just a formal abstraction, this description in the categorical language might contribute to free the analysis of data from an approach based on the "absolute role" of the objects, and to propose an "observer-oriented" approach focused on the signals produced by those objects and on the operators that act on them. In this framework, the observer could be seen as a collection $\{F_i^j : \Phi_i \to \Phi_i\}$ of operators acting on the elements of the family

$\{\Phi_i\}$ of the sets of possible signals, and shape comparison could be interpreted as a metric property of the pair $\left(\{\Phi_i\}, \{F_i^j\}\right)$, instead of a discrimination of absolute features of the objects.

Acknowledgment

The authors thank Marian Mrozek for his suggestions and advice. The research described in this article has been partially supported by GNSAGA-INdAM (Italy), and is based on the work realized by the authors within the ESF-PESC Networking Programme "Applied and Computational Algebraic Topology". The second author is supported by National Science Centre (Poland) DEC-2013/09/N/ST6/02995 grant and by the TOPOSYS project FP7-ICT-318493-STREP. The first author is grateful to the city of Viareggio for its inspiring hospitality.

References

1. S. Biasotti, L. De Floriani, B. Falcidieno, P. Frosini, D. Giorgi, C. Landi, L. Papaleo, M. Spagnuolo, *Describing shapes by geometrical-topological properties of real functions*, ACM Computing Surveys, vol. 40, n. 4, 12:1–12:87 (2008).
2. F. Cagliari, B. Di Fabio and C. Landi, *The natural pseudo-distance as a quotient pseudo-metric, and applications*, Forum Mathematicum, vol. 27, n. 3, 1729–1742 (2015).
3. G. Carlsson, A. Zomorodian, *The theory of multidimensional persistence*, Discrete and Computational Geometry, vol. 42, no. 1, 71–93 (2009).
4. N. Cavazza, M. Ethier, P. Frosini, T. Kaczynski, C. Landi, *Comparison of persistent homologies for vector functions: from continuous to discrete and back*, Computers and Mathematics with Applications, vol. 66, 560–573 (2013).
5. A. Cerri, B. Di Fabio, M. Ferri, P. Frosini, and C. Landi, *Betti numbers in multidimensional persistent homology are stable functions*, Mathematical Methods in the Applied Sciences, vol. 36, 1543–1557 (2013).
6. F. Chazal, D. Cohen-Steiner, M. Glisse, L.J. Guibas, S.Y. Oudot, *Proximity of persistence modules and their diagrams*, SCG '09: Proceedings of the 25th annual symposium on Computational geometry, Aarhus, Denmark, 237–246 (2009).
7. C. Chen, H. Edelsbrunner, *Diffusion runs low on persistence fast*, Proceedings of the 2011 International Conference on Computer Vision, 423–430 (2011).
8. D. Cohen-Steiner, H. Edelsbrunner, J. L. Harer, *Stability of persistence diagrams*, Discrete and Computational Geometry, vol. 37, n. 1, 103–120 (2007).
9. V. de Silva, P. Skraba and M. Vejdemo-Johansson, *Topological analysis of recurrent systems*, Workshop on Algebraic Topology and Machine Learning, NIPS 2012 (2012). Preprint available at http://sites.google.com/site/nips2012topology/contributed-talks.
10. B. Di Fabio, P. Frosini, *Filtrations induced by continuous functions*, Topology and its Applications, vol. 160, 1413–1422 (2013).
11. P. Donatini, P. Frosini, *Natural pseudodistances between closed manifolds*, Forum Mathematicum, vol. 16, n. 5, 695–715 (2004).
12. P. Donatini, P. Frosini, *Natural pseudodistances between closed surfaces*, Journal of the European Mathematical Society, vol. 9, n. 2, 231–253 (2007).
13. P. Donatini, P. Frosini, *Natural pseudodistances between closed curves*, Forum Mathematicum, vol. 21, n. 6, 981–999 (2009).
14. H. Doraiswamy, V. Natarajan, *Efficient algorithms for computing Reeb graphs*, Computational Geometry, vol. 42, 606–616 (2009).
15. H. Edelsbrunner, J. L. Harer, *Persistent homology—a survey*, Contemp. Math., vol. 453, 257–282 (2008).

16. H. Edelsbrunner, J. L. Harer, Computational Topology. An introduction, American Mathematical Society, Providence, RI, 2010.
17. P. Frosini, *G-invariant persistent homology*, Mathematical Methods in the Applied Sciences, vol. 38, n. 6, 1190–1199 (2015). DOI: 10.1002/mma.3139.
18. P. Frosini, *Measuring shapes by size functions*, Proc. of SPIE, Intelligent Robots and Computer Vision X: Algorithms and Techniques, Boston, MA 1607, 122–133 (1991).
19. P. Frosini, C. Landi, *Uniqueness of models in persistent homology: the case of curves*, Inverse Problems 27.12 (2011): 124005.
20. P. Frosini, M. Mulazzani, *Size homotopy groups for computation of natural size distances*, Bulletin of the Belgian Mathematical Society, vol. 6, n. 3, 455–464 (1999).
21. A. Hatcher, Algebraic Topology, Cambridge University Press, 2002.
22. S. Nelaturi, V. Shapiro, *Solving inverse configuration space problems by adaptive sampling*, Computer-Aided Design, vol. 45, 373–382 (2013).
23. J. Perea, J. L. Harer, *Sliding windows and persistence: an application of topological methods to signal analysis*, Foundations of Computational Mathematics (2013).
24. J. Reininghaus, N. Kotava, D. Gunther, J. Kasten, H. Hagen, I. Hotz, *A scale space based persistence measure for critical points in 2D scalar fields*, IEEE Transactions on Visualization and Computer Graphics, vol. 17, n. 12, 2045–2052 (2011).
25. C. Uras, A. Verri, *Computing size functions from edge maps*, International Journal of Computer Vision, vol. 23, n. 2, 169–183 (1997).
26. A. Verri, C. Uras, *Metric-topological approach to shape representation and recognition*, Image and Vision Computing, vol. 14, n. 3, 189–207 (1996).
27. A. Verri, C. Uras, P. Frosini, M. Ferri, *On the use of size functions for shape analysis*, Biological Cybernetics, vol. 70, 99–107 (1993).

A Proof of Proposition 7

Proof 1. The value $d_{\mathcal{F}}(F_1, F_2)$ is finite for every $F_1, F_2 \in \mathcal{F}$, because Φ is bounded. Indeed, a finite constant L exists such that $d_\infty(\varphi, \mathbf{0}) := \|\varphi\|_\infty \leq L$ for every $\varphi \in \Phi$. Hence $\|F_1(\varphi) - F_2(\varphi)\|_\infty \leq \|F_1(\varphi)\|_\infty + \|F_2(\varphi)\|_\infty \leq 2L$ for any $\varphi \in \Phi$ and any $F_1, F_2 \in \mathcal{F}$, since $F_1(\varphi), F_2(\varphi) \in \Phi$. This implies that $d_{\mathcal{F}}(F_1, F_2) \leq 2L < \infty$ for every $F_1, F_2 \in \mathcal{F}$.
2. $d_{\mathcal{F}}$ is obviously symmetrical.
3. The triangle inequality holds, since

$$\begin{aligned}
&d_{\mathcal{F}}(F_1, F_2) := \sup_{\varphi \in \Phi} \|F_1(\varphi) - F_2(\varphi)\|_\infty \leq \\
&\sup_{\varphi \in \Phi} \left(\|F_1(\varphi) - F_3(\varphi)\|_\infty + \|F_3(\varphi) - F_2(\varphi)\|_\infty\right) \leq \\
&\sup_{\varphi \in \Phi} \|F_1(\varphi) - F_3(\varphi)\|_\infty + \sup_{\varphi \in \Phi} \|F_3(\varphi) - F_2(\varphi)\|_\infty = \\
&d_{\mathcal{F}}(F_1, F_3) + d_{\mathcal{F}}(F_3, F_2)
\end{aligned}$$

for any $F_1, F_2, F_3 \in \mathcal{F}$.
4. The definition of $d_{\mathcal{F}}$ immediately implies that $d_{\mathcal{F}}(F, F) = 0$ for any $F \in \mathcal{F}$.
5. If $d_{\mathcal{F}}(F_1, F_2) = 0$, then the definition of $d_{\mathcal{F}}$ implies that $\|F_1(\varphi) - F_2(\varphi)\|_\infty = 0$ for every $\varphi \in \Phi$, and hence $F_1(\varphi) = F_2(\varphi)$ for every $\varphi \in \Phi$. Therefore $F_1 \equiv F_2$.
□

B Remark

If X and Y are two homeomorphic spaces and $h : Y \to X$ is a homeomorphism, then the persistent homology group with respect to the function $\varphi : X \to \mathbb{R}$ and the persistent homology group with respect to the function $\varphi \circ h : Y \to \mathbb{R}$ are isomorphic at each point (u, v) in the domain. The isomorphism between the two persistent homology groups is the one taking each homology class $[c = \sum_{i=1}^{r} a_i \cdot \sigma_i] \in PH_k^{\varphi}(u, v)$ to the homology class $[c' = \sum_{i=1}^{r} a_i \cdot (h^{-1} \circ \sigma_i)] \in PH_k^{\varphi \circ h}(u, v)$, where each σ_i is a singular simplex involved in the representation of the cycle c.

C Approximation of the non-expansive G_1-operator $F_{\hat{w},\hat{\gamma}}$

The operator $F_{\hat{w},\hat{c}}(\varphi)(x) := \sup_{r \in \mathbb{R}} \sum_{i=1}^{n} w_i \cdot \varphi(x + rc_i)$ can be approximated by substituting the supremum in its definition with a maximum for r belonging to a finite set.

In order to show this, first of all we observe that if $\hat{w} = \mathbf{0}$ then $F_{\hat{w},\hat{c}}$ is just the null operator, while if $\hat{c} = \mathbf{0}$ then $F_{\hat{w},\hat{c}}$ is just the operator taking each function φ to the function $\left(\sum_{i=1}^{n} w_i\right) \cdot \varphi$. Therefore, we can restrict ourselves to consider the case $\hat{w} \neq \mathbf{0}$, $\hat{c} \neq \mathbf{0}$.

Secondly, let us consider the function $\psi := \varphi \circ h$, where $h : \mathbb{R} \to \mathbb{R}$ is the linear homeomorphism taking 0 to -1 and 1 to 2 (i.e., $h(y) = 3y - 1$). Given that every function $\varphi \in \Phi_{ds}$ is assumed to be Lipschitz, with Lipschitz constant C, the function ψ is Lipschitz, with Lipschitz constant $3C$. The support of ψ is the interval $[0, 1]$. It is easy to check that $F_{\hat{w},\hat{c}}(\varphi)(x) = F_{\hat{w},\hat{\gamma}}(\psi)(y)$, where $y = \frac{1}{3}x + \frac{1}{3}$ and $\hat{\gamma} = (\gamma_1, \ldots, \gamma_n) := \frac{1}{3}\hat{c}$.

Furthermore, we can assume that $\gamma_p \neq \gamma_q$ for $p < q$. Indeed, if $\gamma_p = \gamma_q$ with $p < q$ we can consider the new vectors $\hat{w}' := (w'_1, \ldots, w'_{n-1})$, $\hat{\gamma}' := (\gamma'_1, \ldots, \gamma'_{n-1})$ obtained by setting for $1 \leq i \leq n - 1$

$$w'_i = \begin{cases} w_i, & \text{if } i < p \\ w_p + w_q, & \text{if } i = p \\ w_i, & \text{if } p < i < q \\ w_{i+1}, & \text{if } i \geq q \end{cases} \quad \text{and} \quad \gamma'_i = \begin{cases} \gamma_i, & \text{if } i < q \\ \gamma_{i+1}, & \text{if } i \geq q \end{cases}.$$

It is easy to check that $F_{\hat{w}',\hat{\gamma}'} = F_{\hat{w},\hat{\gamma}}$.

Hence, we can assume $\mu := \min\{|\gamma_p - \gamma_q| : p \neq q\} > 0$. Let us set $M := \max_i |\gamma_i|$.

We start by observing that if $|y| > \frac{M}{\mu} + 1$ the value $F_{\hat{w},\hat{\gamma}}(\psi)(y)$ is easily computable, because the condition $0 \leq y + r\gamma_p, y + r\gamma_q \leq 1$ cannot hold for $p \neq q$.

Indeed, if $y < -\frac{M}{\mu} - 1$, then $0 \leq y + r\gamma_i$ implies $|r\gamma_i| \geq r\gamma_i \geq -y > \frac{M+\mu}{\mu} > 0$, and hence $|r| > \frac{M+\mu}{\mu|\gamma_i|} \geq \frac{1}{\mu}$. If $y > \frac{M}{\mu} + 1$, then $y + r\gamma_i \leq 1$ implies $r\gamma_i \leq 1 - y < -\frac{M}{\mu} < 0$, and hence $|r\gamma_i| > \frac{M}{\mu}$, so that $|r| > \frac{M}{\mu|\gamma_i|} \geq \frac{1}{\mu}$. As a consequence, in both cases $|(y + r\gamma_p) - (y + r\gamma_q)| = |r| \cdot |\gamma_p - \gamma_q| \geq |r| \cdot \mu > 1$, for $p \neq q$. Therefore, the condition $0 \leq y + r\gamma_p, y + r\gamma_q \leq 1$ cannot hold, so that at most one of the points $y + r\gamma_p, y + r\gamma_q$ can belong to $[0, 1]$.

Now, let us consider the value $\sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i)$ as a function of $r \in \mathbb{R}$, under the assumption that $|y| > \frac{M}{\mu} + 1$.

When $r = 0$, for every index i we have that $y + r\gamma_i = y \notin [0, 1]$, so implying that $\sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i) = 0$ because the support of ψ is contained in $[0, 1]$.

When $r \neq 0$ we have that at most one of the points in the set $\{y + rc_1, \ldots, y + rc_n\}$ can belong to the interval $[0, 1]$. Moreover, for every index i such that $\gamma_i \neq 0$ and every $\eta \in [0, 1]$ exactly one value $r \neq 0$ exists, such that $y + rc_i = \eta$.

It follows that $\sup_{r\in\mathbb{R}} \sum_{i=1}^{n} w_i \cdot \psi\,(y + rc_i) = \max\{0, \max(w_1 \cdot \psi), \ldots, \max(w_n \cdot \psi)\}$.

In conclusion, if $|y| > \frac{M}{\mu} + 1$, $F_{\hat{w},\hat{\gamma}}(\psi)(y) = \max\{0, \max(w_1 \cdot \psi), \ldots, \max(w_n \cdot \psi)\}$.

It follows that if $|y| > \frac{M}{\mu} + 1$ we can easily approximate $F_{\hat{w},\hat{\gamma}}(\psi)(y)$ because the values $\max(w_i \cdot \psi)$ can be approximated with arbitrary precision. Indeed, we know that the function ψ is Lipschitz, with Lipschitz constant $3C$. This implies that we can approximate $\max(w_i \cdot \psi)$ by computing the value $\max_s(w_i \cdot \psi\,(r_s))$, where $\{r_1, \ldots, r_m\}$ is a sufficiently dense finite subset of the interval $[0, 1]$.

Let us now consider the case $|y| \leq \frac{M}{\mu} + 1$. In this case, if $\gamma_i \neq 0$ and $|r| > \frac{M+2\mu}{\mu|\gamma_i|}$ then $\psi(y + r\gamma_i) = 0$. Indeed, the support of ψ is contained in $[0, 1]$ and $|y + r\gamma_i| \geq \big||r\gamma_i| - |y|\big| = |r\gamma_i| - |y| > 1$ because $|r\gamma_i| > \frac{M}{\mu} + 2$ and $|y| \leq \frac{M}{\mu} + 1$.

Setting $R := \frac{M+2\mu}{\mu\cdot\min\{|\gamma_i|:\gamma_i\neq 0\}}$, it follows that $|y| \leq \frac{M}{\mu} + 1$ implies

$$\sup_{|r|>R} \sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i) = \begin{cases} w_j \cdot \psi(y), & \text{if an index } j \text{ exists s.t. } \gamma_j = 0 \\ 0, & \text{if } \gamma_i \neq 0 \text{ for every index } i \end{cases}$$

so that

$$\begin{aligned}
&\sup_{r\in\mathbb{R}} \sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i) = \\
&\max\left\{ \sup_{|r|\leq R} \sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i)\,,\ \sup_{|r|>R} \sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i) \right\} = \\
&\begin{cases} \max\left\{ \sup_{|r|\leq R} \sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i)\,, w_j \cdot \psi(y) \right\}, & \text{if an index } j \text{ exists s.t. } \gamma_j = 0 \\ \max\left\{ \sup_{|r|\leq R} \sum_{i=1}^{n} w_i \cdot \psi\,(y + r\gamma_i)\,, 0 \right\}, & \text{if } \gamma_i \neq 0 \text{ for every index } i. \end{cases}
\end{aligned}$$

Now, we have

$$\begin{aligned}
&\left|\sum_{i=1}^{n} w_i \cdot \psi\left(y + r\gamma_i\right) - \sum_{i=1}^{n} w_i \cdot \psi\left(y + r'\gamma_i\right)\right| = \\
&\left|\sum_{i=1}^{n} w_i \cdot \left(\psi\left(y + r\gamma_i\right) - \psi\left(y + r'\gamma_i\right)\right)\right| \leq \\
&\sum_{i=1}^{n} |w_i| \cdot \left|\psi\left(y + r\gamma_i\right) - \psi\left(y + r'\gamma_i\right)\right| \leq \\
&\sum_{i=1}^{n} |w_i| \cdot 3C \cdot |(r - r') \cdot \gamma_i| = \\
&\sum_{i=1}^{n} |w_i| \cdot 3C \cdot |r - r'| \cdot |\gamma_i| \leq \\
&\sum_{i=1}^{n} |w_i| \cdot 3C \cdot |r - r'| \cdot M = \\
&3C \cdot M \cdot |r - r'|
\end{aligned}$$

because $\sum_{i=1}^{n} |w_i| = 1$ and ψ is Lipschitz, with Lipschitz constant $3C$. Hence we can approximate $\sup_{|r|\leq R} \sum_{i=1}^{n} w_i \cdot \psi\left(y + r\gamma_i\right)$ by computing the value $\max_s \sum_{i=1}^{n} w_i \cdot \psi\left(y + r_s\gamma_i\right)$, where $\{r_1, \ldots, r_m\}$ is a sufficiently dense finite subset of the interval $[-R, R]$. It follows that we can easily approximate $F_{\hat{w},\hat{\gamma}}(\psi)(y)$ also in the case $|y| \leq \frac{M}{\mu} + 1$.

Therefore our statement is proven.

D Definition of the first five operators used in Section 5

The formal definition of our operators is:

$$F_\beta(\varphi)(\boldsymbol{x}) := \int_B \varphi(\boldsymbol{x} - \boldsymbol{y}) \cdot \beta\left(\|\boldsymbol{y}\|_2\right)\ d\boldsymbol{y}$$

where β is an integrable function defined on a ball $B \subset \mathbb{R}^2$ such that $\int_B \left|\beta\left(\|\boldsymbol{y}\|_2\right)\right|\ d\boldsymbol{y} \leq 1$ (here, $\|\boldsymbol{y}\|_2$ denotes the Euclidean norm of the vector $\boldsymbol{y}$). This condition is necessary in the proof of non-expansiveness, and β can be considered as a kernel function. We used the following four kernel functions:

1. $\beta(t) = \begin{cases} 16/\pi, & \text{if } 0 \leq t \leq 1/4 \\ 0, & \text{if } t < 0 \vee t > 1/4 \end{cases}.$
2. $\beta(t) = \begin{cases} 16/\pi, & \text{if } 0 \leq t < 1/8 \\ -16/\pi, & \text{if } 1/8 \leq t \leq 1/4 \\ 0, & \text{if } t < 0 \vee t > 1/4 \end{cases}.$
3. $\beta(t) = \begin{cases} 16/\pi, & \text{if } 0 \leq t < 1/16 \\ -16/\pi, & \text{if } 1/16 \leq t < 1/8 \\ 16/\pi, & \text{if } 1/8 \leq t < 3/16 \\ -16/\pi, & \text{if } 3/16 \leq t \leq 1/4 \\ 0, & \text{if } t < 0 \vee t > 1/4 \end{cases}.$
4. $\beta(t) = \begin{cases} 4/\pi, & \text{if } 0 \leq t < 1/4 \\ -4/\pi, & \text{if } 1/4 \leq t \leq 1/2 \\ 0, & \text{if } t < 0 \vee t > 1/2 \end{cases}.$