SYMMETRIES OF DATA SETS AND FUNCTORIALITY OF PERSISTENT HOMOLOGY

WOJCIECH CHACHÓLSKI, ALESSANDRO DE GREGORIO, NICOLA QUERCIOLI AND FRANCESCA TOMBARI

Abstract. The aim of this article is to describe a new perspective on functoriality of persistent homology and explain its intrinsic symmetry that is often overlooked. A data set for us is a finite collection of functions, called measurements, with a finite domain. Such a data set might contain internal symmetries which are effectively captured by the action of a set of the domain endomorphisms. Different choices of the set of endomorphisms encode different symmetries of the data set. We describe various category structures on such enriched data sets and prove some of their properties such as decompositions and morphism formations. We also describe a data structure, based on coloured directed graphs, which is convenient to encode the mentioned enrichment. We show that persistent homology preserves only some aspects of these collection of enriched data sets however not all. In other words persistent homology is not a functor on the entire category of enriched data sets. Nevertheless we show that persistent homology is functorial locally. We use the concept of set equivariant operator (SEO) to capture some of the information missed by persistent homology. Moreover, we provide examples and give ways to construct such SEOs.

1. Introduction

In this article we give an answer to the question: what is persistent homology a functor of?

We consider data sets given by finite sets of functions on a finite set X with real values. There are several important consequences of data sets having this form. For example, they endow X with a pseudometric, enabling us to extract non-trivial homological information in form of persistent homology, one of the key invariants studied in Topological Data Analysis. A single measurement does not contain any higher non-trivial homological information. Sets of measurements, however, do. Thus it is essential that measurements, on a given set X, are grouped together to form various data sets. In this case persistent homology becomes a non-expansive (Lipschitz continuous with Lipschitz constant less or equal to 1) function $\mathrm{PH}_d^{\Phi}\colon \Phi \to \mathrm{Tame}([0,\infty)\times \mathbf{R}, \mathrm{Vect})$, assigning to each measurement in the data set Φ a tame vector space parametrized by $[0,\infty)\times \mathbf{R}$. It is important to notice that the choice of a set of measurements on X affects the pseudometric defined on it. One can use this fact to change the metric on X in order to extract more meaningful information from persistent homology. For example consider X to be a finite sample of points on a circle. If Φ consists of only one function given by the $x-$coordinate, then the persistent homology of this measurement is trivial in degrees greater than 0. If we enlarge the data set by adding to the $x-$coordinate the function given by precomposing x with rotation by 90 degrees, then the persistent homology of the function x with respect to this bigger data set gains a non-trivial homology in degree 1. This illustrates how our knowledge of an object is affected by the number and the type of measurements done on it. Furthermore in this example we gain additional information by enlarging the set of measurements through the action of some of the endomorphisms of X on the existing measurements. We can then take advantage of these actions to inject geometrical features of our choice on a given data set. For exhibiting and extracting interesting homological features of data sets, such actions are therefore important.

A data set Φ is naturally equipped with an action of the monoid of its operations $\mathrm{End}_\Phi(X)$, which are endomorphisms of X preserving Φ by composition on the right. We introduce the notion of Grothendieck graph to encode this action and represent it as graph. Persistent homology turns out to be a functor indexed by this graph, rather than simply a function. Thus not only persistent homology can be assigned to individual measurements in a data set, but operations can be used to compare persistent homologies of different sets of measurements. That is what we call local functorial properties of persistent homology.

Persistent homology also has certain global functorial properties. There are various ways of representing data in the form of sets of measurements. We might choose different units or different parametrizations of the domain of the admissible measurements, or we might need to focus only on certain operations, such as rotations. Furthermore, the same measurements might be part of different data sets. These are some of the reasons why it is essential to be able to compare data sets equipped with different structures. Thus, instead of studying a data set by itself, we introduce the notion of perception pair encoding intrinsic symmetries of the data set as action of a subset $M \subset \mathrm{End}_\Phi(X)$. A SEO (set equivariant operator) between two perception pairs (Φ, M) and (Ψ, N) is a pair (α, T) consisting of a map $T : M \to N$ and an equivariant (with respect to T) function $\alpha : \Phi \to \Psi$. The use of this kind of operators for the comparison of perception pairs of data sets has been inspired by [1, 2] , where GENEOs (Group Equivariant Non-Expansive Operators) are introduced and used for applications to neural networks. The novelty in our approach is to study equivariance with respect to a set, instead of a group. This weakening of hypothesis is supported by this need in applications. In fact, the set of operations that one might want to perform on a data set likely does not form a group in general; think, for example, about a figure that can be rotated only up to a certain angle in order not to change its interpretation. As a first step towards the understanding of such set equivariant operators, we illustrate some cardinal examples and show techniques to construct them, starting from a collection of generators.

Parallely, SEOs play a fundamental role in describing the functoriality of persistent homology. In particular, if a SEO is geometric, then there is a comparison map between persistent homologies of the perception pairs connected by the SEO. However, if a SEO is not geometric, such as the change of units SEO, there is no direct comparison of persistent homologies of the involved perception pairs. Such SEOs therefore exhibit diverse homological features of data sets enhancing the analysis. This suggests complementarity of such operators and persistent homology for a geometric analysis of a data set. Consider the change of unit as an example. If we think about measurements as real-valued functions, the change of unit, in general, is the SEO obtained by composing them with a given real valued function defined on the real numbers. Multiplication by -1 is an example of such SEO. It has the effect of turning the sub-level sets persistent homology of a measurement into its super-level sets persistent homology, leading, in general, to a completely different information about the data set. The outcome consists of two different points of view on the same object, that are not functorially comparable, but together may enhance the accuracy of the analysis of the object of interest.

2. Data sets

In this article, a data set is given by a finite set of real valued functions on a finite set X, also called measurements:

$$\Phi = \{\phi_i \colon X \to \mathbf{R} \mid i = 1, \cdots, m\}.$$

The **domain** of the data set Φ, $\text{dom}(\Phi)$, is the set X, which is the domain of all the functions in Φ. Without additional hypothesis, the collection of such data sets is just a subcategory of the category of sets. Our purpose is to add more intricate, but meaningful, structure to this setting based on the metric that the measurements of the data set induce on the set X.

In this most primitive landscape, however, we can already perform products and coproducts. Let $\phi \colon X \to \mathbf{R}$ and $\psi \colon Y \to \mathbf{R}$ be functions. Define $\phi + \psi \colon X \coprod Y \to \mathbf{R}$ to be the function that maps x in X to $\phi(x)$ and y in Y to $\psi(y)$. The **coproduct** of two data sets Φ and Ψ, denoted by $\Phi \coprod \Psi$, is defined to be the data set given by the measurements $\{\phi + 0 \mid \phi \in \Phi\} \cup \{0 + \psi \mid \psi \in \Psi\}$ on $X \coprod Y$. Their **product**, denoted by $\Phi \times \Psi$, is defined to be the data set given by the measurements $\{\phi + \psi \mid \phi \in \Phi \text{ and } \psi \in \Psi\}$ on $X \coprod Y$. The functions:

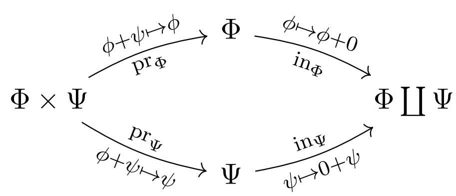

satisfy the following universal properties, which justify the names coproduct and product:

- for any data set Π, and any two functions $\alpha \colon \Phi \to \Pi$ and $\beta \colon \Psi \to \Pi$, there is a unique function $\mu \colon \Phi \coprod \Psi \to \Pi$ for which $\mu \, \text{in}_\Phi = \alpha$ and $\mu \, \text{in}_\Psi = \beta$;
- for any data set Π, and any two functions $\alpha \colon \Pi \to \Phi$ and $\beta \colon \Pi \to \Psi$, there is a unique function $\mu \colon \Pi \to \Phi \times \Psi$ for which $\text{pr}_\Phi \mu = \alpha$ and $\text{pr}_\Psi \mu = \beta$.

Let $f \colon \mathbf{R} \to \mathbf{R}$ be a function. By composing with f, a data set Φ is transformed into a new data set $f\Phi := \{f\phi \mid \phi \in \Phi\}$. This operation is called **change of units** along f. The symbol $f- \colon \Phi \to f\Phi$ denotes the function mapping ϕ to $f\phi$. As the following example shows, $f-$ may not be a functor. Let $f \colon \mathbf{R} \to \mathbf{R}$ map $\{r \in \mathbf{R} \mid r < 0\}$ to -1 and $\{r \in \mathbf{R} \mid r \geq 0\}$ to 1. Consider $X = \{x_1, x_2\}$, two data sets $\{1, 2\}$ and $\{-1, 1\}$ given by the constant functions $-1, 1, 2 \colon X \to \mathbf{R}$, and a function $\alpha \colon \{1, 2\} \to \{-1, 1\}$ mapping 1 to -1

and 2 to 1. Then $f\{1,2\} = \{1\}$ and $f-\colon \{-1,1\} \to f\{-1,1\}$ is the identity. Thus there is no function $\beta\colon f\{1,2\} \to f\{-1,1\}$ making the following diagram commutative:

$$\begin{array}{ccc} \{1,2\} & \xrightarrow{\ f-\ } & f\{1,2\} = \{1\} \\ {\scriptstyle\alpha}\downarrow & & \not\downarrow \\ \{-1,1\} & \xrightarrow{f-=\mathrm{id}} & f\{-1,1\} = \{-1,1\} \end{array}$$

Consequently, the association $f-$ is not a functor. If f is invertible, then $f-\colon \Phi \to f\Phi$ is a bijection whose inverse is given by $f^{-1}-$. Thus, $f-$ is a functor and $\beta = (f-)\alpha(f^{-1}-)$. Changing the units along any function preserves products and coproducts i.e., $f(\Phi \coprod \Psi)$ is isomorphic to $f(\Phi) \coprod f(\Psi)$, and $f(\Phi \times \Psi)$ is isomorphic to $f(\Phi) \times f(\Psi)$. A similar reasoning is used in [7] to study brain data, in order to obtain results that are invariant under transformations given by change of units with invertible functions, and in [8] to study metric spaces that are isometric up to a rescaling of the distance functions.

Let Φ be a data set with domain X. By composing a function $f\colon Y \to X$ with the measurements in Φ, we obtain a new data set $\Phi f := \{\phi f \mid \phi \in \Phi\}$ with domain Y. This operation is called **domain change** along f. The symbol $-f\colon \Phi \to \Phi f$ denotes the function that maps ϕ to ϕf. Let $f_1\colon Z_1 \to X$ and $f_2\colon Z_2 \to Y$ be two functions. Their coproduct is $f_1 \coprod f_2\colon Z_1 \coprod Z_2 \to X \coprod Y$. For any datasets Φ and Ψ with $\mathrm{dom}(\Phi) = X$ and $\mathrm{dom}(\Psi) = Y$, the following equalities hold:

$$(\Phi \coprod \Psi)(f_1 \coprod f_2) = \Phi f_1 \coprod \Psi f_2, \qquad (\Phi \times \Psi)(f_1 \coprod f_2) = \Phi f_1 \times \Psi f_2.$$

3. Metrics and persistent homology

We can think about a data set Φ as a subset $\Phi \subset \mathbf{R}^{|X|}$. Via this inclusion Φ inherits a metric induced by the infinity norm $\|v\|_\infty = \max\{|v_i|\}$ on $\mathbf{R}^{|X|}$. We use the symbol $\|\phi - \psi\|_\infty$ to denote the distance between ϕ and ψ in Φ. The considered data sets are not just sets anymore but metric spaces. Therefore non-expansive (1-Lipschitz) functions between data sets play a special role. For example, let $f\colon \mathbf{R} \to \mathbf{R}$ be a function. If f is non-expansive, then so is the change of units along f, $f-\colon \Phi \to f\Phi$, that maps ϕ to $f\phi$. The domain change $-h\colon \Phi \to \Phi h$ is non-expansive along any h. Non-expansiveness is an important assumption to prove some stability results in [1] and it is also reasonable in applications, since it is important that these functions between data sets do not alter the information too much.

By taking all the measurements of Φ together, we can form a function $[\phi_1 \cdots \phi_m]\colon X \to \mathbf{R}^m$. Via this function, X inherits a pseudometric d_Φ induced by the infinity norm on $\mathbf{R}^m$. Explicitly, $d_\Phi(x,y) := \max_{1\le i\le m}|\phi_i(x) - \phi_i(y)|$. This metric plays a fundamental role as it permits us to extract persistent homologies (see [3, 6]). In this article, the **persistent homology** of a data set Φ with coefficients in a field and in a given degree d assigns a vector space $\mathrm{PH}_d^\Phi(\phi)_{r,s}$ to each measurement ϕ in Φ, for every (r,s) in $[0,\infty) \times \mathbf{R}$, and it is defined as:

$$\mathrm{PH}_d^\Phi(\phi)_{r,s} := H_d\left(\mathrm{VR}_r(\phi \le s, d_\Phi)\right), \text{ where:}$$

- $\phi \le s := \phi^{-1}(-\infty, s]$;

- $\mathrm{VR}_r(\phi \leq s, d_\Phi)$ is the **Vietoris-Rips** complex whose simplices are given by the subsets $\sigma \subset (\phi \leq s)$ of diameter not exceeding r with respect to d_Φ;
- H_d is the homology in degree d with coefficients in a given field.

If $s \leq s'$ and $r \leq r'$, then $(\phi \leq s) \subset (\phi \leq s')$ and therefore $\mathrm{VR}_r(\phi \leq s) \subset \mathrm{VR}_{r'}(\phi \leq s')$. The linear function induced on homology by this inclusion is denoted by:

$$\mathrm{PH}_d^\Phi(\phi)_{(r,s)\leq(r',s')} \colon \mathrm{PH}_d^\Phi(\phi)_{r,s} \to \mathrm{PH}_d^\Phi(\phi)_{r',s'}.$$

These functions form a functor $\mathrm{PH}_d^\Phi(\phi)$ indexed by the poset $[0,\infty) \times \mathbf{R}$ with values in the category of vector spaces. Since X is finite, $\mathrm{PH}_d^\Phi(\phi)$ is **tame** (see [11]). This means that the values of $\mathrm{PH}_d^\Phi(\phi)$ are finite dimensional, and there are two finite sequences $0 = r_0 < r_1 < \cdots < r_m$ in $[0,\infty)$ and $s_0 < s_1 < \cdots < s_l = \infty$ in $\mathbf{R}$ such that $\mathrm{PH}_d^\Phi(\phi)$, restricted to subposets of the form $[r_i, r_{i+1}) \times (-\infty, s_0) \subset [0,\infty) \times \mathbf{R}$ and $[r_i, r_{i+1}) \times [s_j, s_{j+1}) \subset [0,\infty) \times \mathbf{R}$, is constant. The category of such functors is denoted by $\mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect})$. Thus a data set Φ leads to a function assigning to each measurement ϕ its persistent homology in a given degree:

$$\mathrm{PH}_d^\Phi \colon \Phi \to \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect}).$$

Next, we recall a definition of the interleaving metric in the direction of the vector $(0,1)$ on $\mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect})$ (see [9]). Let P and Q be in $\mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect})$.

- P and Q are ε-**interleaved** if, for all (r,s) in $[0,\infty) \times \mathbf{R}$, there are linear functions $f_{r,s} \colon P_{r,s} \to Q_{r,s+\varepsilon}$ and $g_{r,s} \colon Q_{r,s} \to P_{r,s+\varepsilon}$ making the following diagram commute

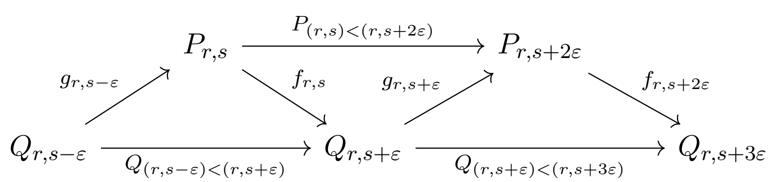

- $d_{\bowtie}(P,Q) := \inf\{\varepsilon \in [0,\infty) \mid P \text{ and } Q \text{ are } \varepsilon\text{-interleaved}\}$.

The function $P, Q \mapsto d_{\bowtie}(P,Q)$ is an extended (∞ is allowed) metric on the set $\mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect})$ called **interleaving metric** in the direction of the vector $(0,1)$.

3.1. **Proposition.** *The function* $\mathrm{PH}_d^\Phi \colon \Phi \to \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect})$ *is non-expansive if the set* Φ *is equipped with the metric* $\|\phi - \psi\|_\infty$ *and the set* $\mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect})$ *is equipped with the interleaving metric in the direction of the vector* $(0,1)$.

Proof. Let $\phi, \psi \colon X \to \mathbf{R}$ be measurements in Φ and $\varepsilon = \|\phi - \psi\|_\infty$. For every s in $\mathbf{R}$, the sublevel set $\phi \leq s$ is a subset of $\psi \leq s + \varepsilon$, and $\psi \leq s$ is a subset of $\phi \leq s + \varepsilon$. This translates into the inclusions

$$\mathrm{VR}_r(\phi \leq s, d_\Phi) \subset \mathrm{VR}_r(\psi \leq s+\varepsilon, d_\Phi) \qquad \mathrm{VR}_r(\psi \leq s, d_\Phi) \subset \mathrm{VR}_r(\phi \leq s+\varepsilon, d_\Phi)$$

leading to the functions:

$$f_{s,r} \colon \mathrm{PH}_d^\Phi(\phi)_{r,s} \to \mathrm{PH}_d^\Phi(\psi)_{r,s+\varepsilon} \qquad g_{s,r} \colon \mathrm{PH}_d^\Phi(\psi)_{r,s} \to \mathrm{PH}_d^\Phi(\phi)_{r,s+\varepsilon}.$$

These functions provide a ε-interleaving between $\mathrm{PH}_d^\Phi(\phi)$ and $\mathrm{PH}_d^\Phi(\psi)$, implying $\|\phi - \psi\|_\infty \geq d_{\bowtie}(\mathrm{PH}_d^\Phi(\phi), \mathrm{PH}_d^\Phi(\psi))$. □

A measurement $\phi\colon X \to \mathbf{R}$ can be part of many data sets and its persistent homology depends on what data set this function is part of. The reason for this is that persistent homology depends on the metric d_Φ, but the metric, in our setting, depends on the data set. For example, let $X = \{x_1, x_2, x_3, x_4\}$ and $\phi, \psi\colon X \to \mathbf{R}$ be measurements defined as follows:

$\phi(x_1) = -1$	$\phi(x_2) = \phi(x_3) = 0$	$\phi(x_4) = 1$
$\psi(x_3) = -1$	$\psi(x_1) = \psi(x_4) = 0$	$\psi(x_2) = 1$

The measurement ϕ is part of two data sets $\Phi = \{\phi\}$ and $\Psi = \{\phi, \psi\}$. The induced pseudometrics d_Φ and d_Ψ on X can be depicted by the following diagrams where the continuous, dashed, and dotted lines indicate distance 0, 1 and 2 respectively:

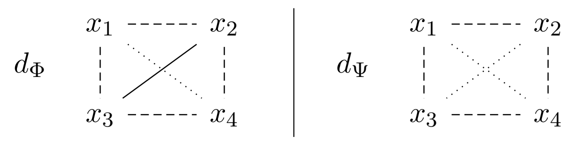

In this case $\mathrm{PH}_1^\Phi(\phi)_{r,s} = 0$ for all r and s, however:

$$\dim \mathrm{PH}_1^\Psi(\phi)_{r,s} = \begin{cases} 1 & \text{if } 1 \leq s \text{ and } 1 \leq r < 2 \\ 0 & \text{otherwise} \end{cases}$$

To understand persistent homology, it is therefore paramount to understand how it changes when data sets change and here functoriality plays an essential role.

Let Φ and Ψ be data sets consisting of measurements on X and Y respectively. A function $\alpha\colon \Phi \to \Psi$ is called **geometric** if there is a function $f\colon Y \to X$, called a **realization** of α, making the following diagram commute for every ϕ in Φ:

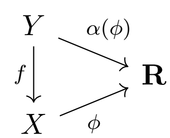

For example, $-f\colon \Phi \to \Phi f$ is geometric, as it is realized by f.

The commutativity of the triangle above has two consequences. First, f is non-expansive with respect to the pseudometrics d_Φ on X and d_Ψ on Y. Second, for s in $\mathbf{R}$ and ϕ in Φ, the subset $(\alpha(\phi) \leq s) \subset Y$ is mapped via f into $(\phi \leq s) \subset X$, i.e., the following diagram commutes

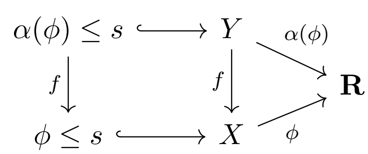

The realization f induces therefore a map of Vietoris-Rips complexes and their homologies:

$$f_{r,s}\colon \mathrm{VR}_r(\alpha(\phi) \leq s, d_\Psi) \to \mathrm{VR}_r(\phi \leq s, d_\Phi);$$

$$\begin{array}{ccc} \mathrm{PH}_d^{\Psi}(\alpha(\phi))_{r,s} & & \mathrm{PH}_d^{\Phi}(\phi)_{r,s} \\ \| & & \| \\ H_d\left(\mathrm{VR}_r(\alpha(\phi)\leq s, d_\Psi)\right) & \xrightarrow{H_d(f_{r,s})} & H_d\left(\mathrm{VR}_r(\phi \leq s, d_\Phi)\right). \end{array}$$

If $f, f'\colon Y \to X$ are two realizations of α, then for y in Y, $d_\Phi(f(y), f'(y)) = 0$, hence they are points of the same simplex in the Vietoris-Rips complex, implying that $f_{r,s}$ and $f'_{r,s}$ are homotopic for all r and s. Consequently, $H_d(f_{r,s}) = H_d(f'_{r,s})$. The linear function $H_d(f_{r,s})$ depends therefore only on α and it is independent on the choice of its realization f. We denote this function by

$$\mathrm{PH}_d^{\alpha}(\phi)_{r,s}\colon \mathrm{PH}_d^{\Psi}(\alpha(\phi))_{r,s} \to \mathrm{PH}_d^{\Phi}(\phi)_{r,s}.$$

These functions are natural in r and s and induce a morphism in the category $\mathrm{Tame}([0,\infty)\times \mathbf{R}, \mathrm{Vect})$ between persistent homologies:

$$\mathrm{PH}_d^{\alpha}(\phi)\colon \mathrm{PH}_d^{\Psi}(\alpha(\phi)) \to \mathrm{PH}_d^{\Phi}(\phi).$$

If $\alpha\colon \Phi \to \Psi$ and $\beta\colon \Psi \to \Xi$ are geometric functions realized by $f\colon Y \to X$ and $g\colon Z \to Y$, then the composition $\beta\alpha\colon \Phi \to \Xi$ is also geometric, and realized by the composition $fg\colon Z \to X$. Consequently, for every measurement ϕ in Φ, $\mathrm{PH}_d^{\beta\alpha}(\phi) = \mathrm{PH}_d^{\alpha}(\phi)\mathrm{PH}_d^{\beta}(\alpha(\phi))$, assuring the commutativity of the diagram

$$\mathrm{PH}_d^{\Xi}(\beta\alpha(\phi)) \xrightarrow{\mathrm{PH}_d^{\beta}(\alpha(\phi))} \mathrm{PH}_d^{\Psi}(\alpha(\phi)) \xrightarrow{\mathrm{PH}_d^{\alpha}(\phi)} \mathrm{PH}_d^{\Phi}(\phi), \qquad \mathrm{PH}_d^{\Xi}(\beta\alpha(\phi)) \xrightarrow{\mathrm{PH}_d^{\beta\alpha}(\phi)} \mathrm{PH}_d^{\Phi}(\phi)$$

For any $\alpha\colon \Phi \to \Psi$, taking persistent homology leads to two functions on Φ:

$$\Phi \xrightarrow{\mathrm{PH}_d^{\Phi}} \mathrm{Tame}([0,\infty)\times \mathbf{R}, \mathrm{Vect}), \qquad \Phi \xrightarrow{\alpha} \Psi \xrightarrow{\mathrm{PH}_d^{\Psi}} \mathrm{Tame}([0,\infty)\times \mathbf{R}, \mathrm{Vect})$$

These functions rarely coincide. However, when α is geometric, we can use the morphisms $\mathrm{PH}_d^{\alpha}(\phi)\colon \mathrm{PH}_d^{\Psi}(\alpha(\phi)) \to \mathrm{PH}_d^{\Phi}(\phi)$ to compare the values of these two functions on Φ. For non-geometric α, we are not equipped with such comparison morphisms and there is no reason for such a comparison to even exist. For example, consider the change of unit along the function $f\colon \mathbf{R} \to \mathbf{R}$, $f(x) := -x$. Then $f-\colon \Phi \to f\Phi$ is an isomorphism. In this case

$$\mathrm{PH}_d^{\Phi}(\phi)_{r,s} := H_d\left(\mathrm{VR}_r(\phi \leq s, d_\Phi)\right) \;\Big|\; (f-)\mathrm{PH}_d^{f\Phi}(\phi) = H_d\left(\mathrm{VR}_r(\phi \geq -s, d_\Phi)\right).$$

Thus PH_d^{Φ} encodes information about sub-level sets of the measurements in Φ and $(f-)\mathrm{PH}_d^{f\Phi}$ encodes information about super-level sets of the measurements. These persistent homologies encode therefore complementary information, analogously to the so called extended persistence (see [4, 10]).

4. Actions

To describe symmetries of a data set Φ with domain X, we consider operations on X that convert measurements into measurements. By definition a **Φ-operation** is a function $g\colon X \to X$ such that, for every measurement ϕ in Φ, the composition ϕg also belongs to Φ. If $g\colon X \to X$ is such an operation, then, for all ϕ and ψ in Φ,

$$\|\phi - \psi\|_\infty = \max_{x\in X}|\phi(x) - \psi(x)| \geq \max_{x\in \mathrm{im}(g)}|\phi(x) - \psi(x)| = \|\phi g - \psi g\|_\infty.$$

Thus the function $-g\colon \Phi \to \Phi$ that maps ϕ to ϕg is non-expansive.

The composition of Φ-operations is again a Φ-operation, and the identity function id_X is also a Φ-operation. In this way the set of Φ-operations with the composition becomes a unitary monoid, called the **structure monoid** of Φ, and denoted by:

$$\mathrm{End}_\Phi(X) = \{g : X \to X \mid \phi g \in \Phi \text{ for every } \phi \in \Phi\} \subset \mathrm{End}(X).$$

A Φ-operation g is invertible if there is a Φ-operation h such that $gh = hg = \mathrm{id}_X$. Since Φ is finite, a Φ-operation is invertible if and only if it is a bijection. The collection of invertible Φ-operations is denoted by

$$\mathrm{Aut}_\Phi(X) = \{g : X \to X \mid \ g \text{ is a bijection, and } \phi g \in \Phi \text{ for every } \phi \in \Phi\}.$$

With the composition, $\mathrm{Aut}_\Phi(X)$ becomes a group for which the inclusion $\mathrm{Aut}_\Phi(X) \subset \mathrm{End}_\Phi(X)$ is a monoid homomorphism.

From now on a data set Φ will be equipped with an associative right action of the monoid $\mathrm{End}_\Phi(X)$:

$$\Phi \times \mathrm{End}_\Phi(X) \to \Phi, \qquad (\phi, g) \mapsto \phi g.$$

A **perception pair** of Φ is a choice of a subset $M \subset \mathrm{End}_\Phi(X)$ (not necessarily a submonoid) encoding the symmetries of the data set Φ induced by this action. A perception pair, or M-perception pair or M-action, is thus denoted by a pair (Φ, M). The pair $(\Phi, \mathrm{End}_\Phi(X))$ is an example of a perception pair called universal. Every choice of an M-action on Φ encodes certain symmetries of Φ. Different choices of M can encode different symmetries of the data set. This flexibility is important in applications. For example, given a data sets that represent images, we might want to focus on rotational symmetries, so we may use an appropriate action on the data set to inject the corresponding geometry.

A perception pair (Φ, M) is called a **monoid perception pair** if $M \subset \mathrm{End}_\Phi$ is a submonoid, containing the identity element. If (Φ, M) is a perception pair, we use the symbol $(\Phi, \langle M\rangle)$ to denote the monoid perception pair where $\langle M\rangle \subset \mathrm{End}_\Phi(X)$ is the submonoid generated by M. If a submonoid $M \subset \mathrm{End}_\Phi(X)$ is also a group, then (Φ, M) is called a **group perception pair**. The perception pair $(\Phi, \mathrm{Aut}_\Phi(X))$ is an example of a group perception pair called universal. A perception pair (Φ, M) for which any element g in M is a bijection is called **group-like perception pair**. For group like perception pairs (Φ, M) the finiteness of X implies that the monoid $\langle M\rangle$ is in fact a subgroup of $\mathrm{Aut}_\Phi(X)$. Thus any group-like perception pair (Φ, M) leads to a group perception pair $(\Phi, \langle M\rangle)$.

Let (Φ, M) be a perception pair. For a subset $\Omega \subset \Phi$, the symbol ΩM denotes the set of all measurements in Φ which either belong to Ω or are of the form $\omega g_1 \cdots g_k$, for some ω in Ω and some sequence of elements $g_1, \ldots g_k$ in M. If $\Omega M = \Phi$, then Ω is said to **generate** the perception pair (Φ, M). In the case (Φ, M) is a monoid perception pair, then any element in ΩM is of the form ωg for some ω in Ω and g in M. Note that $\Omega M = \Omega\langle M\rangle$ for every perception pair (Φ, M).

If ψ belongs to $\phi M := \{\phi\}M$, then ψ is said to be a **deformation** of ϕ. If (Φ, M) is a group perception pair, then the relation of being a deformation is an equivalence relation. For a general perception pair, however, being a deformation can fail to be even a symmetric relation. Two measurements in Φ are said to be **connected** if they are related by the equivalence relation generated by the relation of being a deformation: ϕ and ψ are connected if either ϕ is a deformation of ψ or ψ is a deformation of ϕ. The symbol Φ/M denotes the partition of Φ induced by this equivalence relation generated by the deformation of the elements of Φ. We refer to Φ/M as the **quotient** of the perception pair (Φ, M). The partitions Φ/M and $\Phi/\langle M\rangle$ coincide. If (Φ, M) is a group perception pair, then Φ/M coincide with the orbit partition of the usual group action of M on Φ. Let (Φ, M) be a perception pair. For a measurement ψ in Φ, the symbol $[\psi]$ denotes the equivalence class in Φ/M containing ψ. Explicitly, $[\psi]$ is the subset of Φ consisting of all the measurements connected to ψ.

A perception pair (Φ, M) is called **transitive** if all the elements in Φ are connected to each other. For example, let M be a submonoid of $\mathrm{End}(X)$. For a given function $\phi\colon X \to \mathbf{R}$, define a data set $\phi M := \{\phi g \mid g \in M\}$ to consist of all functions of the form $x \mapsto \phi(g(x))$ for all g in M. Then every $g\colon X \to X$ in M is a ϕM-operation. The obtained perception pair $(\phi M, M)$ is transitive. Any transitive group perception pair is of such form. For all measurements ϕ in any perception pair (Φ, M), the perception pair $([\phi], M)$, with $M \subset \mathrm{End}_{[\psi]}(X)$, is transitive. Any transitive perception pair is of this form.

Let (Φ, M) be an perception pair. A subset $\Omega \subset \Phi$ is called **independent** if no element in Ω is a deformation of any other element in Ω, explicitly: $\omega \notin \omega' M$ for all $\omega \neq \omega'$ in Ω.

A **basis** of (Φ, M) is an independent subset $\Omega \subset \Phi$ such that $\Omega M = \Phi$ (Ω generates (Φ, M)).

Two measurements ψ and ϕ are called **indistinguishable** if ψ is a deformation of ϕ and ϕ is a deformation of ψ. If (Φ, M) is a group perception pair, then ψ and ϕ are indistinguishable if and only if $\psi = \phi g$ for some g in M, i.e., if ψ is a deformation of ϕ.

4.1. **Proposition.** (1) *Every perception pair has a basis.*
(2) *Let $\Omega, \Omega' \subset \Phi$ be two bases of an perception pair (Φ, M). Then there is a bijection $\sigma\colon \Omega \to \Omega'$ such that ω and $\sigma(\omega)$ are indistingishable for every ω in Ω.*

Proof. (1): Let (Φ, M) be a perception pair. Choose $\Omega \subset \Phi$ to be an independent subset for which ΩM is maximal. The existence of Ω is guaranteed by the finiteness of Φ. We claim that $\Omega M = \Phi$ and hence Ω is a basis. If this is not the case, consider ψ in $\Phi \setminus \Omega M$. Define $\Omega' = \{\psi\} \cup \{\omega \in \Omega \mid \omega \notin \{\psi\}M\}$. Then $\Omega' M$ contains Ω and hence ΩM. It also contains ψ. Since Ω' is independent, we get a contradiction to the maximality of ΩM, and thus the claim holds.

(2): Let ω be in Ω. Since $\Omega M = \Phi = \Omega' M$, there is ω' in Ω' such that $\omega \in \omega' M$. Let ω_1 in Ω be such that $\omega' \in \omega_1 M$. Then $\omega \in \omega' M \subset \omega_1 M$, and hence $\omega = \omega_1$ by the independence of Ω. The desired bijection is then given by $\omega \mapsto \omega'$. □

According to Proposition 4.1, any two bases of a perception pair have the same number of elements. We define the **dimension** of a perception pair to be the cardinality of its bases. For example, a transitive group perception pair has dimension 1. In fact for a transitive group perception pair any single measurement forms a basis. More generally, the dimension of a group perception pair (Φ, M) equals the cardinality of Φ/M. In this case $\Omega \subset \Phi$ is a basis if and only if, for every equivalence class $[\psi]$ in Φ/M, the intersection $\Omega \cap [\psi]$ has only one element. Since being a basis depends only on the monoid $\langle M\rangle$, the dimension of a group-like perception pair (Φ, M) equals also the cardinality of Φ/M, and similarly a subset $\Omega \subset \Phi$ is a basis if and only if, for every equivalence class $[\psi]$ in the partition Φ/M, the intersection $\Omega \cap [\psi]$ has only one element.

The dimension of a transitive monoid perception pair can be bigger than 1. For example, let $X = \{x_1, x_2, x_3\}$ and consider functions $\phi_1, \phi_2, \phi_3 \colon X \to \mathbf{R}$ and $g_1, g_2, g_3 \colon X \to X$ defined as follows:

$$\begin{array}{c|c|c|c|c|c}
\phi_1(x_1)=2 & \phi_2(x_1)=2 & \phi_3(x_1)=1 & g_1(x_1)=x_2 & g_2(x_1)=x_2 & g_3(x_1)=x_1\\
\phi_1(x_2)=2 & \phi_2(x_2)=2 & \phi_3(x_2)=2 & g_1(x_2)=x_2 & g_2(x_2)=x_2 & g_3(x_2)=x_2\\
\phi_1(x_3)=3 & \phi_2(x_3)=2 & \phi_3(x_3)=2 & g_1(x_3)=x_3 & g_2(x_3)=x_2 & g_3(x_3)=x_2
\end{array}$$

The compositions $g_i g_j$ and $\phi_i g_j$ are described by the following tables:

	g_1	g_2	g_3
g_1	g_1	g_2	g_2
g_2	g_2	g_2	g_2
g_3	g_2	g_2	g_3

	g_1	g_2	g_3
ϕ_1	ϕ_1	ϕ_2	ϕ_2
ϕ_2	ϕ_2	ϕ_2	ϕ_2
ϕ_3	ϕ_2	ϕ_2	ϕ_3

Thus the functions g_1, g_2, and g_3 are $\Phi := \{\phi_1, \phi_2, \phi_3\}$-operations. Furthermore the subset $M := \{\mathrm{id}, g_1, g_2, g_3\} \subset \mathrm{End}_\Phi(X)$ is a submonoid. The perception pair (Φ, M) is a transitive monoid perception pair. Since the set $\{\phi_1, \phi_3\}$ is independent and generates (Φ, M), it is a basis. Thus (Φ, M) is an example of a transitive monoid perception pair of dimension 2.

5. Perception space

To compare perception pairs of various data sets we are going to use SEOs (set equivariant operators). A **SEO** from a perception pair (Φ, M) to another (Ψ, N), denoted as $(\alpha, T) \colon (\Phi, M) \to (\Psi, N)$, is a pair of functions $(\alpha \colon \Phi \to \Psi, T \colon M \to N)$ for which the following diagram commutes:

$$\begin{array}{ccccc}
\Phi \times M & \hookrightarrow & \Phi \times \mathrm{End}_\Phi(X) & \xrightarrow{\text{action}} & \Phi\\
\downarrow{\scriptstyle \alpha\times T} & & & & \downarrow{\scriptstyle \alpha}\\
\Psi \times N & \hookrightarrow & \Psi \times \mathrm{End}_\Psi(Y) & \xrightarrow{\text{action}} & \Psi
\end{array}$$

Explicitly, for ϕ in Φ and g in M, it holds $\alpha(\phi g) = \alpha(\phi)T(g)$. This implies that, for ϕ in Φ and a sequence $g_1, \ldots, g_k$ in M,

$$\alpha(\phi g_1 \cdots g_k) = \alpha(\phi)T(g_1)\cdots T(g_k).$$

Be, however, aware that, in general, there may not be a homomorphism $T\colon \langle M\rangle \to \langle N\rangle$ of monoids which extends $T\colon M \to N$ and makes the following diagram commute

$$\begin{array}{ccccccc} \Phi\times M & \hookrightarrow & \Phi\times\langle M\rangle & \hookrightarrow & \Phi\times \mathrm{End}_\Phi(X) & \xrightarrow{\text{action}} & \Phi \\ \downarrow{\scriptstyle \alpha\times T} & & \downarrow{\scriptstyle \alpha\times T} & & & & \downarrow{\scriptstyle\alpha} \\ \Psi\times N & \hookrightarrow & \Psi\times\langle N\rangle & \hookrightarrow & \Psi\times \mathrm{End}_\Psi(Y) & \xrightarrow{\text{action}} & \Psi \end{array}$$

A SEO between monoid perception pairs $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ is called a **MEO** (monoid equivariant operators) if $T\colon M \to N$ is a monoid homomorphism. A MEO between group perception pairs is also called a **GEO** (group equivariant operators).

Let $(\alpha_0, T_0)\colon (\Phi_0, M_0) \to (\Phi_1, M_1)$ and $(\alpha_1, T_1)\colon (\Phi_1, M_1) \to (\Phi_2, M_2)$ be SEOs. Then the compositions $(\alpha_1\alpha_0, T_1T_0)$ form a SEO. Furthermore the pair $(\mathrm{id}_\Phi, \mathrm{id}_M)\colon (\Phi, M) \to (\Phi, M)$ is also a SEO. The composition of SEOs is an associative operation and defines a category structure on the collection of perception pairs with SEOs as morphisms. This category is called **perception space**.

A SEO $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ is an isomorphism if and only if both of the functions α and T are bijections. Isomorphisms preserve independence and being a basis:

5.1. **Proposition.** *If $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ is an isomorphism, then a subset $\Omega \subset \Phi$ is independent or a basis if and only if its image $\alpha(\Omega) \subset \Psi$ is independent or a basis.*

Proof. Assume α and T are bijections. This assumption imply that ϕ_1 belongs to $\phi_2 M$ if and only if $\alpha(\phi_1)$ belongs to $\alpha(\phi_2)N$. It follows that two elements in Φ are (in)dependent if and only if their images via α are (in)dependent in Ψ. By the same argument, $\Omega M = \Phi$ if and only $\alpha(\Omega)T(M) = \alpha(\Phi)$. □

According to Proposition 5.1 two isomorphic perception pairs have the same dimension.

The universal perception pairs $(\Phi, \mathrm{End}_\Phi(X))$ and $(\Phi, \mathrm{Aut}_\Phi(X))$ are special in the perception space. For any (Φ, M), the pair $(\mathrm{id}, i\colon M \hookrightarrow \mathrm{End}_\Phi(X))$ defines a SEO $(\Phi, M) \to (\Phi, \mathrm{End}_\Phi(X))$ called **canonical**. If (Φ, M) is a group perception pair, then the pair $(\mathrm{id}, i\colon M \hookrightarrow \mathrm{Aut}_\phi(X))$ defines a GEO $(\Phi, M) \to (\Phi, \mathrm{Aut}_\phi(X))$ also called canonical.

The rest of this section is devoted to present three ways of constructing SEOs.

Change of units. Choose a function $f\colon \mathbf{R} \to \mathbf{R}$. For any perception pair (Φ, M), consider the data set $f\Phi$ (see Section 2). If g is a Φ-operation, then it is also a $f\Phi$-operation. Thus there is an inclusion $\mathrm{End}_\Phi(X) \subset \mathrm{End}_{f\Phi}(X)$, which is an equality if f is invertible, therefore we have a perception pair $(f\Phi, M)$. If (Φ, M) is a monoid or a group perception pair, then so is $(f\Phi, M)$. The pair $(f-, \mathrm{id}_M)\colon (\Phi, M) \to (f\Phi, M)$ is a SEO called the change of units along f.

Assume now that f is invertible. If $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ is a SEO, then the pair of functions $\big((f-)\alpha(f^{-1}-), T\big)$ forms a SEO between $(f\Phi, M)$ and $(f\Psi, N)$. The assignment $C(f)\colon (\alpha, T) \mapsto ((f-)\alpha(f^{-1}-), T)$ is a functor from the perception space to itself which is also called change of units along f. It is an equivalence of categories. Indeed,

$$\begin{aligned}\mathrm{C}(f)\mathrm{C}(f^{-1})((\Phi,M)) &= \mathrm{C}(f)(f^{-1}\Phi,M) = (\Phi,M)\\ \mathrm{C}(f)\mathrm{C}(f^{-1})((\alpha,T)) &= \mathrm{C}(f)((f^{-1}-)\alpha(f-),T))\\ &= ((f-)(f^{-1}-)\alpha(f-)(f^{-1}-),T) = (\alpha,T),\end{aligned}$$

and the same holds for $\mathrm{C}(f^{-1})\mathrm{C}(f)$. The SEOs $(f-,\mathrm{id}_M)\colon(\Phi,M)\to(f\Phi,M)$, for all perception pairs (Φ,M), form a natural transformation between the identity functor on the perception space and the change of units along f functor.

Domain change. Let (Φ,M) and (Ψ,N) be perception pairs of data sets consisting of measurements on X and Y respectively. A SEO $(\alpha,T)\colon(\Phi,M)\to(\Psi,N)$ is called **geometric** if there is a function $f\colon Y\to X$, called a realization of (α,T), making the following diagram commute for every ϕ in Φ and g in M

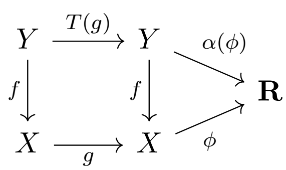

For example, let (Φ,M) be a perception pair of a data set consisting of measurements on X. Then the SEO $(\mathrm{id}_\Phi,\mathrm{id}_M)\colon(\Phi,M)\to(\Phi,M)$ is geometric. The identity function $\mathrm{id}_X\colon X\to X$ is one of its realizations.

Let $Y\subset X$ have the following property: $g(y)$ belongs to Y for all y in Y and g in M. Consider the data set $\Phi|_Y$ given by the domain change along the inclusion $Y\subset X$. The restriction of g to Y is a $\Phi|_Y$-operation for every g in M. We use the symbol $T_Y\colon M\to\mathrm{End}_{\Phi|_Y}(Y)$ to denote the function that maps g in M to the restriction of g to Y. The perception pair $(\Phi|_Y,T_Y(M))$ is called the **restriction** of (Φ,M) to the subset Y. The pair $(\Phi\twoheadrightarrow\Phi|_Y,T_Y)$ forms a geometric SEO. The inclusion $i_Y\colon Y\hookrightarrow X$ is one of its realizations.

Let $f\colon Y\to X$ be a bijection. Consider the data set Φf. For any g in M, the function $f^{-1}gf\colon Y\to Y$ is a Φf-operation. Define $T\colon M\to\mathrm{End}_{\Phi f}(Y)$ to map g in M to $f^{-1}gf$. The perception pair $(\Phi f,T(M))$ is called the domain change of (Φ,M) along f. The pair $(-f\colon\Phi\to\Phi f,T)$ forms a geometric SEO and $f\colon Y\to X$ is one of its realizations.

Extending from a basis. SEOs can be effectively constructed using bases.

5.2. **Proposition.** *Let (Φ,M) and (Ψ,N) be perception pairs and Ω be a basis of (Φ,M). Then two SEOs $(\alpha,T),(\alpha',T')\colon(\Phi,M)\to(\Psi,N)$ are equal if and only if $T=T'$ and $\alpha(\omega)=\alpha'(\omega)$ for any ω in Ω.*

Proof. The only non trivial thing to prove in the statement of the proposition is that $\alpha=\alpha'$ when their restrictions to Ω are equal. Assume $T=T'$ and $\alpha(\omega)=\alpha'(\omega)$ for any ω in Ω. Since Ω generates (Φ,M), any element in Φ is of the form $\phi=\omega g_1\cdots g_k$ for some ω in Ω and a sequence of elements $g_1,\ldots,g_k$ in M. The assumption and the fact that (α,T) and

(α', T) are SEOs imply:

$$\alpha(\phi) = \alpha(\omega g_1 \cdots g_k) = \alpha(\omega)T(g_1)\cdots T(g_k) =$$
$$= \alpha'(\omega)T(g_1)\cdots T(g_k) = \alpha'(\omega g_1 \cdots g_k) = \alpha'(\phi).$$

Consequently $\alpha = \alpha'$. □

According to Proposition 5.2, a SEO is determined by what it does on a basis of the domain. This is analogous to a linear map between vector spaces being determined by its values on a basis. However, unlike linear maps, we cannot freely map elements of a basis of a perception pair to obtain a SEO, but we need to respect certain additional relations. If (Φ, M) is a perception pair, a **relation** between measurements ϕ and ψ in Φ is a pair of sequences $((g_1, \dots, g_k), (h_1, \dots, h_l))$ of elements in M for which $\phi g_1 \cdots g_k = \psi h_1 \cdots h_l$.

5.3. **Proposition.** *Let (Φ, M) and (Ψ, N) be perception pairs, Ω be a basis of (Φ, M), and $\bar{\alpha}\colon \Omega \to \Psi$ and $T\colon M \to N$ be functions.*

1. *Assume that for every relation $((g_1, \dots, g_k), (h_1, \dots, h_l))$ between two elements ω, ω' in Ω, the pair $((T(g_1), \dots, T(g_k)), (T(h_1), \dots, T(h_l)))$ is a relation between $\alpha(\omega)$ and $\alpha(\omega')$ in Ψ. Under this assumption, there is a unique SEO, $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$, for which the restriction of $\alpha\colon \Phi \to \Psi$ to Ω is $\bar{\alpha}$.*
2. *Assume (Φ, M) and (Ψ, N), are monoid perception pairs, T is a monoid homomorphism, and if $\omega g = \omega' h$ for some ω, ω' in Ω and g, h in M, then $\alpha(\omega)T(g) = \alpha(\omega')T(h)$. Then there is a unique MEO, $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$, for which the restriction of $\alpha\colon \Phi \to \Psi$ to Ω is $\bar{\alpha}$.*
3. *Assume (Φ, M) and (Ψ, N) are group perception pairs, T is a group homomorphism, and if $\omega = \omega g$, for some ω in Ω and g in M, then $\alpha(\omega) = \alpha(\omega)T(g)$. Then there is a unique GEO, $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$, for which the restriction of $\alpha\colon \Phi \to \Psi$ to Ω is $\bar{\alpha}$.*

Proof. Since the proofs are analogous, we illustrate only how to show statement (2). For every ϕ in Φ, there exist (not necessarily unique) ω in Ω and g in M such that $\phi = \omega g$. The assumption implies that the expression $\alpha(\omega)T(g)$ depends on ϕ and not on the choices of ω and g for which $\phi = \omega g$. Thus by mapping ϕ in Φ to $\alpha(\omega)T(g)$ in Ψ, we obtain a well defined function also denoted by $\alpha\colon \Phi \to \Psi$. The pair (α, T) is the desired MEO. The uniqueness is a consequence of Proposition 5.2. □

An example of Proposition 5.3.3 is the following. Assume (Φ, M) is a transitive group perception pair and (Ψ, N) is a group perception pair. Choose an element ω in Φ and recall that any such element is a basis of (Φ, M). If $T\colon M \to N$ is a group homomorphism, then any GEO $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ is uniquely determined by the element $\alpha(\omega)$ in Ψ. Thus by choosing a basis element ω in Φ, we can identify the collection of GEOs of the form $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ with a subset of Ψ. To describe this subset explicitly, we apply Proposition 5.3.3. It states that there is a GEO $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ (necessarily unique) such that $\alpha(\omega) = \psi$ if and only if the following implication holds: if $\omega = \omega g$, then $\psi = \psi T(g)$. Since the collection $M_\omega := \{g \in M \mid \omega = \omega g\}$ is the isotropy subgroup of ω consisting of all the elements in M that fix ω, GEOs of the form $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$

can be identified with the subset of all the elements in Ψ whose isotropy group contains $T(M_\omega)$.

6. Decomposition

Let (Φ, M) be a perception pair of a data set Φ. Consider its quotient Φ/M, which is a partition of Φ, and the perception pairs $([\psi]], M)$ for every equivalence class $[\psi]$ in Φ/M (see Section 4). Let X be the domain of Φ. Recall that the domain of the data set $\coprod_{[\psi]\in\Phi/M}[\psi]$ is given by the disjoint union $\coprod_{[\psi]\in\Phi/M} X$, and that this data set consists of functions $\coprod_{[\psi]\in\Phi/M} X \to \mathbf{R}$ whose restrictions to all but one summands X in $\coprod_{[\psi]\in\Phi/M} X$ is the 0 function and the restriction to the remaining summand belongs to the corresponding equivalence class of the partition Φ/M. Define:

$$M' = \left\{ \coprod_{[\psi]\in\Phi/M} g\colon \coprod_{[\psi]\in\Phi/M} X \to \coprod_{[\psi]\in\Phi/M} X \quad | \quad g \in M \right\}.$$

Then $M' \subset \mathrm{End}_{\coprod_{[\psi]\in\Phi/M}[\psi]}(\coprod_{[\psi]\in\Phi/M} X)$. We call $(\coprod_{[\psi]\in\Phi/M}[\psi], M')$ the diagonal perception pair. Define $T\colon M \to M'$ to map $g\colon X \to X$ in M to $\coprod_{[\psi]\in\Phi/M} g$ in M'. Define $\alpha\colon \Phi \to \coprod_{[\psi]\in\Phi/M}[\psi]$ to map ϕ to the function $\coprod_{[\psi]\in\Phi/M} X \to \mathbf{R}$ whose restriction to the summand X corresponding to the equivalence class $[\phi]$ is ϕ and that maps all other summands to 0. Note that both of the functions α and T are bijections. Furthermore they form a SEO between (Φ, M) and $(\coprod_{[\psi]\in\Phi/M}[\psi], M')$.

6.1. **Proposition.** *The SEO* $(\alpha, T)\colon (\Phi, M) \to (\coprod_{[\psi]\in\Phi/M}[\psi], M')$ *is an isomorphism.*

7. Grothendieck graphs

In this section we explain a convenient data structure to encode perception pairs of data sets.

A **Grothendieck graph** is a triple (V, M, E) consisting of a finite set V whose elements are called vertices, a finite set M whose elements are called colors or operations, and a subset $E \subset V \times M \times V$ whose elements are called edges, such that, for every vertex v in V, the following composition is a bijection:

$$(\{v\} \times M \times V) \cap E \hookrightarrow E \hookrightarrow V \times M \times V \xrightarrow{\mathrm{pr}_M} M.$$

This condition assures that, for every v in V and g in M, there is a unique element in V, denoted by vg, such that (v, g, vg) is an edge in E. For example, let (Φ, M) be a perception pair of a data set Φ. Define

$$E_{\Phi,M} := \{(\phi, g, \psi) \in \Phi \times M \times \Phi \mid \phi g = \psi\}.$$

Then the triple $(\Phi, M, E_{\Phi,M})$ is a Grothendieck graph. We think about this graph as a convenient data structure representing the perception pair (Φ, M).

Grothendieck graphs are also convenient to represent SEOs. Define a **morphism between Grothendieck graphs** (V, M, E) and (W, N, F) to be a pair of functions $\alpha\colon V \to$

W and $T\colon M \to N$ such that, if (v, g, w) belongs to E, then $(\alpha(v), T(g), \alpha(w))$ belongs to F. Such a morphism is denoted as $(\alpha, T)\colon (V, M, E) \to (W, N, F)$. The componentwise composition defines a category structure on the collection of Grothendieck graphs and we use the symbol GGraph to denote this category. If $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$ is a SEO, then $(\alpha, T)\colon (\Phi, M, E_{\Phi,M}) \to (\Psi, N, E_{\Psi,N})$ is a morphism between the associated Grothendieck graphs. By assigning to a SEO (α, T) the graph morphism given by the same pair (α, T), we obtain a fully faithful functor from the perception space to GGraph.

Grothendieck graphs can also be used to encode pseudometric information on perception pairs. A pseudometric on a Grothendieck graph (V, M, E) is a pseudometric d on V such that $d(v, w) \geq d(vg, wg)$ for all v and w in V, and g in M. For example, the pseudometric $\|\phi - \psi\|_\infty$ on Φ is a pseudometric on the graph $(\Phi, M, E_{\Phi,M})$.

A Grothendieck graph (V, M, E) is said to be compatible with a monoid structure on M if $(v, 1, v)$ is in E, and whenever (v_0, g_0, v_1) and (v_1, g_1, v_2) belong to E, then so does $(v_0, g_1 g_0, v_2)$. In this case the composition operation given by the association $(v_0, g_0, v_1)(v_1, g_1, v_2) \mapsto (v_0, g_1 g_0, v_2)$ defines a category structure, denoted by $\mathrm{Gr}_M V$, with V as the set of objects and E as the set of morphisms. This category is a familiar Grothendieck construction [5, 12]. For example, the Grothendieck graph associated with a monoid perception pair (Φ, M) is compatible with the monoid structure on M. We think about $\mathrm{Gr}_M \Phi$ as an additional categorical structure on the data set Φ, where objects are the measurements in Φ, morphisms are triples $(\phi, g, \phi g)$, where ϕ is in Φ, g is in M, and the composition of $(\phi, g, \phi g)$ and $(\phi g, h, \phi gh)$ is given by $(\phi, gh, \phi gh)$.

A **contravariant functor** indexed by a Grothendieck graph (V, M, E) with values in a category $\mathcal{C}$, denoted by $P\colon (V, M, E) \to \mathcal{C}$, is by definition a sequence of objects $\{P(v) \mid v \in V\}$ and a sequence of morphisms $\{P(v_0, g, v_1)\colon P(v_1) \to P(v_0) \mid (v_0, g, v_1) \in E\}$ in $\mathcal{C}$ such that the following holds: if (v_0, g_0, v_1), (v_1, g_1, v_2), and (v_0, h, v_2) are edges in E, then $P(v_2, h, v_0) = P(v_2, g_1, v_1) P(v_1, g_0, v_0)$. If (V, M, E) is compatible with a monoid structure on M, then a contravariant functor indexed by (V, M, E) is simply a contravariant functor indexed by the category $\mathrm{Gr}_M V$.

Let (Φ, M) be a perception pair of a data set Φ consisting of measurements on X, and let $(\Phi, M, E_{\Phi,M})$ be the associated Grothendieck graph. For every g in M, the function $-g\colon \Phi \to \Phi$, mapping ϕ to ϕg, is geometric and realized by $g\colon X \to X$ (see Section 3). Persistent homology leads therefore to the following collections of objects and morphisms in $\mathrm{Tame}([0, \infty) \times \mathbf{R}, \mathrm{Vect})$ as explained in Section 3:

$$\left\{\mathrm{PH}_d^{\Phi}(\phi) \mid \phi \in \Phi\right\},$$

$$\left\{\mathrm{PH}_d^{-g}(\phi)\colon \mathrm{PH}_d^{\Phi}(\phi g) \to \mathrm{PH}_d^{\Phi}(\phi) \mid (\phi, g, \phi g) \in E_{\Phi,M}\right\}.$$

These sequences form a functor $PH_d^{\Phi}\colon (\Phi, M, E_{\Phi,M}) \to \mathrm{Tame}([0, \infty) \times \mathbf{R}, \mathrm{Vect})$ also referred to as the persistent homology functor of the perception pair (Φ, M).

Let $(\alpha, T)\colon (W, N, F) \to (V, M, E)$ be a morphism and $P\colon (V, M, E) \to \mathcal{C}$ be a functor. The following sequences of objects and morphisms in $\mathcal{C}$ form a contravariant functor denoted by $P(\alpha, T)\colon (W, N, F) \to \mathcal{C}$ and called the **composition** of (α, T) with P:

$$\{P(\alpha(w)) \mid w \in W\},$$

$$\{P(\alpha(w_0), T(g), \alpha(w_1))\colon P(\alpha(w_1)) \to P(\alpha(w_0)) \mid (w_0, g, w_1) \in F\}\,.$$

For example, let $(\mathrm{id}_\Phi, i)\colon (\Phi, M) \to (\Phi, \mathrm{End}_\Phi(X))$ be the canonical SEO (see Section 5). Consider the induced morphism of the associated Grothendieck graphs:

$$(\mathrm{id}_\Phi, i_M)\colon (\Phi, M, E_{\Phi,M}) \to (\Phi, \mathrm{End}_\Phi(X), E_{\Phi,\mathrm{End}_\Phi(X)}).$$

Consider also the persistent homology of the universal perception pair:

$$PH_d^\Phi\colon (\Phi, \mathrm{End}_\Phi(X), E_{\Phi,\mathrm{End}_\Phi(X)}) \to \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect}).$$

The composition of these two functors coincides with the persistent homology of the perception pair (Φ, M):

$$PH_d^\Phi\colon (\Phi, M, E_{\Phi,M}) \to \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect}).$$

In this way we obtain a commutative diagram:

$$\begin{array}{ccccc} & & (\Phi, \mathrm{End}_\Phi(X), E_{\Phi,\mathrm{End}_\Phi(X)}) & & \\ & \overset{(\mathrm{id}_\Phi, i_M)}{\nearrow} & & \overset{PH_d^\Phi}{\searrow} & \\ (\Phi, M, E_{\Phi,M}) & & \xrightarrow{\qquad PH_d^\Phi \qquad} & & \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect}) \end{array}$$

There is no such commutative diagram for arbitrary SEOs.

Consider a SEO $(\alpha, T)\colon (\Phi, M) \to (\Psi, N)$. We can form two functors indexed by the graph $(\Phi, M, E_{\Phi,M})$:

$$\begin{array}{ccccc} & & \xrightarrow{\qquad \mathrm{PH}_d^\Phi \qquad} & & \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect}) \\ (\Phi, M, E_{\Phi,M}) & & & & \\ & \xrightarrow{\alpha} & (\Psi, N, E_{\Psi,N}) & \xrightarrow{\mathrm{PH}_d^\Psi} & \mathrm{Tame}([0,\infty) \times \mathbf{R}, \mathrm{Vect}) \end{array}$$

These functors rarely coincide. However, in the case (α, T) is a geometric SEO, the morphisms $\mathrm{PH}_d^\alpha(\phi)\colon \mathrm{PH}_d^\Psi(\alpha(\phi)) \to \mathrm{PH}_d^\Phi(\phi)$ (see Section 3), for all ϕ in Φ, form a natural transformation.

8. Conclusions

In Figure1 we give a graphical representation of some of the concepts introduced in this article. Data sets can be equipped with three structures: a pseudometric, a perception pair describing an action, and a Grothendieck graph. We imagine the perception space as the collection of all possible perception pairs of data sets, represented by the shaded region in the figure. Each point in the perception space has a lot of internal structure allowing the extraction of persistent homology. In this landscape the black arrows represent geometric SEOs and the grey ones non-geometric SEOs. The main idea is that geometric SEOs enable us to compare relevant persistent homologies, whereas non-geometric SEOs contain complementary information to persistence.

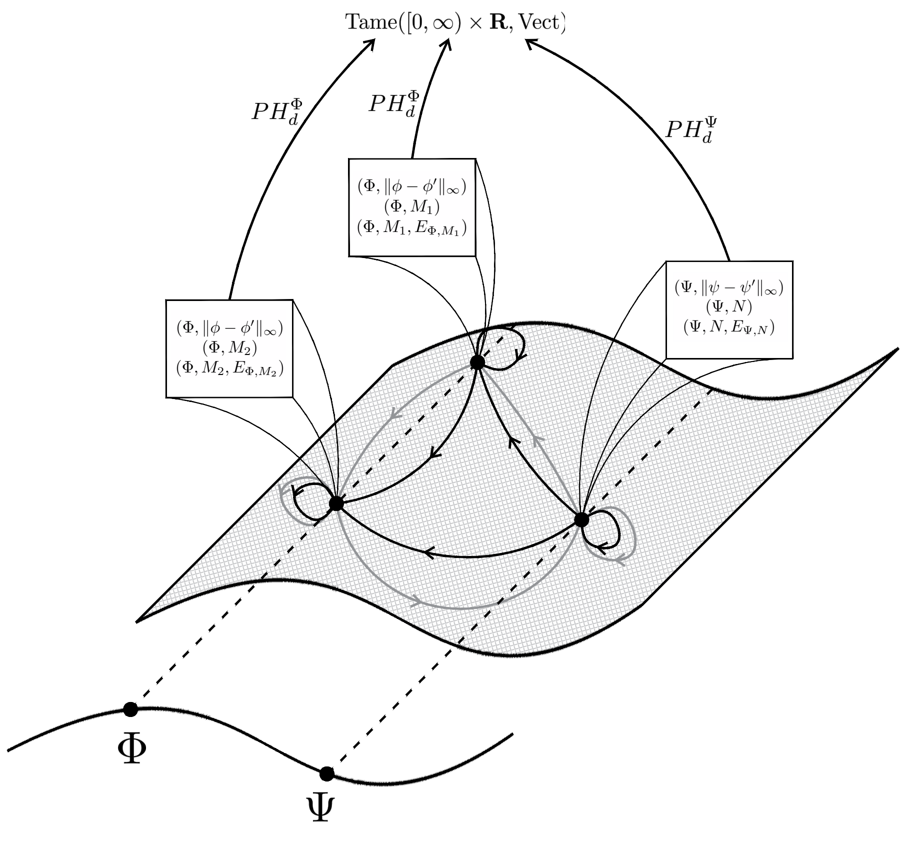

Figure 1

References

[1] Mattia G. Bergomi, Patrizio Frosini, Daniela Giorgi, and Nicola Quercioli. Towards a topological-geometrical theory of group equivariant non-expansive operators for data analysis and machine learning. *CoRR*, abs/1812.11832, 2018.

[2] Mattia G. Bergomi, Patrizio Frosini, Daniela Giorgi, and Nicola Quercioli. Towards a topological–geometrical theory of group equivariant non-expansive operators for data analysis and machine learning. *Nature Machine Intelligence*, 1(9):423–433, 2019.

[3] Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete Comput. Geom.*, 42(1):71–93, 2009.

[4] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.*, 9(1):79–103, 2009.

[5] William G. Dwyer and Hans-Werner Henn. *Homotopy theoretic methods in group cohomology*. Advanced Courses in Mathematics. CRM Barcelona. Birkhäuser Verlag, Basel, 2001.

[6] Herbert Edelsbrunner and John Harer. Persistent homology—a survey. In *Surveys on discrete and computational geometry*, volume 453 of *Contemp. Math.*, pages 257–282. Amer. Math. Soc., Providence, RI, 2008.

[7] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. 112(44):13455–13460, 2015.

[8] Alessandro De Gregorio, Ulderico Fugacci, Facundo Memoli, and Francesco Vaccarino. On the notion of weak isometry for finite metric spaces, 2020.

[9] Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Found. Comput. Math.*, 15(3):613–650, 2015.

[10] Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.

[11] Martina Scolamiero, Wojciech Chachólski, Anders Lundman, Ryan Ramanujam, and Sebastian Öberg. Multidimensional persistence and noise. *Foundations of Computational Mathematics*, 17(6):1367–1406, 2017.

[12] R. W. Thomason. Homotopy colimits in the category of small categories. *Math. Proc. Cambridge Philos. Soc.*, 85(1):91–109, 1979.