Some new results in 2D persistent homology

Patrizio Frosini

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

(based on a joint work with Andrea Cerri and Marc Ethier)

International Workshop and Conference on Topology & Applications, Rajagiri School of Engineering & Technology, Kochi, Kerala, INDIA, December 5-11, 2018

Outline

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

Let us recall what a persistence diagram is

Let $\varphi : X \to \mathbb{R}$ be a continuous function. Let us consider the sublevel sets $X_t := \{x \in X : \varphi(x) \leq t\}$ for t varying in $\mathbb{R}$, where the parameter t is seen as the time.

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

The points of the persistence diagram are endowed with multiplicity.

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

We can replace the points on the diagonal with just one point Δ, counted with infinite multiplicity.

Let us recall what a persistence diagram is

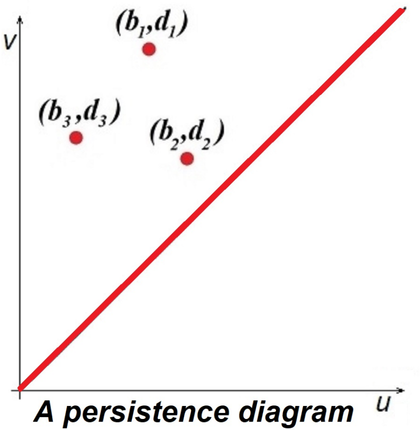

A persistence diagram

What can we do for vector-valued functions?

In topological data analysis we often have to consider vector-valued functions instead of scalar functions. As an example, we could have to compare two closed curves in $\mathbb{R}^2$, i.e. two functions $\varphi_1, \varphi_2 : S^1 \to \mathbb{R}^2$. How can we extend the concept of persistence diagram to this setting?

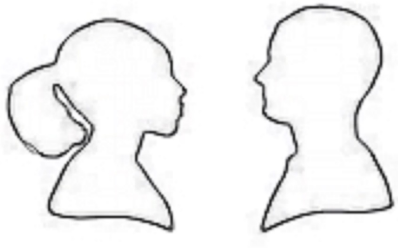

The foliation method

A (WRONG) IDEA: Let us study the two components f_1, f_2 separately.

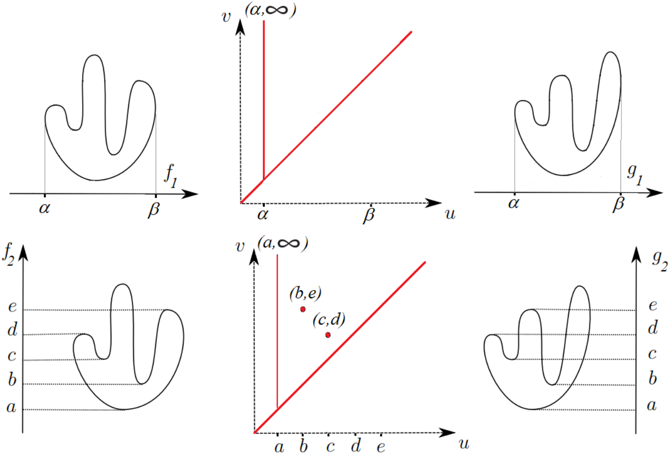

The foliation method

A BETTER IDEA:
Let us consider every filtration associated with a positive slope line.

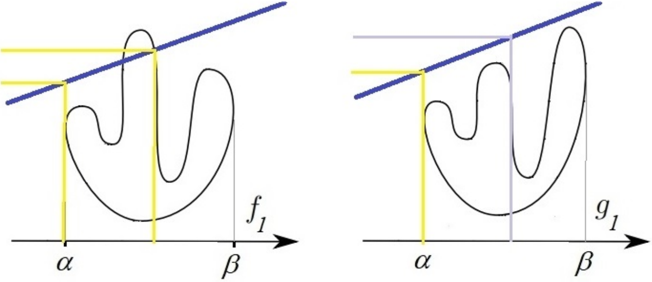

The chosen line allows to distinguish between these curves.

Bifiltrations and positive slope lines

If we have a bifiltration given by a function $f = (f_1, f_2) : X \to \mathbb{R}^2$, we can consider a unit vector (w.r.t. $\|\cdot\|_1$) $w = (a, 1-a)$ with a positive slope, and a point $P = (b, -b)$. Every choice of P and w defines a filtration $\{X_t\}$ of X, where X_t is the set of points of X whose image by f is both under and on the left of the point $P + tw$. As a consequence, each choice of P and w defines a persistence diagram.

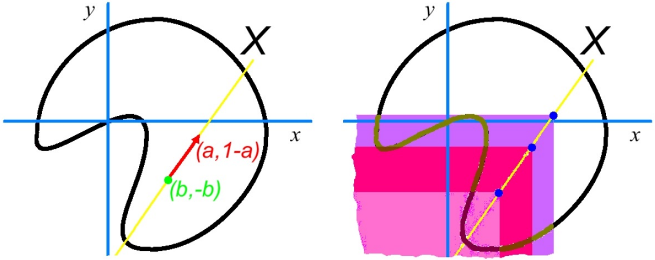

Bifiltrations and positive slope lines

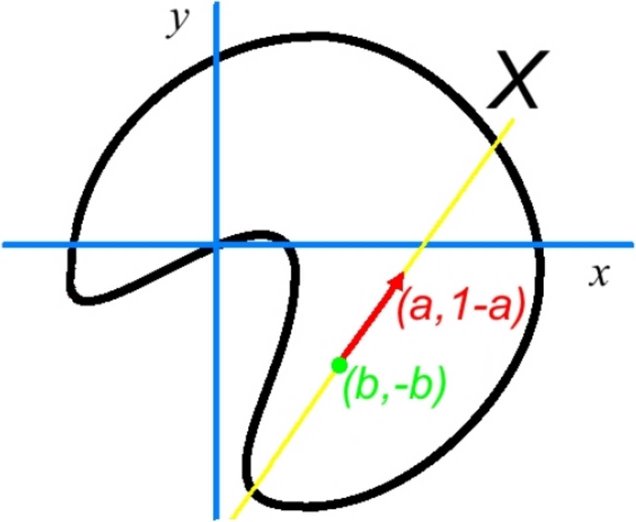

Bifiltrations and positive slope lines

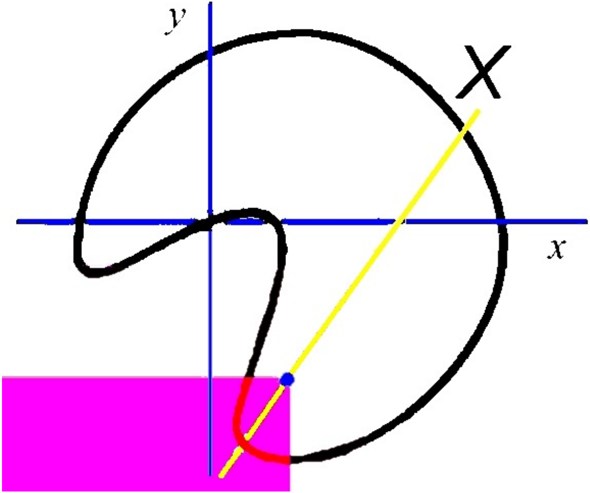

Bifiltrations and positive slope lines

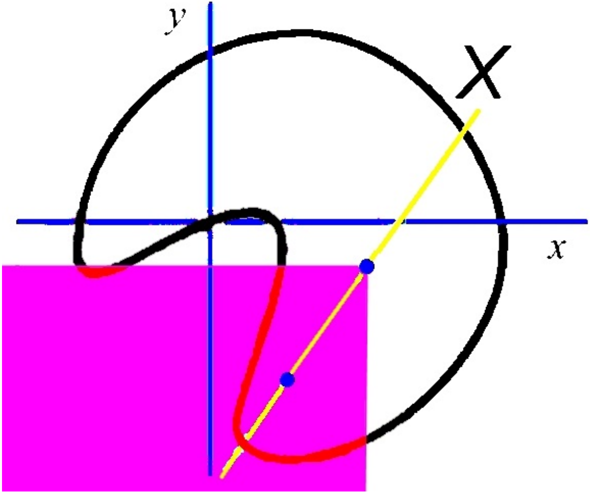

Bifiltrations and positive slope lines

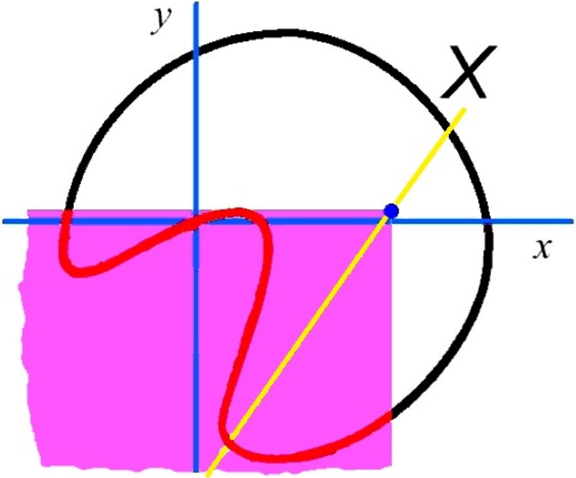

The normalized function $f^*_{(a,b)}$

If we set $(x,y) = P + tw = (at+b, (1-a)t-b)$ and define the function $f_{(a,b)}(p) := \max\left\{\frac{f_1(p)-b}{a}, \frac{f_2(p)+b}{1-a}\right\}$, we can write $X_t = \{p \in X : f_1(p) \leq x, f_2(p) \leq y\}$ as the set $\{p \in X : f_{(a,b)}(p) \leq t\}$.
As a consequence, the filtration $\{X_t\}$ of X leads us to consider the persistence diagram $\mathrm{Dgm}(f_{(a,b)})$ of the function $f_{(a,b)}$.
In order to get a stability theorem we have to normalize $f_{(a,b)}$ by setting

$$f^*_{(a,b)}(p) := \min\{a, 1-a\} \cdot f_{(a,b)}(p).$$

The persistence diagram $\mathrm{Dgm}(f^*_{(a,b)})$ can be obtained by multiplying the persistence diagram $\mathrm{Dgm}(f_{(a,b)})$ by the value $\min\{a, 1-a\}$.

Stability of D_{match}

We can define a 2D matching distance $D_{\mathrm{match}}(f,g)$ by setting
$D_{\mathrm{match}}(f,g) := \sup_{(a,b)\in]0,1[\times\mathbb{R}} d_{\mathrm{match}}\left(\mathrm{Dgm}(f^*_{(a,b)}),\mathrm{Dgm}(g^*_{(a,b)})\right)$.

Theorem (Stability Theorem)

$D_{\mathrm{match}}(f,g) \leq \|f-g\|_\infty$.

Remark: The previous theorem strongly depends on the normalization of persistence diagrams.

The distance D_{match} has been introduced in the paper *[S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, C. Landi, Multidimensional size functions for shape comparison, Journal of Mathematical Imaging and Vision, vol. 32 (2008), n. 2, 161-179.]*

Computation of the 2D matching distance

An algorithm to compute the 2D matching distance exists:

- *S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, A new algorithm for computing the 2-dimensional matching distance between size functions, Pattern Recognition Letters, vol. 32 (2011), n. 14, 1735-1746.*
- *A. Cerri, P. Frosini, A new approximation algorithm for the matching distance in multidimensional persistence, AMS Acta, 2971 (2011).*

Computation of the 2D matching distance

Remarks:

- In principle, 2D persistent homologies can be also compared by using the interleaving distance (*[M. Lesnick, The theory of the interleaving distance on multidimensional persistence modules, Foundations of Computational Mathematics, vol. 15 (2015), n. 3, 613-650]*). Unfortunately, the interleaving distance on multidimensional persistence modules is not easy to compute (NP-hard, *[H. B. Bjerkevik and M. B. Botnan, Computational Complexity of the Interleaving Distance, arXiv:1712.04281]*). This fact justifies the interest in the 2D matching distance D_{match}, which is more easily computable.
- As for the visualization of 2D persistence modules we point out the interesting paper *[M. Lesnick and M. Wright, Interactive Visualization of 2-D Persistence Modules, arXiv:1512.00180]*.

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

Some experiments in 3D shape comparison

Let us have a look at some pictures illustrating some results that we have obtained in 3D shape comparison by means of the 2D matching distance. The objects that we compare are displayed on the left of each figure. The color at the point (a, b) represents the value $d_{\mathrm{match}}\left(\mathrm{Dgm}(f^*_{(a,b)}), \mathrm{Dgm}(g^*_{(a,b)})\right)$. The largest values are in **red** and **brown**, the lowest ones are in **blue**. The values at points (a, b) with $|b|$ large are not displayed, since they just represent the 1-dimensional persistence of the two components of $f = (f_1, f_2)$ and are not relevant for our exposition. We recall that we are interested in $D_{\mathrm{match}}(f, g)$, i.e. the supremum of $d_{\mathrm{match}}\left(\mathrm{Dgm}(f^*_{(a,b)}), \mathrm{Dgm}(g^*_{(a,b)})\right)$, for $0 < a < 1$ and $b \in \mathbb{R}$.

Some experiments in 3D shape comparison

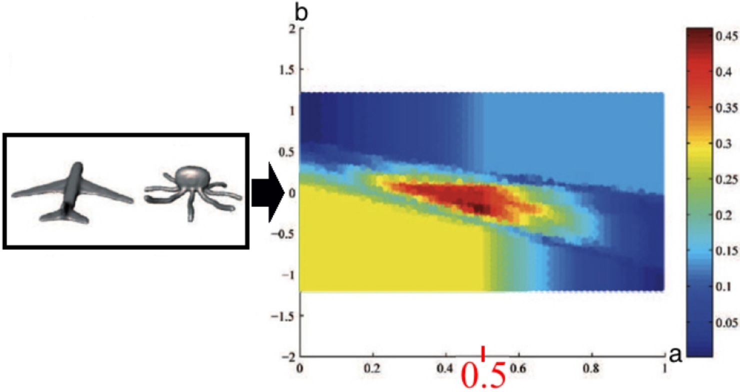

Please note that the largest value (i.e. the 2-dimensional matching distance) is taken at a point (a, b) with $a \approx 1/2$.

Some experiments in 3D shape comparison

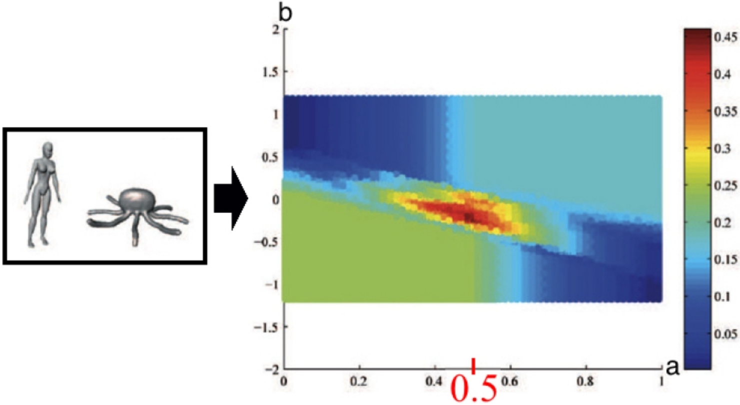

Please note that the largest value (i.e. the 2-dimensional matching distance) is taken at a point (a, b) with $a \approx 1/2$.

Some experiments in 3D shape comparison

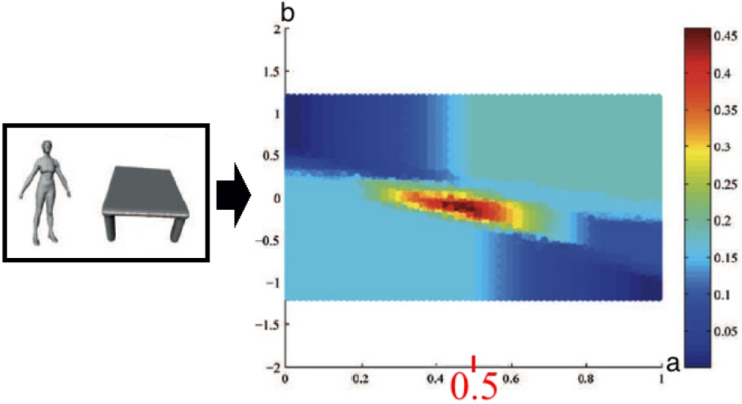

Please note that the largest value (i.e. the 2-dimensional matching distance) is taken at a point (a, b) with $a \approx 1/2$.

An interesting question

In our experiments, it seems that all the relevant information is given at points (a, b) with $a = 1/2$, i.e. in correspondence of lines with slope 1.

At the beginning we thought that this phenomenon was just a coincidence. So we looked for other examples, showing global maxima at points (a, b) with $a \neq 1/2$, but our search was unsuccessful.

A natural question arises: Is there any interesting principle hidden in the results of our experiments?

WHAT IS SPECIAL WITH SLOPE 1 LINES?

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

Some difficulties in our problem

Unfortunately, we are not able to directly answer this question, because of the lack of geometrical properties in the definition of D_{match}. Furthermore, we observe that while the metric D_{match} is rather simple to define and compute by considering a suitable family of filtering functions associated with lines having a positive slope, it has two main drawbacks.

- First, it forgets the natural link between the homological properties of filtrations associated with lines that are close to each other, so that part of the interesting homological information is lost.
- Second, its intrinsically discontinuous definition makes it difficult to study its properties.

A naïve (and wrong) solution to our problem

Our first (and naïve) idea was the one of defining a new distance, focusing on COHERENT matchings, i.e. matchings $\sigma_{(a,b)}$ that depend continuously on a and b. In order to build coherent matchings we tried to follow this approach:

- We should fix a pair $(\bar{a},\bar{b})$ and a matching $\sigma_{(\bar{a},\bar{b})}$ between the persistence diagrams $\mathrm{Dgm}(f^*_{(\bar{a},\bar{b})}),\mathrm{Dgm}(g^*_{(\bar{a},\bar{b})})$;
- Then we should construct a coherent family of matchings $\sigma_{(a,b)}$ between the persistence diagrams $\mathrm{Dgm}(f^*_{(a,b)}),\mathrm{Dgm}(g^*_{(a,b)})$, by "transporting $\sigma_{(\bar{a},\bar{b})}$ to any other point (a,b) in the parameter space by continuity";
- This transportation would use the stability of the normalized persistence diagrams $\mathrm{Dgm}(f^*_{(a,b)}),\mathrm{Dgm}(g^*_{(a,b)})$, i.e. the fact that they depend continuously on a and b.

A naïve (and wrong) solution to our problem

The previously proposed approach has a problem. If along the transportation we meet a point (a', b') at which one of the two normalized persistence diagrams contains a multiple point, then our transportation is not well-defined. Indeed, during "collisions" the identity of points is not preserved and we are not able to follow them.

The consequent natural idea is the one of following the movements of points of $\mathrm{Dgm}(f^*_{(a,b)})$ and $\mathrm{Dgm}(g^*_{(a,b)})$ avoiding the pairs (a', b') for which points with multiplicity greater than one exist. We call these pairs singular. The other pairs are called regular.

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

A key difficulty in our naïve solution

Unfortunately, the previous approach still has a problem.
Indeed, the definition of transport depends on the choice of the path γ that we follow in the parameter space. Precisely, this definition depends on the homotopy class of γ relative to the startpoint $(\bar{a}, \bar{b})$ and the endpoint (a, b).

We call this fact the monodromy phenomenon in 2-dimensional persistent homology.

For more details about the monodromy phenomenon in 2-dimensional persistent homology we refer to the paper
[A. Cerri, M. Ethier, P. Frosini, A study of monodromy in the computation of multidimensional persistence, Proceedings of the 17th IAPR International Conference on Discrete Geometry for Computer Imagery, LNCS 7749, 2013, 192-202].

An example of monodromy in 2D persistent homology

Let us see what happens if we consider the function $f = (f_1, f_2) : \mathbb{R}^2 \to \mathbb{R}^2$ obtained by setting

$$f_1(x,y) = x$$

$$f_2(x,y) := \begin{cases} -x, & \text{for } y = 0 \\ -x+1, & \text{for } y = 1 \\ -2x, & \text{for } y = 2 \\ -2x+\frac{5}{4}, & \text{for } y = 3 \end{cases}$$

$f_2(x,y)$ then being extended linearly for every x on the segment joining $(x,0)$ with $(x,1)$, $(x,1)$ with $(x,2)$, and $(x,2)$ to $(x,3)$. On the half-lines $\{(x,y) \in \mathbb{R}^2 : y < 0\}$ and $\{(x,y) \in \mathbb{R}^2 : y > 3\}$ f_2 is then being taken with constant slope 1 in the variable y.

An example of monodromy in 2D persistent homology

The graph of the function f_2:

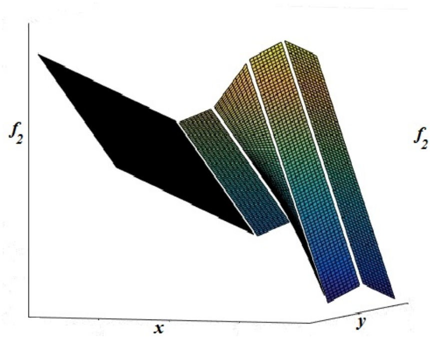

An example of monodromy in 2D persistent homology

We see that if we turn around the point $(1/4, 0)$ in the parameter space $]0,1[\times\mathbb{R}$, two points in the persistence diagram $\mathrm{Dgm}(f^*_{(a,b)})$ exchange their position. In other words, a loop around the singular point $(1/4, 0)$ induces a permutation on the persistence diagram.

Therefore, a monodromy group is associated with the function f. In order to properly define this group, we have to give a precise definition of the path followed by a point $p \in \mathrm{Dgm}(f^*_{(a,b)})$ when (a, b) moves. We will make this point precise in the next slides.

$a = 1/2$?

The existence of monodromy implies that different paths going from $(\bar{a}, \bar{b})$ to (a, b) can "transport matchings in different ways".

Does monodromy prevent us from proceeding in our research and from defining a coherent 2-dimensional matching distance?

Fortunately, the answer is no, as we will show in the next slides.

In the next slides we will also come back to our question

$$\boxed{a = 1/2?}$$

In order to face this problem we will have to introduce some mathematical machinery.

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

Some technical assumptions

Let $f = (f_1, f_2)$ be a smooth map from a closed C^∞-manifold M of dimension $r \geq 2$ to the real plane $\mathbb{R}^2$. Choose a Riemannian metric on M so that we can define gradients for f_1 and f_2.

The Jacobi set $\mathbb{J}(f)$ is the set of all points $p \in M$ at which the gradients of f_1 and f_2 are linearly dependent, namely $\nabla f_1(p) = \lambda \nabla f_2(p)$ or $\nabla f_2(p) = \lambda \nabla f_1(p)$ for some $\lambda \in \mathbb{R}$. In particular, if $\lambda \leq 0$ the point $p \in M$ is said to be a critical Pareto point for f. The set of all critical Pareto points of f is denoted by $\mathbb{J}_P(f)$.

Some technical assumptions

We will assume that $f : M \to \mathbb{R}^2$ is regular enough in a suitable sense (here we skip the technical details).

It follows that the Jacobi set is a smooth 1-submanifold of M, consisting of finitely many components, each one diffeomorphic to a circle.

Furthermore, the set of critical Pareto points at which the gradients of f_1 and f_2 are not orthogonal to the Jacobi set is made of a finite family $\{\alpha_i\}$ of arcs. Along these arcs, one of f_1 and f_2 is strictly increasing and the other is strictly decreasing. Each arc can meet critical points for f_1, f_2 only at its endpoints.

For more details: *[Y.H. Wan, Morse theory for two functions, Topology 14 (1975), no. 3, 217-228.]*

The Jacobi set

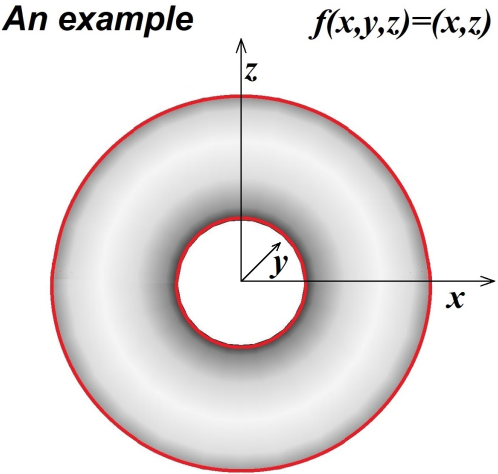

Critical Pareto points

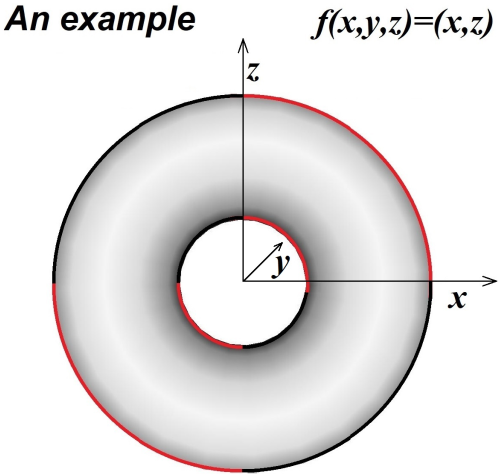

The extended Pareto grid

Our purpose is to establish a formal link between the position of points of $\mathrm{Dgm}(f^*_{(a,b)})$ for a function f and the intersections of the positive slope line $r_{(a,b)}$ with a particular subset of the plane $\mathbb{R}^2$, called the extended Pareto grid of f.

To define the extended Pareto grid, we start by listing the critical points $p_1, \ldots, p_h$ of f_1 and the critical points $q_1, \ldots, q_k$ of f_2 (we assume that $\{p_1, \ldots, p_h\} \cap \{q_1, \ldots, q_k\} = \emptyset$). Then we consider the following half-lines: for each critical point p_i of f_1 (resp. each critical point q_j of f_2), the half-line $\{(x,y) \in \mathbb{R}^2 | x = f_1(p_i), y \geq f_2(p_i)\}$ (resp. the half-line $\{(x,y) \in \mathbb{R}^2 | x \geq f_1(q_j), y = f_2(q_j)\}$).

The extended Pareto grid

The extended Pareto grid $\boldsymbol{\Gamma(f)}$ will be the union of $f(\mathbb{J}_P(f))$ with these half-lines. The closures of the images of the arcs α_i will be called proper contours of f, while the half-lines will be known as improper contours of f. We observe that every contour is a closed set.

Let $S(f)$ be the set of all points of $\Gamma(f)$ that belong to more than one (proper or improper) contour. This leads to endow the points of $\Gamma(f)$ with a concept of multiplicity. We assume that $S(f)$ is finite, and that the extended Pareto set is generic in a suitable technical sense (e.g., we assume no triple point exists).
Let $\mathscr{D}(f)$ be the set of double points in $\Gamma(f)$.
Each connected component of $\Gamma(f) \setminus \mathscr{D}(f)$ is called a contour-arc of f.

The extended Pareto grid: An example

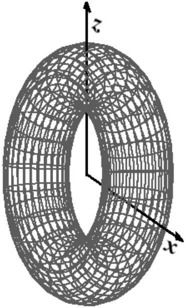

The torus endowed with the filtering function $f(p) := (x(p), z(p))$.

The extended Pareto grid: An example

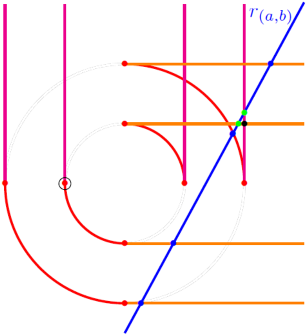

The extended Pareto grid for the torus endowed with the filtering function $f(p) := (x(p), z(p))$. A blue positive slope line $r_{(a,b)}$ is also represented.

The Position Theorem

With the concept of extended Pareto grid at hand, we can state and prove the following result, which gives a necessary condition for P to be a point of $\mathrm{Dgm}(f^*_{(a,b)})$.

We recall that

$$f^*_{(a,b)} := \max\left\{\frac{\min\{a,1-a\}}{a}\cdot(f_1-b), \frac{\min\{a,1-a\}}{1-a}\cdot(f_2+b)\right\}.$$

We set $\Delta := \{(u,v)\in\mathbb{R}^2 : u = v\}$.

Theorem (Position Theorem)

Let $(a,b)\in]0,1[\times\mathbb{R}$, $P\in\mathrm{Dgm}(f^*_{(a,b)})\setminus\Delta$. *Then, for each finite coordinate* c *of* P *a point* $(x,y)\in r_{(a,b)}\cap\Gamma(f)$ *exists, such that* $c = \frac{\min\{a,1-a\}}{a}\cdot(x-b) = \frac{\min\{a,1-a\}}{1-a}\cdot(y+b)$.

Using the extended Pareto grid

The Position Theorem suggests a way to find the possible positions for points of $\mathrm{Dgm}(f^*_{(a,b)})$. It consists in drawing the extended Pareto grid $\Gamma(f)$ and considering its intersections $(x_1, y_1), \ldots, (x_l, y_l)$ with the positive slope line $r_{(a,b)}$. For each proper point of $\mathrm{Dgm}(f^*_{(a,b)})$, both its coordinates belong to the set

$$\left\{ \frac{\min\{a, 1-a\}}{a} \cdot (x_i - b) = \frac{\min\{a, 1-a\}}{1-a} \cdot (y_i + b) \right\}_{1 \le i \le l} \cup \{\infty\}.$$

In other words, the Position Theorem allows us to follow the movements of the points in the persistence diagram $\mathrm{Dgm}(f^*_{(a,b)})$ by following the intersection points of the line $r_{(a,b)}$ with the extended Pareto grid of f.

Using the extended Pareto grid

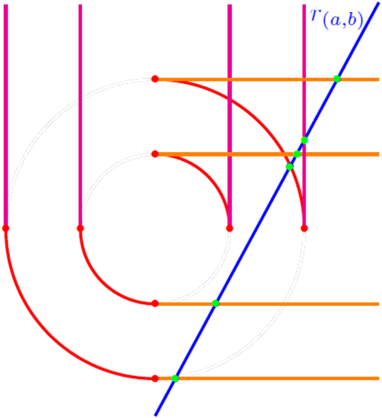

Each coordinate of a point in $\mathrm{Dgm}(f^*_{(a,b)})$ equals $\frac{\min\{a,1-a\}}{a} \cdot (x-b)$, where (x,y) is a green point.

Using the extended Pareto grid

Note that when $b < 0$ and $|b|$ is sufficiently large, the positive slope line $r_{(a,b)}$ may intersect $\Gamma(f)$ only at the vertical half-lines. In this case, $f^*_{(a,b)} := \frac{\min\{a,1-a\}}{a} \cdot (f_1 - b)$, and the values $x_1, \ldots, x_l$ are the critical values of f_1. Similarly, when $b > 0$ and $|b|$ is large enough, $r_{(a,b)}$ intersects $\Gamma(f)$ only at the horizontal half-lines. Then $f^*_{(a,b)} := \frac{\min\{a,1-a\}}{1-a} \cdot (f_2 + b)$, and the values $y_1, \ldots, y_l$ are the critical values of f_2. (See next slide)

Using the extended Pareto grid

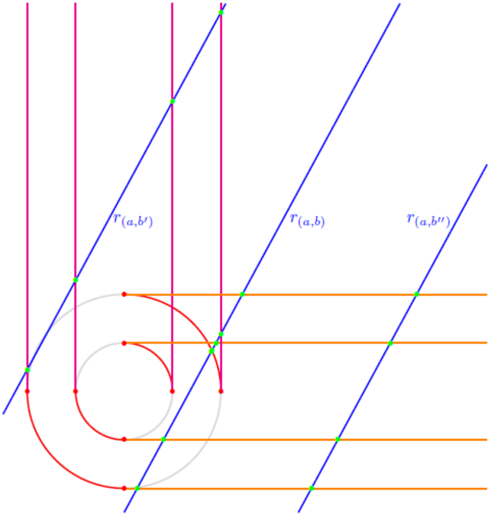

Using the extended Pareto grid

The Position Theorem also allows us to deduce where singular pairs can be in $]0,1[\times\mathbb{R}$.

Proposition

Let $(\bar{a},\bar{b})\in]0,1[\times\mathbb{R}$ *be a singular pair for* f*. If* $\mathrm{Dgm}\left(f^*_{(\bar{a},\bar{b})}\right)$ *contains a proper multiple point, then* $r_{(\bar{a},\bar{b})}$ *contains two points of* $\mathscr{D}(f)$*. If* $\mathrm{Dgm}\left(f^*_{(\bar{a},\bar{b})}\right)$ *contains an improper multiple point, then* $r_{(\bar{a},\bar{b})}$ *contains at least one point of* $\mathscr{D}(f)$*.*

Singular pairs

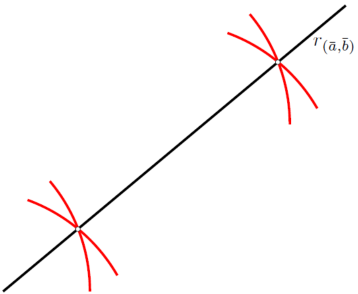

Figure: A line $r_{(\bar{a},\bar{b})}$ associated with a singular pair $(\bar{a},\bar{b}) \in]0,1[\times\mathbb{R}$, in case $\mathrm{Dgm}\left(f^*_{(\bar{a},\bar{b})}\right)$ contains a proper double point. Parts of four contours (split in eight proper contour-arcs) are displayed in red.

Creation and distruction of points in $\mathrm{Dgm}(f^*_{(a,b)})$ when (a,b) varies in $]0,1[\times\mathbb{R}$

The Position Theorem allows us to deduce at which points of Δ points of $\mathrm{Dgm}(f^*_{(a,b)})$ can be created or destroyed.

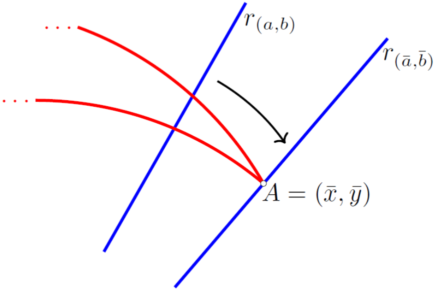

Figure: A point of $\mathrm{Dgm}(f^*_{(a,b)})$ reaches the diagonal Δ and disappears.

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

The coherent 2-dimensional matching distance

Our next step is the definition of the coherent 2-dimensional matching distance.

The existence of monodromy implies that each loop in the set $\mathrm{Reg}(f)$ of all regular pairs for f induces a permutation on $\mathrm{Dgm}(f^*_{(a,b)})$. In other words, it is not possible to establish which point in $\mathrm{Dgm}(f^*_{(a,b)})$ corresponds to which point in $\mathrm{Dgm}(f^*_{(a',b')})$ for $(a,b) \neq (a',b')$, since the answer depends on the path that is considered from (a,b) to (a',b') in the set $\mathrm{Reg}(f)$. As a consequence, different paths going from (a,b) to (a',b') might produce different results while "transporting" a matching $\sigma_{(a,b)} : \mathrm{Dgm}(f^*_{(a,b)}) \to \mathrm{Dgm}(g^*_{(a,b)})$ to another point $(a',b') \in]0,1[\times\mathbb{R}$.

Despite this problem, it is possible to define a notion of coherent 2-dimensional matching distance.

Choice of the functional set $\mathscr{F}_{U,c}$

In order to proceed we fix a connected open subset U of $]0,1[\times\mathbb{R}$, choose a $c>0$ and define $\mathscr{F}_{U,c}$ as the set of all normal functions $f: M\to\mathbb{R}^2$ such that $\mathrm{Reg}(f)\supseteq U$ and the distance $d(P,Q)$ between any two distinct points P,Q of $\mathrm{Dgm}\left(f^*_{(a,b)}\right)$ is strictly greater than $2c>0$ for every $(a,b)\in U$. Incidentally, we observe that for any $(a,b)\in U$ and any $f\in\mathscr{F}_{U,c}$, the sup-norm distance between the points in $\mathrm{Dgm}\left(f^*_{(a,b)}\right)\setminus\{\Delta\}$ and the diagonal Δ is strictly larger than $2c$. The definition of our coherent matching distance will depend on the choice of this set U. In the following we will assume that the functions f,g belong to the set $\mathscr{F}_{U,c}$. This assumption guarantees that the points in the persistence diagrams $\mathrm{Dgm}\left(f^*_{(a,b)}\right)$ are far enough from each other, and that the same holds for $\mathrm{Dgm}\left(g^*_{(a,b)}\right)$.

Transporting a matching along a path

First, we need to specify the concept of transporting a point $X \in \mathrm{Dgm}(f^*_{(a(0),b(0))})$ along a path $(a(t), b(t))$ in $\mathrm{Reg}(f)$.
We set $\Delta^* := \{(u,v) \in \mathbb{R} \times (\mathbb{R} \cup \{\infty\}) : u < v\}$.

Definition (Induced path)

A continuous path $P : [0,1] \to \Delta^* \cup \{\Delta\}$ is said to be *induced by the path* $\pi : [0,1] \to U$ if $P(\tau) \in \mathrm{Dgm}\left(f^*_{\pi(\tau)}\right)$ for every $\tau \in [0,1]$.

Proposition

Let $\pi : [0,1] \to U$ *be a continuous path. For every point* $X \in \mathrm{Dgm}\left(f^*_{\pi(0)}\right)$*, a unique continuous path* $P : [0,1] \to \Delta^* \cup \{\Delta\}$ *induced by* π *exists, such that* $P(0) = X$*. If* $X = \Delta$ *then* $P([0,1]) = \{\Delta\}$*, otherwise* $P([0,1]) \subseteq \Delta^*$*.*

Transport of points and matchings

With reference to the previous Proposition, we say that π transports X to $X' = P(1)$ with respect to f and write $T^f_\pi(X) = X'$. Now, we need to define the concept of transporting a matching along a path $\pi : [0,1] \to U$ with $\pi(0) = (a,b)$. Let $\sigma_{(a,b)}$ be a matching between $\mathrm{Dgm}(f^*_{(a,b)})$ and $\mathrm{Dgm}(g^*_{(a,b)})$, with $(a,b) \in U$. We can naturally associate to $\sigma_{(a,b)}$ a matching $\sigma_{\pi(1)} : \mathrm{Dgm}(f^*_{\pi(1)}) \to \mathrm{Dgm}(g^*_{\pi(1)})$. Suppose that $\sigma_{(a,b)}(X) = Y$. We set $\sigma_{\pi(1)}(X') = Y'$ if and only if π transports X to X' with respect to f and Y to Y' with respect to g. We also say that π transports $\sigma_{(a,b)}$ to $\sigma_{\pi(1)}$ along π with respect to the pair (f,g). The transported matching will be denoted by the symbol $T^{(f,g)}_\pi(\sigma_{(a,b)})$.

Transport of points and matchings

In other words, the transport of matchings is defined by the following commutative diagram:

$$\begin{tikzcd}
\mathrm{Dgm}\left(f^*_{\pi(0)}\right) \arrow[r, "\sigma_{\pi(0)}"] \arrow[d, "T^f_\pi"] & \mathrm{Dgm}\left(g^*_{\pi(0)}\right) \arrow[d, "T^g_\pi"] \\
\mathrm{Dgm}\left(f^*_{\pi(1)}\right) \arrow[r, "T^{(f,g)}_\pi\left(\sigma_{\pi(0)}\right) := T^g_\pi \circ \sigma_{\pi(0)} \circ \left(T^f_\pi\right)^{-1}"] & \mathrm{Dgm}\left(g^*_{\pi(1)}\right)
\end{tikzcd}$$

Transport of points

The next property trivially follows from the definition of transport.

Proposition

Let π_1, π_2 *be two continuous paths in* U*, with* $\pi_1(1) = \pi_2(0)$*. Let* $\pi_1 * \pi_2$ *be their composition, i.e. the loop* $\pi_1 * \pi_2 : [0,1] \to U$ *defined by setting* $\pi_1 * \pi_2(t) := \pi_1(2t)$ *for* $0 \leq t \leq 1/2$ *and* $\pi_1 * \pi_2(t) := \pi_2(2t-1)$ *for* $1/2 \leq t \leq 1$*. Then* $T^f_{\pi_2} \circ T^f_{\pi_1} = T^f_{\pi_1 * \pi_2}$*.*

Continuity of the transport w.r.t. the path

By using the 1-dimensional Stability Theorem and the Position Theorem, we can prove that the transport along a path in $\mathrm{Reg}(f)$ is continuous with respect to the path, as stated by the following proposition.

Proposition

Let $X \in \mathrm{Dgm}(f^*_{\pi(0)})$*. The function* $T^f_\pi(X)$ *is continuous in the variable* π*, when* π *varies in the set* $S^f_{(\bar{a},\bar{b})}$ *of the paths in* U *starting from a fixed point* $(\bar{a},\bar{b})$ *and* $S^f_{(\bar{a},\bar{b})}$ *is endowed with the uniform convergence metric.*

Each loop in U induces a permutation on $\mathrm{Dgm}(f^*_{(\bar{a},\bar{b})})$

From the previous proposition the next result immediately follows.

Proposition

If two paths π, π' *in* U *are homotopic to each other relatively to their common extrema, then* $T^f_\pi \equiv T^f_{\pi'}$*. The map* T^f *taking each equivalence class* $[\pi]$ *to the permutation* T^f_π *is a well-defined homomorphism from the fundamental group of* U *at* $(a,b) \in U$ *to the group of permutations of* $\mathrm{Dgm}\left(f^*_{(a,b)}\right)$*.*

Definition

The image of the group homomorphism T^f will be called the persistent monodromy group of the filtering function f with respect to U.

The coherent cost of a matching

Definition

Let $\Pi_{(a,b)}(U)$ be the set of all continuous paths $\pi : [0,1] \to U$ with $\pi(0) = (a,b)$. If $\sigma_{(a,b)} \in \Sigma_{(a,b)}^{(f,g)}$, the *coherent cost* of $\sigma_{(a,b)}$ is the value

$$\mathrm{cohcost}_U\left(\sigma_{(a,b)}\right) := \sup_{\pi \in \Pi_{(a,b)}(U)} \mathrm{cost}\left(T_{\pi}^{(f,g)}\left(\sigma_{(a,b)}\right)\right).$$

Proposition ($\mathrm{cohcost}_U$ is invariant under transport)

Let $\sigma_{(a,b)} : \mathrm{Dgm}\left(f^{*}_{(a,b)}\right) \to \mathrm{Dgm}\left(g^{*}_{(a,b)}\right)$ *be a matching, with* $(a,b) \in U$*. If* $\pi' : [0,1] \to U$ *is a continuous path with* $\pi'(0) = (a,b)$*, then* $\mathrm{cohcost}_U\left(T_{\pi'}^{(f,g)}\left(\sigma_{(a,b)}\right)\right) = \mathrm{cohcost}_U\left(\sigma_{(a,b)}\right)$.

The definition of the coherent matching distance

Let $\Sigma_{(a,b)}^{(f,g)}$ be the set of all matchings from $\mathrm{Dgm}\left(f^*_{(a,b)}\right)$ to $\mathrm{Dgm}\left(g^*_{(a,b)}\right)$, with $(a,b) \in U$.

Definition

Let $(a,b) \in U$. The coherent 2-dimensional matching distance between f and g is defined as

$$CD_U(f,g) = \min_{\sigma_{(a,b)} \in \Sigma_{(a,b)}^{(f,g)}} \mathrm{cohcost}_U\left(\sigma_{(a,b)}\right).$$

The definition of the coherent matching distance

Proposition

$CD_U(f,g)$ does not depend on the basepoint (a,b).

Proposition

$CD_U(f,g)$ is a pseudo-distance.

Stability of the coherent 2-dimensional matching distance

The next result shows that the coherent 2-dimensional matching distance is stable, in a suitable sense.

Theorem

If $f, g \in \mathscr{F}_{U,c}$ *and* $\|f-g\|_\infty < c$, *then* $CD_U(f,g) \leq \|f-g\|_\infty$.

A maximum principle for the coherent transport

Theorem (Maximum Principle)

Let $\sigma_{(\bar{a},\bar{b})} \in \Sigma^{(f,g)}_{(\bar{a},\bar{b})}$ *with* $\mathrm{cost}\left(\sigma_{(\bar{a},\bar{b})}\right) < \infty$. *If* $V \subseteq U$ *is an open neighborhood of* $(\bar{a},\bar{b})$ *such that*

$$\mathrm{cost}\left(\sigma_{(\bar{a},\bar{b})}\right) > \mathrm{cost}\left(T^{(f,g)}_{\pi}\left(\sigma_{(\bar{a},\bar{b})}\right)\right)$$

for all paths π *in* V *from* $\left(\bar{a},\bar{b}\right)$ *to* (a,b) *with* $a \neq \bar{a}$, *then* $\bar{a} = \frac{1}{2}$.

Topological bifiltrations

The importance of lines with slope 1 in 2D persistent homology

The need for coherent matchings

The phenomenon of monodromy

Extended Pareto Grid

The coherent 2-dimensional matching distance CD_U

The distance CD_U is achieved at $a = 1/2$

The distance CD_U is achieved at $a = 1/2$

Theorem

Let $(\bar{a}, \bar{b}) \in U$ *and* $f, g \in \mathscr{F}_{U,c}$. *Assume that the closure* $\overline{U}$ *of* U *in* $\mathbb{R}^2$ *is a compact set contained in the open set* $]0,1[\times\mathbb{R}$, *and that* ∂U *is a* C^1*-submanifold of* $\mathbb{R}^2$. *Then a matching* $\sigma_{(\bar{a},\bar{b})} \in \Sigma^{(f,g)}_{(\bar{a},\bar{b})}$, *a point* $(\hat{a}, \hat{b}) \in \overline{U}$ *and a continuous path* $\hat{\pi} : [0,1] \to \overline{U}$ *from* $(\bar{a}, \bar{b})$ *to* $(\hat{a}, \hat{b})$ *exist such that*

1. $\hat{\pi}([0,1[) \subseteq U$;
2. $\mathrm{cost}\left(T^{(f,g)}_{\hat{\pi}}\left(\sigma_{(\bar{a},\bar{b})}\right)\right) = \mathrm{cohcost}_U\left(\sigma_{(\bar{a},\bar{b})}\right) = CD_U(f,g)$;
3. $(\hat{a}, \hat{b}) \in \partial U$ *or* $\hat{a} = \frac{1}{2}$.

Further research

In our opinion, many problems should deserve further research:

- First of all, it would be interesting to extend the presented concepts to filtering functions taking values in $\mathbb{R}^m$ with $m > 2$.
- Secondly, the genericity of our assumptions concerning the extended Pareto grid should be possibly proved (joint research with Francesca Tombari).
- Finally, methods for the efficient computation of the coherent matching distance should be developed.

For more information about the approach described in these slides use the following link: http://arxiv.org/pdf/1801.06636v2.pdf.

Conclusions

In this talk we have presented a new line of research in 2D persistent homology, introducing the concept of coherent matching distance and studying some of its properties. In order to do that, we have also introduced the concept of extended Pareto grid and shown its use to manage the phenomenon of monodromy. Finally, we have proved a maximum principle for the coherent transport of matchings and a theorem that makes clear the importance of filtrations associated with lines of slope 1 in 2D persistent homology.


Thanks for your attention!